



## hmmTMB: Hidden Markov Models with Flexible Covariate Effects in R

Théo Michelot   
Dalhousie University

---

### Abstract

Hidden Markov models (HMMs) are widely applied in studies where a discrete-valued process of interest is observed indirectly. They have for example been used to model behavior from human and animal tracking data, disease status from medical data, and financial market volatility from stock prices. The model has two main sets of parameters: transition probabilities, which drive the latent state process, and observation parameters, which characterize the state-dependent distributions of observed variables. One particularly useful extension of HMMs is the inclusion of covariates on those parameters, to investigate the drivers of state transitions or to implement Markov-switching regression models. We present the new R package **hmmTMB** for HMM analyses, with flexible covariate models in both the hidden state and observation parameters. In particular, non-linear effects are implemented using penalized splines, including multiple univariate and multivariate splines, with automatic smoothness selection. The package allows for various random effect formulations (including random intercepts and slopes), to capture between-group heterogeneity. **hmmTMB** can be applied to multivariate observations, and it accommodates various types of response data, including continuous (bounded or not), discrete, and binary variables. Parameter constraints can be used to implement non-standard dependence structures, such as semi-Markov, higher-order Markov, and autoregressive models. Here, we summarize the relevant statistical methodology, we describe the structure of the package, and we present an example analysis of animal tracking data to showcase the workflow of the package.

*Keywords:* hidden Markov models, latent state, regime switching, non-parametric, penalized splines, non-homogeneous, mixed model, random effects, R.

---

## 1. Introduction

Hidden Markov models (HMMs) have been applied in many areas, including medicine (Altman and Petkau 2005), ecology (McClintock, Langrock, Gimenez, Cam, Borchers, Glennie, and

Patterson 2020), and finance (Bulla and Bulla 2006). They are useful when an observed phenomenon is driven by several unobserved regimes, or states, and when it is of interest to model the state-switching dynamics. We first describe the basic mathematical formulation of discrete-time HMMs, and we refer the reader to Zucchini, MacDonald, and Langrock (2017) for more detail.

An HMM consists of two stochastic processes: a state process  $(S_t)$ , which can take on a finite number of values  $\{1, 2, \dots, K\}$ , and an observation process  $(Z_t)$ . At each time step,  $Z_t$  is assumed to depend on the current state  $S_t$  through a state-dependent distribution,

$$p(Z_t | Z_1, \dots, Z_{t-1}, S_1, \dots, S_{t-1}, S_t) = p(Z_t | S_t).$$

The observation distribution (also called “emission distribution”) is often taken from some parametric family, where the parameters are state-dependent, i.e.,  $Z_t | \{S_t = j\} \sim \mathcal{D}(\omega_j)$  for  $j \in \{1, \dots, K\}$ . For example, if  $\mathcal{D}$  is the normal distribution,  $\omega_j$  might be a vector of the mean and standard deviation of  $Z_t$  in state  $j$ . In the case where there are several observed variables  $Z_t = (Z_{1t}, Z_{2t}, \dots, Z_{dt})$ , it is most common to assume that they are conditionally independent given the state, i.e.,  $p(Z_t | S_t) = p(Z_{1t} | S_t) \times p(Z_{2t} | S_t) \times \dots \times p(Z_{dt} | S_t)$ . This assumption of contemporaneous conditional independence is often reasonable because dependence between the variables is induced by the state process.

The state process is specified as a first-order Markov chain, i.e.,

$$p(S_t | S_1, \dots, S_{t-1}) = p(S_t | S_{t-1}). \quad (1)$$

It is parameterized in terms of an  $K \times K$  transition probability matrix

$$\Gamma = \begin{pmatrix} \gamma_{11} & \gamma_{12} & \cdots & \gamma_{1K} \\ \gamma_{21} & \gamma_{22} & \cdots & \gamma_{2K} \\ \vdots & \vdots & \ddots & \vdots \\ \gamma_{K1} & \gamma_{K2} & \cdots & \gamma_{KK} \end{pmatrix}$$

where  $\gamma_{ij} = P(S_t = j | S_{t-1} = i)$ , subject to the constraints that elements from each row sum to 1, i.e.,  $\sum_{k=1}^K \gamma_{ik} = 1$  for any  $i \in \{1, \dots, K\}$ . The specification of the Markov chain also requires a vector of  $K$  initial probabilities,  $\delta^{(1)} = (P(S_1 = 1), P(S_1 = 2), \dots, P(S_1 = K))$ , which must sum to 1.

Given observations  $z_1, z_2, \dots, z_n$  from the process  $(Z_t)$ , the problem of inference is typically to estimate all model parameters: the observation parameters  $\{\omega_1, \dots, \omega_K\}$ , the transition probability matrix  $\Gamma$ , and the initial distribution  $\delta^{(1)}$ . This is a difficult problem because the state process is unobserved, but a computationally efficient procedure exists to evaluate the likelihood of the data by marginalizing over all possible state sequences (the “forward algorithm”; Zucchini *et al.* 2017). The likelihood can usually be optimized numerically to obtain parameter estimates. This can be extended directly to the case of several independent time series: the full likelihood is the product of the likelihoods of the individual time series obtained from the forward algorithm.

In this paper, we introduce the new package **hmmTMB** for R (R Core Team 2025), which provides tools to implement HMMs where parameters of interest can be specified as flexible functions of covariates, such as non-linear effects modeled using splines, and random effects. The package implements a wide variety of model formulations, including mixed HMMs

(Altman 2007), non-homogeneous HMMs (Hughes, Guttorp, and Charles 1999), and Markov-switching regression and generalized additive models (Kim, Piger, and Startz 2008; Langrock, Kneib, Glennie, and Michelot 2017). We anticipate that it will be useful to many researchers in fields where HMMs are applied.

## 2. HMMs with non-parametric and random effects

The main focus of **hmmTMB** is to provide inferential tools for hidden Markov models that include flexible covariate dependencies on the transition probabilities or on the state-dependent observation distribution parameters. In this section, we give a brief overview of the modeling approach and of its implementation in **hmmTMB**.

### 2.1. Model formulation

#### *Linear predictor*

Consider a generic parameter  $\theta$ , which can be either a transition probability or an observation parameter, and the corresponding linear predictor  $\eta = h(\theta)$ , where  $h$  is a link function defined so that  $\eta$  is real-valued. (We discuss the choice of link function below.) To capture effects of covariates on  $\theta$ , we write the linear predictor using a mixed linear model approach,

$$\eta = X\alpha + R\beta, \quad \beta \sim N(0, Q), \quad (2)$$

where  $X$  and  $R$  are the design matrices for fixed and random effects, respectively,  $\alpha$  is the vector of fixed effects, and  $\beta$  is the vector of random effects with covariance matrix  $Q$ .

Here, we use the term “random effect” loosely, to refer to any term that enters the model as specified above, and it encompasses not only models for inter-group variability (e.g., random intercept), but also non-parametric covariate effects based on penalized splines (Wand 2003). This broad definition of random effects is for example discussed by Hodges (2013). In the case of i.i.d. normal random effects, the corresponding columns of  $R$  consist of 0s and 1s, indicating which group each row belongs to, and the  $\beta$ s are levels of the random effect (e.g., group-specific intercept), which are typically not of direct interest. In contrast, when using splines, each column of  $R$  corresponds to a basis function, where each row is its evaluation for given covariate values. The corresponding coefficients  $\beta$  are weights for the basis functions, which determine the shape of the resulting spline, and they are constrained by the random effect distribution to impose smoothness. For penalized splines, the value of  $\beta$  is of interest, as it is used to infer the shape of the non-linear relationship. The distribution of  $\beta$  can also be viewed as a (possibly improper) Bayesian prior, which captures the assumption that the resulting functions should be smooth (Wood 2017; Miller 2025).

The structure of the matrix  $Q$  reflects the dependence between entries of  $\beta$ . It is typically a block-diagonal matrix, where each block corresponds to a different random effect (or, equivalently, penalized spline) included in the model. The block for an i.i.d. normal random effect is diagonal, with diagonal elements equal to some variance parameter  $\sigma^2$ . The block for a penalized spline can be written as  $\lambda^{-1}S^{-}$ , where  $S^{-}$  is the pseudo-inverse of the penalty matrix for the chosen basis (Wood 2017), and  $\lambda > 0$  is a smoothness parameter. Note that the i.i.d. normal case can then be viewed as a special case of a penalized spline, where  $S$  is

the identity matrix, and where  $\lambda = 1/\sigma^2$ . In the following, we do not distinguish between variance and smoothness parameters.

An HMM includes many parameters (transition probabilities and state-dependent observation parameters), and each has an associated linear predictor as described above. There is therefore potentially one  $X$ ,  $\alpha$ ,  $R$ ,  $\beta$ , and  $Q$  for each HMM parameter. In practice, these can be combined and stored conveniently as block-diagonal matrices.

### *Transition probabilities*

When covariates affect the dynamics of the state process, we denote as  $\Gamma^{(t)}$  the transition probability matrix at time  $t$ , and the transition probabilities are related to the linear predictor using the multinomial logit (softmax) link,

$$\gamma_{ij}^{(t)} = \frac{\exp(\eta_{ij}^{(t)})}{\sum_{k=1}^K \exp(\eta_{ik}^{(t)})},$$

where  $\eta_{ij}^{(t)}$  is defined by Equation 2 for  $i \neq j$ , with the convention that  $\eta_{ii}^{(t)} = 0$ . Note that there are only  $K \times (K - 1)$  linear predictors for a  $K$ -state model, due to row constraints of the transition probability matrix. Here, we assume that the non-diagonal elements are estimated, but a different reference element could be chosen for each row. This gives rise to a non-homogeneous HMM, a useful framework to investigate the effects of covariates on the dynamics of the state process (Hughes *et al.* 1999; Patterson, Basson, Bravington, and Gunn 2009).

### *Observation parameters*

Covariate effects can also be included in the state-dependent parameters of the observation distributions. The state-dependent distribution is then time-varying, and we write  $Z_t \mid \{S_t = j\} \sim \mathcal{D}(\omega_{j1}^{(t)}, \dots, \omega_{jL}^{(t)})$ , where each  $h(\omega_{jl}^{(t)})$  is modeled as in Equation 2. In this case, the link function  $h$  depends on the domain of definition of each parameter  $\omega_{jl}$ ; for example, an identity link might be used for a real-valued parameter, or a log link for a positive-valued parameter. This formulation includes as a special case a rich family of Markov-switching regression models, including Markov-switching generalized linear and additive models (Kim *et al.* 2008; Langrock *et al.* 2017). In addition, because any distribution parameter can be specified as functions of covariates (e.g., variance or scale), this formulation includes as a special case the Markov-switching version of the generalized additive model for location, shape, and scale (GAMLSS; Rigby and Stasinopoulos 2005; Adam, Mayr, and Kneib 2022).

## **2.2. Implementation using TMB and mgcv**

### *Marginal likelihood with TMB*

Given (possibly multivariate) observations  $z_1, z_2, \dots, z_n$  from the process  $(Z_t)$ , the primary aim of the analysis is to estimate the fixed effect parameters  $\alpha$  (including intercepts and linear effects), and the smoothness parameters  $\lambda$ , which parameterize the covariance matrix  $Q$  of the random effect distribution. (A secondary aim might be to predict the random effect parameters  $\beta$  themselves, which we discuss later.) We propose using maximum likelihood

estimation, based on the marginal likelihood of the fixed effect and smoothness parameters (Skaug and Fournier 2006).

The distribution of  $\beta$  (Equation 2) can be captured by penalizing the joint log-likelihood of fixed and random effects, as,

$$l_p(\alpha, \beta, \lambda \mid z_1, \dots, z_n) = \log\{L(\alpha, \beta \mid z_1, \dots, z_n)\} - \sum_i \lambda_i \beta_i^\top S_i \beta_i$$

where  $L(\alpha, \beta \mid z_1, \dots, z_n)$  is the unpenalized HMM likelihood of parameters  $\alpha$  and  $\beta$  (e.g., computed using the forward algorithm), and the penalty is obtained as the log of the density function of a multivariate normal distribution (excluding additive constants). The sum is over all random effects included in the model.

For model fitting, we consider the marginal likelihood of  $\alpha$  and  $\lambda$ , obtained by integrating out the random effects  $\beta$ ,

$$\begin{aligned} L(\alpha, \lambda \mid z_1, \dots, z_n) &= \int \exp\{l_p(\alpha, \beta, \lambda \mid z_1, \dots, z_n)\} d\beta \\ &= \int p(z_1, \dots, z_n \mid \alpha, \beta) p(\beta \mid \lambda) d\beta \end{aligned}$$

where  $p(z_1, \dots, z_n \mid \alpha, \beta) = L(\alpha, \beta \mid z_1, \dots, z_n)$  is the HMM likelihood, and  $p(\beta \mid \lambda)$  is a multivariate normal density function with mean zero and block-diagonal precision matrix with blocks  $\lambda_i S_i$ . We use the Template Model Builder (**TMB**) R package to compute the marginal likelihood by integrating the random effects  $\beta$  based on the Laplace approximation (Kristensen, Nielsen, Berg, Skaug, and Bell 2016). The marginal likelihood can then be optimized numerically to obtain estimates of  $\alpha$  and  $\lambda$ . In cases where the levels of the random effects  $\beta$  are of direct interest, predicted values can be obtained, akin to best linear unbiased predictions. This is for example key for penalized splines, where these predictions are needed to derive the estimated non-linear function.

As an alternative to the marginal likelihood approach, Markov chain Monte Carlo could be used to sample over random effects in a Bayesian framework, and posterior inference could be carried out on all model parameters, as discussed in Section 4.8.

Although they do not focus on mixed effect models like the ones presented in this paper, Bacri, Berentsen, Bulla, and Hølleland (2022) and Bacri, Berentsen, Bulla, and Støve (2023) describe in detail the implementation of basic HMMs using **TMB**. These are good additional resources for users who would like to better understand the capabilities of **TMB**, or who might want to implement HMMs from scratch.

### *Spline specification with **mgcv***

The approach described so far for penalized splines assumes that the design matrix of basis functions  $R$  and the penalty matrix  $S$  are known. In practice, both matrices depend on the choice of the type of spline (e.g., cubic spline, or thin plate regression spline), and on the number of basis functions. We use the R package **mgcv** to specify  $R$  and  $S$  (Wood 2017), which provides great flexibility to define a wide range of basis-penalty models.

### *Confidence intervals*

**TMB** can output point estimates for all model parameters (including predicted values of the random effects), that we denote as  $(\hat{\alpha}, \hat{\beta}, \hat{\lambda})$ . It also computes the estimated joint precision

matrix of the parameters (in `TMB::sdreport()`), and we can take its inverse to get the estimated covariance matrix  $\hat{\Sigma}$ . Following asymptotic maximum likelihood theory, the estimators approximately follow  $(\hat{\alpha}, \hat{\beta}, \hat{\lambda}) \sim N[(\alpha, \beta, \lambda), \hat{\Sigma}]$ . Based on this result, we apply the following simulation-based procedure to create confidence intervals:

1. Generate a large number  $J$  of samples  $(\alpha^{(j)}, \beta^{(j)}, \lambda^{(j)}) \sim N[(\hat{\alpha}, \hat{\beta}, \hat{\lambda}), \hat{\Sigma}]$ ,  $j \in \{1, \dots, J\}$ .
2. For each  $(\alpha^{(j)}, \beta^{(j)})$ , derive a sample  $\theta^{(j)}$  of the HMM parameter of interest (i.e., transition probability or observation parameter), based on Equation 2.
3. Use quantiles of  $\{\theta^{(1)}, \theta^{(2)}, \dots, \theta^{(J)}\}$  as the bounds of a confidence interval for  $\theta$  (e.g., the 2.5% and 97.5% quantiles for a 95% confidence interval).

Note that, even though the  $\lambda^{(j)}$  are not used in Steps 2 and 3, it is important to sample from the joint distribution in Step 1 to account for the uncertainty in the smoothness parameters. Choosing a large value for  $J$  leads to a better approximation, but also to increased computational effort. In **hmmTMB**,  $J = 1000$  is used as the default, but it can be changed by the user in plotting and prediction functions.

In cases where  $\theta$  is a function of covariates, it can be derived over a grid of covariate values in Step 2, and the procedure then returns a pointwise confidence interval in Step 3. Under the view that the smoothness penalty captures a prior distribution for  $\beta$ , these are approximate credible intervals for penalized splines, which have the correct coverage probability “across the function” (as defined by Wood 2017, Section 6.10).

### 3. hmmTMB overview

The **hmmTMB** package is available on the Comprehensive R Archive Network (CRAN) at <https://CRAN.R-project.org/package=hmmTMB>, and its functionalities are described in several detailed vignettes, including example applications to financial and ecological data. Each function is also documented using **roxygen2** (Wickham, Danenberg, Csárdi, and Eugster 2022). Here, we provide an overview of the key features of the package (as of version 1.0.2), but we refer to the documentation and vignettes for more details. A simulation study is described in Appendix D, to investigate the performance of the method to estimate non-linear covariate effects and random effects.

#### 3.1. Installation

The package can be installed using `install.packages("hmmTMB")`. It depends on the following R packages, all available from CRAN: **TMB** (Kristensen *et al.* 2016), **mgcv** (Wood 2017), **RcppEigen** (Bates and Eddelbuettel 2013), **R6** (Chang 2021), and **ggplot2** (Wickham 2016). For all code examples in this paper, we used R version 4.4.0, **TMB** version 1.9.17, **mgcv** version 1.9-3, **RcppEigen** version 0.3.4.0.2, **R6** version 2.6.1, and **ggplot2** version 3.5.2.

#### 3.2. Data requirements

The data set for an HMM analysis needs to be provided as a `data.frame` object with one row for each observation time  $t \in \{1, 2, \dots, n\}$ . The data frame should include named columns for

all variables that are included in the model, either as response variables or as covariates. It should also contain a column named “ID”, to identify the time series in cases where multiple time series are analysed jointly. The likelihood of the full data set is computed as the product of the likelihoods of the individual time series (assuming independence). If ID is not provided, all observations are assumed to belong to the same time series. There is one other reserved column name: “**state**” should only be included when some states are known (Section 5.2). As an example, the following code chunk shows the required structure for a model with two response variables Z1 (real-valued) and Z2 (integer-valued), and one covariate cov1.

	ID	Z1	Z2	cov1
1	1	0.6187002	2	-0.10873270
2	1	-0.2424762	1	-0.02203936
3	1	-1.6081877	2	-0.02267833
4	1	-4.4758467	2	0.11035220
5	1	-1.9505381	4	0.11112987
6	1	0.7924493	4	0.02495790

The package can accommodate missing values in the response variables, and these should be entered as NA in the data frame passed as input. However, model fitting does not allow for missing values in the covariates. If a covariate includes NAs, these are replaced by the last non-missing covariate value (or the *next* non-missing value if there is no previous non-missing value). This is a crude solution, and each user should think about the best method to interpolate the covariate for their application. Variables included as random effects should be formatted as factors, as required by **mgcv**, rather than character strings or integers.

### 3.3. R6 syntax

We use an object-oriented framework based on the package **R6**. In this context, a class refers to a type of data structure, for which are defined some attributes (data and parameters) and some methods (functions specific to this class), and an object is an instance from a class. The general syntax is shown in the following code chunk.

```
R> object <- Class$new()
R> object$method()
```

In **hmmTMB**, there are three main classes, each representing a statistical model: ‘**MarkovChain**’, ‘**Observation**’, and ‘**HMM**’ (described in more detail below). The general workflow is to first create an object (representing a model), and then apply methods to update its attributes (e.g., to fit the model). The object is always updated directly, and so results are always stored in the same object. In the rest of this paper, we refer to a particular method as `Class$method()` (to make it clear which class it belongs to) even though, in practice, the method would be called on an object of that class.

One advantage of the object-oriented approach is that the hidden state model and observation model are created and stored separately in **hmmTMB**. As a result, one can for example simulate from a Markov chain, or plot state-dependent observation distributions, without the need to create a full HMM.



### 3.4. Main classes

We present the three main classes in **hmmTMB**, which are used to specify a hidden Markov model object.

#### *'MarkovChain' class*

The **'MarkovChain'** class stores attributes related to the state process model, including the number of states, the transition probability matrices, and the formulas specifying the dependence of transition probabilities on covariates. The constructor **MarkovChain\$new()** takes the following arguments.

- **data**: Data frame, required to define model matrices. This is needed even when no covariates are included, to determine the number of time points in the data.
- **formula**: Optional argument for model formulas. This can be either a single formula, which is then used for all transition probabilities, or a matrix of characters where each non-diagonal entry is the formula for a transition probability (and diagonal entries are set to ".").
- **n\_states**: Number of states  $\geq 2$  (optional if **tpm** is provided).
- **tpm**: Transition probability matrix used to initialize the model. If not provided, the matrix is initialized with 0.9 along the diagonal and  $0.1/(\mathbf{n\_states}-1)$  elsewhere.
- **initial\_state**: Character string indicating the model assumption used for the initial state distribution  $\delta^{(1)}$  of the Markov chain. Two important options are "estimated" (default), if  $\delta^{(1)}$  should be estimated, and "stationary", if  $\delta^{(1)}$  should be fixed to the stationary distribution of the transition probability matrix  $\Gamma^{(1)}$ .

#### *'Observation' class*

An object of the **'Observation'** class represents an observation model, defined by the list of observation variables and associated distributions, the state-dependent observation parameters, and relevant formulas if observation parameters dependent on covariates. To create an **Observation** object, the constructor **Observation\$new()** requires the following arguments.

- **data**: Data frame, containing at least the observation variables and covariates.
- **dists**: Named list of observation distributions for response variables.
- **formulas**: Optional nested list of formulas for the observation parameters. Each element corresponds to one of the observation variables, and it is itself a list, with one element for each parameter of the distribution. If this is not provided, no covariate dependence is assumed.
- **n\_states**: Number of states  $\geq 2$ .
- **par**: Nested list of initial parameter values, with similar structure to **formulas**.



Families of distributions that are implemented for the observation model include distributions for continuous unbounded data (e.g., normal `"norm"`), continuous bounded data (e.g., exponential `"exp"`, gamma `"gamma"`, beta `"beta"`), discrete or binary data (e.g., Poisson `"pois"`, binomial `"binom"`), among others. The full list of distributions is given in Appendix A, and additional distributions will be added in the future.

#### *'HMM' class*

The two main components of an `'HMM'` object are a `'MarkovChain'` object and an `'Observation'` object, which jointly describe the full specification of the model. These are the only two required arguments of the constructor `HMM$new()`. `'HMM'` is the main class that users will interact with, as it includes methods to perform most steps of a typical HMM analysis: model fitting, state decoding, uncertainty quantification, model checking, and plotting.

## 4. Statistical inference

We summarize the main steps of a typical HMM analysis, and the key functions in the `hmmTMB` package.

### 4.1. Model specification

The state process model and the observation model are specified separately, as `'MarkovChain'` and `'Observation'` objects, respectively. An HMM is then created by combining them.

#### *Choice of number of states*

The choice of the number of states needs to be made prior to model fitting. There is no general statistical method to estimate the “best” number of states, and standard model selection criteria (e.g., AIC, BIC) often favor models with many states, even when they cannot be interpreted (Langrock, Kneib, Sohn, and DeRuiter 2015). Pohle, Langrock, Van Beest, and Schmidt (2017) suggest using domain expertise and model checking to choose the number of states for a given study.

#### *Formula syntax*

Model formulas can be defined when creating a `'MarkovChain'` or an `'Observation'` object, to allow for covariate effects. The formulas do not have a left-hand side, and the syntax for the right-hand side is borrowed from `mgcv`, as that package is used to create model matrices. Generally, base R formula syntax can be used for linear terms, e.g., `~ x1 + x2 * x3` to include the linear effects of `x1`, `x2`, `x3`, and the interaction of `x2` and `x3`. Similarly, `poly()` can be used directly for polynomial terms.

For non-linear (“smooth”) terms, we use the function `s()` from `mgcv`, and we refer to its documentation for a detailed description. For example, we might use `~ s(x1, k = 10, bs = "cs")` to model the non-linear relationship between a parameter and `x1`, where `k` and `bs` determine the dimension and type of basis used to define the spline. Some relevant choices of `bs` are `"cs"` (cubic spline), `"ts"` (thin plate regression spline), and `"cc"` (cyclical spline). All three are “shrinkage” bases: the spline is shrunk to zero when  $\lambda \rightarrow \infty$ , such that it can be effectively excluded from the model as part of smoothness estimation (Wood 2017,

Section 5.4.3). The choice of  $k$  presents a trade-off between increased flexibility (large  $k$ ) and reduced computational cost (small  $k$ ). Past a certain point, increasing  $k$  will not change the model, as the smoothness of the spline is induced by the penalty.

Independent normal random intercepts are `s(x, bs = "re")`, where  $x$  is the factor variable used as group. The `mgcv` syntax also allows for interactions between covariates, including interactions between continuous and categorical variables (`s(x, by = y, ...)` where  $y$  is a factor) and two-dimensional smooths with equal smoothness along each dimension (`s(x, by = y, ...)` where  $y$  is numeric). Tensor products for two-dimensional smooths with different smoothness along each covariate are currently not supported.

For formulas in the observation model, **hmmTMB** allows for special functions of the form `state1()`, `state2()`, etc., to include covariate effects only in certain states. For example, one might want to estimate the effect of some covariate on the mean of the observation distribution in state 1, but not in state 2, using something like `~ state1(s(x, k = 10, bs = "cs"))`.

## 4.2. Parameter estimation

The main method for parameter estimation is `HMM$fit()`, which loads the (marginal) negative log-likelihood function from **TMB**, and optimizes it numerically using the function `optimx::optimx()` (Nash and Varadhan 2011). By default, `optimx()` uses a Nelder-Mead optimization method, but other approaches are implemented, and this can be specified in `HMM$fit()` with the argument `method` (see **optimx** documentation for details). Initial parameter values are required as a starting point for the optimizer, and their choice can be difficult for complex models; this is discussed in Section 4.7.

### *Accessing the parameter estimates*

After the model has been fitted by calling `HMM$fit()`, all model parameters are automatically updated to their estimates in the ‘**MarkovChain**’ and ‘**Observation**’ components. There are multiple methods in **hmmTMB** to access the different model parameters:

- `HMM$par()` returns estimates of the observation parameters  $\omega$  and transition probabilities  $\Gamma$ . By default, it computes parameters for the first row of data when covariates are included, but this can be changed with the argument `t` (a vector of row indices for which the parameters should be computed).
- `HMM$coeff_fe()` returns the fixed effect parameters  $\alpha$ , such as intercepts and linear effects.
- `HMM$lambda()` returns the smoothness parameters  $\lambda$ . Conversely, `HMM$sd_re()` returns the standard deviations of random effect terms, i.e.,  $1/\sqrt{\lambda_i}$ .
- `HMM$coeff_re()` returns predictions of the random effect parameters  $\beta$ , e.g., levels of a random intercept, or basis coefficients for a spline. These are for example used to generate plots of the non-linear relationships.

### *Uncertainty quantification and prediction*

The method `HMM$predict()` can be used to predict the HMM parameters  $\omega$  and  $\Gamma$ , includ-

ing confidence intervals, possibly for user-defined covariate values. It takes the following arguments.

- **what**: The model component that should be predicted, either **"tpm"** (transition probabilities), **"delta"** (stationary distribution of state process, computed from transition probability matrix), or **"obspar"** (observation parameters).
- **t**: Time point, or vector of time points, for which the parameters should be predicted. This is only used if **newdata** is not specified, i.e., if the original data set is used.
- **newdata**: Data frame with covariate values for which the parameters should be predicted. If this is not provided, then the original data set is used.
- **n\_post**: Number of posterior samples to use to compute confidence intervals (as described in Section 2.2). This defaults to 0, i.e., confidence intervals are not returned.

If confidence intervals are computed, the output is a list with three elements: **mean** (point estimate), **lcl** (lower bound), and **ucl** (upper bound). Each element is an array with one layer for each time point, and this output can for example be used to create plots of model parameters as functions of covariates, with confidence bands.

Confidence intervals for the fixed effect parameters  $\alpha$  and smoothness parameters  $\lambda$  can be computed with `HMM$confint()`, for example to quantify uncertainty on the linear effect of a covariate. These are derived by `TMB::sdreport()`, and the computational details are described in the documentation for that function.

### 4.3. State decoding

In many analyses, a key output is the classification of observations into states, i.e., an estimate of the hidden state  $S_t$  for each  $t \in \{1, 2, \dots, n\}$ . This procedure is sometimes called state decoding (Zucchini *et al.* 2017), and two different approaches are implemented in **hmmTMB**: global and local decoding. Given parameter estimates and the data, global decoding returns the most likely sequence of states. This is the sequence  $(\hat{s}_1, \hat{s}_2, \dots, \hat{s}_n)$  which maximizes the likelihood of the data, where  $\hat{s}_t \in \{1, \dots, N\}$  is the estimated state for time  $t$ . The most common way to find this sequence is the Viterbi algorithm, and this is implemented in the method `HMM$viterbi()` in **hmmTMB** (Viterbi 1967). In contrast, local decoding returns the probability of being in each state at each time point,  $P(S_t = j | Z_1, \dots, Z_n)$ . The method `HMM$state_probs()` computes these probabilities based on the forward-backward algorithm (Zucchini *et al.* 2017), and returns them as a matrix with one row for each time point and one column for each state.

### 4.4. Model selection

An HMM analysis requires many different modeling decisions, including the choice of the number of states, the choice of the observation distribution (e.g., exponential or gamma for a positive-valued variable), and the choice of covariate effects in the observation and hidden state components of the model. In most cases, we recommend using expert knowledge of the study system to choose the model formulation for a given data set and research question, rather than blindly follow model selection criteria.

Nonetheless, **hmmTMB** implements two different versions of the Akaike information criterion (AIC), which differ in how they measure complexity of models with random or non-linear effects. `HMM$AIC_marginal()` implements the marginal AIC, wherein the number of parameters of the model is defined as the number of fixed effect coefficients ( $\alpha$  in Section 2.1) plus the number of smoothness and variance parameters ( $\lambda$  in Section 2.1). By contrast, the conditional AIC implemented in `HMM$AIC_conditional()` (and accessible through the generic S3 method `AIC()`) also includes random effect coefficients ( $\beta$  in Section 2.1) in the number of parameters, accounting for smoothness constraints to derive an “effective” number of degrees of freedom. These two different approaches are for example described in Section 6.11 of Wood (2017). For models with non-linear and random effects, the marginal AIC tends to be overly conservative and the conditional AIC tends to be overly anti-conservative, so they should be used with caution (Wood 2017).

#### 4.5. Model checking

Model checking is crucial to determine whether the assumptions of the model are appropriate. A popular approach to model checking for HMMs is the use of pseudo-residuals, which are constructed such that they are independent and follow a standard normal distribution if the model formulation is correct (Zucchini *et al.* 2017). Similarly to linear regression residuals, deviations from normality indicate lack of fit, which can point to a failure of the observation distributions to model the data. In **hmmTMB**, pseudo-residuals can be computed with the method `HMM$pseudores()`. This returns a matrix with one row for each response variable and one column for each observation. To assess goodness-of-fit for the  $j$ -th variable in the model, we can then extract the  $j$ -th row of the matrix, and use tools such as quantile-quantile plots (`qqnorm()`) and autocorrelation function plots (`acf()`) to investigate deviations from normality and independence, respectively.

Alternatively, we can use posterior predictive checks for model assessment. The general idea is to simulate from the fitted model, and compare the simulated and observed data, to identify features that are not captured by the model. The method `HMM$check()` computes statistics, defined by the user through the argument `check_fn`, for the observed data and for a large number of simulated data sets. If the observed statistic falls in the far tail of the simulated distribution, this suggests that the corresponding feature of the data-generating process was not captured well by the model.

#### 4.6. Model visualization

In most cases, visualization of the fitted model can greatly help with interpretation. Several methods are implemented in **hmmTMB** to plot the results of an analysis. We describe them briefly here, but they are presented in more detail in the example analysis of Section 7, including figures showcasing their outputs. All plots are created with **ggplot2**, such that they can easily be edited by the user, e.g., to adjust captions, color palettes, or axis scales.

`HMM$plot()` is a flexible function to plot model parameters as functions of covariates, including confidence bands. The only two required arguments are `what` (which can be “obspar” for observation parameters, “tpm” for transition probabilities, or “delta” for stationary state probabilities), and `var` (name of covariate). `HMM$plot_ts()` creates one-dimensional or two-dimensional plots of data variables colored by the most likely state sequence returned by the Viterbi algorithm. `HMM$plot_dist()` produces a graph of the state-dependent distribu-

tions, combined with a histogram of the observations. The state-dependent distributions are weighted by the proportion of time spent in each state in the Viterbi state sequence, so that the sum of the weighted distributions can be compared to the histogram.

#### 4.7. Numerical errors and starting parameters

Numerical errors can arise during model fitting if the optimizer fails to converge to the global maximum of the likelihood. For example, it might return an error if the log-likelihood is evaluated to be infinite, suggesting that the likelihood is very close to zero. In other cases, the package might fail to estimate uncertainty for the model parameters. Each situation is unique, and it is impossible to provide general guidelines to resolve such problems, but we would like to suggest two main approaches: simplifying the model formulation, and trying different sets of initial parameters.

Complex model formulations require the estimation of a large number of parameters, and this can be a difficult task for the optimizer. In our experience, numerical errors are rare for 2-state covariate-free models, and this may therefore be a good place to start. Then, complexity may be increased incrementally, for example adding one covariate or one state at a time. Each time, the argument `init` of `HMM$new()` can be used to initialize the parameters of the new (more complex) model to the estimated values from the previous (simpler) model. This approach reduces the risk of numerical errors, helps with choosing a parsimonious model formulation, and makes troubleshooting easier.

The performance of the optimizer depends on the choice of initial parameters, where starting values closer to the truth will generally perform better. It is therefore important to choose initial values that are plausible, based on inspection of the data, and we recommend trying several sets of starting values to ensure that the optimizer always converges to the same model. In **hmmTMB**, the method `Observation$suggest_initial()` can help with selecting initial parameters. This function uses the  $K$  means algorithm, a simple clustering procedure which ignores the temporal structure of the data, to group observations into tentative “states” (James, Witten, Hastie, and Tibshirani 2013). It is then often straightforward to estimate the observation parameters from data within each group, and these are the suggested values. This is of course a crude approach but, in our experience, it usually performs well enough to circumvent numerical problems.

#### 4.8. Bayesian inference using Stan

The main workflow of **hmmTMB** presented in this section is based on the numerical optimization of the likelihood function implemented in `HMM$fit()`. As described in Section 2.2, **TMB** marginalizes the likelihood over random effects using the Laplace approximation, and it is the marginal likelihood that we optimize.

An alternative approach is to use Markov chain Monte Carlo to sample over random effects (rather than integrate them out), and this can be done with `HMM$fit_stan()`. This method relies on the package **tmbstan** (Monnahan and Kristensen 2018), which automatically uses a **TMB** model object to generate the equivalent Stan code (Carpenter *et al.* 2017). The model is then fitted using Hamiltonian Monte Carlo (HMC) methods, and posterior samples for all model parameters are returned. An example Bayesian analysis is included in Appendix C.

*Prior specification*

By default, improper flat priors are set for all model parameters, but these can be changed by the user. Specifically, normal priors can be specified for all fixed effect parameters  $\alpha$  and smoothness parameters  $\lambda$ , using the method `HMM$set_priors()`. This function takes as input a list of matrices, where each matrix corresponds to a different model component (e.g., fixed effect parameters for hidden state model, fixed effect parameters for observation model). Each matrix has one row for each parameter of the corresponding model component, and two columns, for the means and standard deviations of the normal priors, respectively. A row set to `NA` corresponds to an improper flat prior. If there is demand, prior distributions could be extended to all observation distributions implemented in **hmmTMB** (listed in Appendix A).

*Model fitting*

The model can then be fitted using the method `HMM$fit_stan()`, which is a wrapper for `tmbstan::tmbstan()`. The main arguments required by this method are `chains`, the number of HMC chains to run, and `iter`, the number of HMC iterations in each chain. The choice of these settings will be highly context dependent, and we refer readers to the **Stan** documentation for more information. By default, **tmbstan** samples both fixed and random effects; alternatively, a hybrid approach can be used where fixed effects are sampled and random effect are integrated using the Laplace approximation, with the argument `laplace = TRUE`.

Once the algorithm has run, the output is an object of class ‘`stanfit`’ (from the **rstan** package [Stan Development Team 2024](#)), and it can be accessed with `HMM$out_stan()`. This can either be transformed into a matrix with `as.matrix()` to inspect the posterior samples, or it can directly be passed to **rstan** functions for visualization. In particular, `rstan::stan_trace()` and `rstan::stan_dens()` create trace and density plots of the posterior samples, respectively. Alternatively, **hmmTMB** automatically computes posterior samples for the HMM parameters directly (rather than linear predictor parameters), and these can be accessed using `HMM$iters()`. This method returns a matrix with one row for each HMC iteration, and one column for each HMM parameter (observation parameters and transition probabilities). These transformed posterior draws are computed for the first time point of the data, and are of limited utility in models where parameters are time-varying, i.e., if covariates are included.

*Comparison to maximum likelihood estimation*

The choice of using the Bayesian or maximum likelihood approach will depend on each user’s preference and analysis. For example, the Bayesian method has the appeal of including prior information when it is available, and allowing for full posterior inference, but this comes at a computational cost which might make the approach infeasible for large data sets or complex models.

In **hmmTMB**, a key difference between the two approaches is the way they deal with random effects. The maximum (marginal) likelihood fitting procedure is based on the Laplace approximation to integrate over random effects. This makes the assumption that the likelihood can be approximated reasonably well by a normal density function. The HMM likelihood is often a complex function with multiple (local) modes ([Zucchini et al. 2017](#)), which suggests that the Laplace approximation might not be appropriate in some cases. Research is needed to investigate this problem, and we hope that **hmmTMB** can greatly help with this, as it allows for both approaches to be applied to the same model, and outputs to be compared.



## 5. Other features

### 5.1. Simulation

The method `HMM$simulate()` simulates observations from an HMM, either using the estimated parameters if the model has been fitted, or using the initial parameters otherwise (e.g., to simulate from a user-defined HMM). The state sequence is first simulated from the Markov chain model, and the observations are then generated based on the simulated states. In models with covariates, the function requires a data frame with one column for each covariate, to be passed as the argument `data`. This function is used in `HMM$check()` to compare user-defined features of the observed data to simulated data, for model assessment.

### 5.2. Supervised learning

In most applications, HMMs are used in an unsupervised setting: the states are not known prior to the analysis, and they are entirely data-driven. However, in some cases, information might be available about the states before the analysis. For example, the analyst might be able to classify observations into pre-defined states manually for a subset of the data, but not for the full data set because of time constraints. Then, the aim of the analysis is two-fold: characterizing the pre-defined states with a statistical model, and classifying the full data set. This approach is sometimes called “semi-supervised” (Leos-Barajas, Photopoulou, Langrock, Patterson, Watanabe, Murgatroyd, and Papastamatiou 2017b).

Supervised (or semi-supervised) learning can be implemented in **hmmTMB**, by including a column named `state` in the input data frame. In a  $K$ -state model, this column should contain numbers between 1 and  $K$  for time points where the state is known, and `NA` elsewhere. It is detected by the package, and used to keep the known states fixed (while the unknown states are inferred).

### 5.3. Fixed parameters

In some analyses, it might be useful to fix a parameter to a given value, rather than estimate it. This might for example be the case if a transition is impossible, i.e., the corresponding transition probability should be fixed to zero. Alternatively, we might sometimes want to constrain several parameters in the model to be estimated to a shared value. Such model formulations can be implemented in **hmmTMB** using the `fixpar` argument of `HMM$new()`. This is based on the parameter mapping functionality offered by **TMB**, with the argument `map` of `TMB::MakeADFun()`, and we refer the reader to the **TMB** documentation for details.

Specifically, `fixpar` is defined as a named list with the following optional elements: `obs` (observation parameters), `hid` (transition probabilities), `lambda_obs` (smoothness parameters for observation model), `lambda_hid` (smoothness parameters for state model), and `delta0` (initial distribution of state process). Each of these entries should be a named vector, in a format similar to that expected for the argument `map` of `TMB::MakeADFun()`. Each element of the vector should be named after a parameter, and its value should be `NA` for a fixed parameter, to indicate that it should be fixed to its initial value. If several parameters have a common value, the corresponding elements in `fixpar` should be set to some shared integer. The relevant parameter names can be found using `HMM$coeff_list()`.



#### 5.4. General dependence structures

The flexible parameter constraints presented in Section 5.3 make it possible to implement HMMs with non-standard dependence structures in **hmmTMB**. These models relax either the Markov assumption of the state process, or the assumption that successive observations are independent conditional on the state. These extensions often come with greatly increased computational cost, because they suffer from the curse of dimensionality (the number of parameters increases quickly with model complexity). Parameters from these more complex methods are also often difficult to interpret, and so we recommend parsimony when choosing a model formulation for a given analysis.

##### *Hidden semi-Markov models*

The dwell times of a Markov chain (sometimes called sojourn times) are the time intervals between state transitions. As a consequence of the Markov assumption (Equation 1), dwell times in state  $j$  follow a geometric distribution with parameter  $1 - \gamma_{jj}$ . Semi-Markov models relax the Markov assumption, and increase flexibility by explicitly modeling the dwell time distribution (Zucchini *et al.* 2017, Chapter 12). It has been shown that a hidden semi-Markov model can be approximated by a standard HMM with expanded state space (Langrock and Zucchini 2011). Briefly, this requires representing each state of the semi-Markov model with several states in a Markov chain, and constraining the transition probabilities to control the time spent in each composite state. The approach of Langrock and Zucchini (2011) can be implemented in **hmmTMB** by forcing the relevant transition probabilities to be fixed to zero. This method can in principle be used to estimate any dwell time distribution with arbitrary precision but, in practice, the memory requirements might be prohibitive.

##### *Higher-order hidden Markov models*

Another approach to relax the Markov assumption is to allow the state  $S_t$  to depend not only on  $S_{t-1}$ , but also on  $S_{t-2}$  (and possibly  $S_{t-3}, S_{t-4}$ , etc.). The resulting model is called a higher-order Markov chain, which can incorporate more “memory” about the history of the state process than the first-order model (Zucchini *et al.* 2017, Section 10.3). A higher-order Markov chain ( $\tilde{S}_t$ ) can be viewed as a first-order Markov chain ( $S_t$ ) by defining  $S_t = (\tilde{S}_t, \tilde{S}_{t-1}, \dots, \tilde{S}_{t-k+1})$ . For example, if  $(\tilde{S}_t)$  is a 2-state second-order Markov chain, then  $(S_t)$  is a 4-state first-order Markov chain, with state space  $\{(1, 1), (1, 2), (2, 1), (2, 2)\}$ . Some transition probabilities in this 4-state model must be fixed to zero; for example, it is impossible to transition from  $S_t = (\tilde{S}_t, \tilde{S}_{t-1}) = (1, 1)$  to  $S_{t+1} = (\tilde{S}_{t+1}, \tilde{S}_t) = (1, 2)$ .

##### *Autoregressive hidden Markov models*

The assumption that successive observations of an HMM are conditionally independent given the state process can be relaxed, for example to capture strong autocorrelation in the observation process. A practical approach is to include the observation  $Z_{t-1}$  (and possibly  $Z_{t-2}$  and so on) as a covariate on the parameters of the distribution of  $Z_t$ . For example, a Markov-switching first-order autoregressive model could be written as  $Z_t \mid \{S_t = j\} \sim N(\alpha_j Z_{t-1}, \sigma_j^2)$ . The implementation in **hmmTMB** would require creating an observation model that includes the linear effect of a lagged version of the observed variable on the mean parameter (and no intercept). The slope of this effect is  $\alpha_j$ . A similar approach can be used to implement Markov-switching Gaussian random walks, correlated random walks, and higher-order autoregressive models.

### *Coupled hidden Markov models*

A coupled HMM models two observed variables  $Z_t^{(1)}$  and  $Z_t^{(2)}$  as being driven by two separate, but dependent, state processes  $(S_t^{(1)})$  and  $(S_t^{(2)})$ , respectively (Pohle, Langrock, Van der Schaar, King, and Jensen 2021). Several possible assumptions can be made about the dependence of the state processes, and here we focus on the Cartesian product model of Pohle *et al.* (2021), which can be implemented in **hmmTMB**. The Cartesian product model is based on the process  $(S_t)$  where  $S_t = (S_t^{(1)}, S_t^{(2)})$ . For example, if  $(S_t^{(1)})$  and  $(S_t^{(2)})$  are 2-state Markov chains, then  $(S_t)$  is a 4-state Markov chain with state space  $\{(1, 1), (1, 2), (2, 1), (2, 2)\}$ . The coupled model can therefore be implemented as a 4-state HMM with some constraints on the observation parameters: the distribution parameters of  $Z_t^{(1)}$  are the same in states  $S_t = 1$  and  $S_t = 2$  (because both imply  $S_t^{(1)} = 1$ ), and so on.

### 5.5. Accessing TMB outputs

Users familiar with **TMB** may be interested to directly access and manipulate its outputs, for example to evaluate the negative log-likelihood of the model, or to find the covariance matrix of the parameters. The method `HMM$tmb_obj()` returns the object created by `TMB::MakeADFun()`. This is a list which includes `fn` (a function object for negative log-likelihood), and `gr` (function object for gradient of negative log-likelihood). After likelihood optimization, **hmmTMB** uses the function `TMB::sdreport()`, which computes covariance (or precision) matrices for all model parameters, used for uncertainty quantification. The output of `TMB::sdreport()` can be accessed using the method `HMM$tmb_rep()`.

## 6. Comparison to other packages

There are many software packages for HMMs, which greatly differ in the model formulations that they include, and their focus on particular data types. Although a comprehensive review is infeasible, we contrast the features of some of the most popular and flexible packages to **hmmTMB**, to place it within its broader context.

In R, two popular packages are **depmixS4** and **msm**, which both provide great flexibility for model formulation (Visser and Speekenbrink 2010; Jackson 2011). In particular, similarly to **hmmTMB**, they can accommodate multivariate data, (linear) covariate effects on transition probabilities, missing data, multiple time series, and parameter constraints. Although **depmixS4** includes only a few observation distributions by default, it allows for user-defined distributions. It might therefore be a good alternative in cases where a distribution which is not included in **hmmTMB** is needed. **depmixS4** also allows for more complex parameter constraints than **hmmTMB**, for example inequalities between model parameters. **msm** has a focus on continuous-time HMMs, which are for example widely used in medicine. Although **hmmTMB** does not currently support continuous-time HMMs, we plan to include this functionality in a future release.

The R packages **moveHMM** and **momentuHMM** are also widely used to apply HMMs, with a focus on animal tracking studies in ecology (Michelot, Langrock, and Patterson 2016; McClintock and Michelot 2018). **moveHMM** is specialized for the analysis of two-dimensional animal movement data, making it easier to use than alternatives, but limiting its flexibility in other analyses. **momentuHMM** extends **moveHMM** to more general model formulations, and it

allows for multivariate data, a wide range of observation distributions, parameter constraints, and (linear) covariate effects on all model parameters. In addition, it allows for discrete random effects, as described for example by [Maruotti and Rydén \(2009\)](#) and [McKellar, Langrock, Walters, and Kesler \(2015\)](#). In this approach, a parameter may take several discrete values, and each individual or group is allocated to one of the values. This stands in contrast with the continuous random effects implemented in **hmmTMB** (Section 2.1).

In Python, the main HMM package is **hmmlearn** ([Lebedev 2022](#)). It includes three families of observation distributions (normal, normal mixture, and multinomial), and requires user-defined functions for other distributions. It can be used on multiple time series, and can include parameter constraints; however, it does not allow for covariate effects on the HMM parameters.

Overall, **hmmTMB** is one of the fastest and most flexible available packages for HMMs. Its most unique feature is the ability to implement very general non-parametric and hierarchical models for the transition probabilities and the observation parameters (Section 2.1). This includes splines and their interactions, (continuous) random intercepts and random slopes, and the large family of hierarchical generalized additive models ([Pedersen, Miller, Simpson, and Ross 2019](#)). The interface with **Stan** will also greatly increase the accessibility of Bayesian inference methods for HMMs (although see [Damiano 2023](#)). Finally, the package includes a few smaller features which we expect to be useful in many HMM analyses, such as model checking using posterior predictive simulations, and customizable visualizations based on **ggplot2**.

## 7. Reproducible example

We illustrate the main features of **hmmTMB** through an example analysis of animal tracking data, including non-linear covariate effects and random effects in the state process model. Two additional examples with code are available in the supplementary material: analysis of human activity data with cyclical covariate effect in a Bayesian framework (Appendix C), and analysis of energy prices with regime-switching generalized additive model (Appendix B).

### 7.1. Petrel movement analysis

We considered time-indexed GPS locations from 10 Antarctic petrels, described by [Descamps \*et al.\* \(2016a\)](#) and available through the Movebank data repository ([Descamps \*et al.\* 2016b](#)). HMMs are commonly used on such data to describe the animals' movement in terms of a few states, often interpreted as behavioral states ([Patterson \*et al.\* 2009](#)). The variables of interest are the step lengths (distances between successive locations) and turning angles (angles between successive segments of the track), and we derived them using the R package **moveHMM** ([Michelot \*et al.\* 2016](#)). These two variables are often informative about animal behavior, because they capture both the speed and tortuosity of movement. In addition to step length and turning angle, we derived distances between each bird's locations and its "central location", taken to be the first location of its track, and interpreted as a colony. We were interested in modeling behavioral states of petrels, and understanding how they are driven by the distance to central location. Observations were obtained at regular 30-min intervals, and covered between one and three weeks for the different birds, resulting in 6296 observations for each variable. The code used to preprocess the Movebank data, and to create

the data frame used below, is given in the replication materials.

```
R> library("hmmTMB")
R> data$ID <- factor(data$ID)
R> head(data)
```

	ID	time	lon	lat	d2c	step
1231	PET-A	2011-12-13 00:22:57	7.231573	-69.89723	0.000000	8.406119
1232	PET-A	2011-12-13 00:52:57	7.258316	-69.82245	8.406119	12.358492
1233	PET-A	2011-12-13 01:22:57	7.311297	-69.71320	20.759693	13.436032
1234	PET-A	2011-12-13 01:52:57	7.389327	-69.59586	34.169551	13.041818
1235	PET-A	2011-12-13 02:22:57	7.488740	-69.48424	47.137940	13.074831
1236	PET-A	2011-12-13 02:52:57	7.586307	-69.37217	60.174068	12.678806

	angle
1231	NA
1232	-0.043899821
1233	-0.061301455
1234	-0.074846274
1235	0.004979084
1236	0.034615473

We modeled step length  $L_t > 0$  with a gamma distribution, parameterized in terms of its state-dependent mean  $\omega_{1j} > 0$  and standard deviation  $\omega_{2j} > 0$ , and turning angle  $\varphi_t \in (-\pi, \pi]$  with a wrapped Cauchy distribution with mean zero, parameterized in terms of a state-dependent concentration parameter  $\omega_{3j} \in [0, 1]$ :

$$\begin{cases} L_t \mid \{S_t = j\} \sim \text{gamma}(\omega_{1j}, \omega_{2j}) \\ \varphi_t \mid \{S_t = j\} \sim \text{wrapped Cauchy}(0, \omega_{3j}) \end{cases}$$

The mean turning angle in each state is fixed to zero (rather than treated as an unknown parameter) because the petrels' movement has strong directional persistence, as is common in high-resolution tracking data.

We used a 3-state model, with the expectation that state 1 would capture slow movement, state 3 would capture fast movement, and state 2 would be intermediate. In the state process model, we included a random intercept for the individual  $I_t \in \{1, 2, \dots, 10\}$ , to capture inter-individual variability (McClintock 2021). We also included a non-linear effect of the distance to center  $d_t > 0$  on the transition probabilities. We constrained the transition probabilities between states 1 and 3 ( $\gamma_{13}$  and  $\gamma_{31}$ ) to be zero, under the assumption that animals need to transition through the intermediate state 2. Finally, the hidden state model can be summarized with the following expressions of the linear predictor,

$$\eta_{ij}^{(t)} = \begin{cases} 0 & \text{if } i = j \\ -\infty & \text{if } (i, j) = (1, 3) \text{ or } (3, 1) \\ \alpha_{ij} + \beta_{ij}^{(I_t)} + f_{ij}(d_t) & \text{otherwise,} \end{cases}$$

where  $\beta_{ij}^{(k)} \sim N(0, 1/\lambda_{ij})$  is the random intercept for track  $k$ , and where  $f_{ij}$  is a smooth function, modeled with a spline.

## 7.2. Model specification

### *Observation model*

Creating an `Observation` object requires defining a list of observation distributions; we use the "gamma2" distribution for step length (gamma distribution parameterized by mean and standard deviation), and the "wrpcauchy" distribution for turning angle (wrapped Cauchy distribution). We also need to define a list of initial parameter values for the optimizer; here, we choose them based on data visualization, but we could also update them based on `Observation$suggest_initial()`.

```
R> step_mean0 <- c(1, 6, 20)
R> step_sd0 <- c(1, 5, 10)
R> angle_mean0 <- c(0, 0, 0)
R> angle_rho0 <- c(0.8, 0.8, 0.9)
R> par0 <- list(step = list(mean = step_mean0, sd = step_sd0),
+   angle = list(mu = angle_mean0, rho = angle_rho0))
R> dists <- list(step = "gamma2", angle = "wrpcauchy")
R> obs <- Observation$new(data = data, dists = dists,
+   n_states = 3, par = par0)
```

### *Hidden state model*

If all transition probabilities had the same model formula (i.e., if they all depended on the same covariates), we could simply pass this formula to `MarkovChain$new()`. However, in this example, the transition probabilities  $\gamma_{13}$  and  $\gamma_{31}$  do not depend on covariates, as they are fixed to zero. So, we need to create a matrix to specify the structure of the model, where each element is the formula (as character string) for the corresponding transition probability. The diagonal elements are used as references for each row, so their formulas are set to ".". We create a matrix of initial values for the transition probabilities, where  $\gamma_{13}$  and  $\gamma_{31}$  are set to zero. The other elements are used to initialize the intercepts for the corresponding transition probabilities.

```
R> f <- "~ s(ID, bs = 're') + s(d2c, k = 10, bs = 'cs')"
```

```
R> tpm_structure <- matrix(
+   c(".", f, "~1",
+     f, ".", f,
+     "~1", f, "."),
+   ncol = 3, byrow = TRUE)
R> tpm0 <- matrix(
+   c(0.9, 0.1, 0,
+     0.1, 0.8, 0.1,
+     0, 0.1, 0.9),
+   ncol = 3, byrow = TRUE)
R> hid <- MarkovChain$new(n_states = 3, formula = tpm_structure,
+   data = data, tpm = tpm0, initial_state = "stationary")
```

We used the option `initial_state = "stationary"` to fix the initial distribution  $\delta^{(1)}$  of the state process to the stationary distribution of  $\Gamma^{(1)}$  for each time series. We do this for computational convenience to circumvent common identifiability problems for the estimation of the initial distribution. However, note that the Markov process is not stationary due to effects of time-varying covariates, and so this does not define a stationary HMM (Zucchini *et al.* 2017). Another approach would be to set `initial_state` to a vector of integers, which would be used as the fixed initial state for each time series, if the initial states were known.

### *Hidden Markov model*

We need to define two different parameter constraints: the means of the turning angle distributions are fixed to zero, and the transition probabilities between states 1 and 3 are fixed to zero. We have initialized the parameters to these values, and we now need to let the package know that they should not be estimated (i.e., kept fixed at the initial values), using the `fixpar` argument of `HMM$new()`. We do this as explained in Section 5.3, and we can then create the ‘HMM’ object.

```
R> fixpar <- list(
+   obs = c("angle.mu.state1.(Intercept)" = NA,
+   "angle.mu.state2.(Intercept)" = NA,
+   "angle.mu.state3.(Intercept)" = NA),
+   hid = c("S1>S3.(Intercept)" = NA,
+   "S3>S1.(Intercept)" = NA))
R> hmm <- HMM$new(obs = obs, hid = hid, fixpar = fixpar)
```

We can now fit the model with `HMM$fit()`. This takes about 1.5 min on a desktop computer with a 13th generation Intel i9-13900 CPU.

```
R> hmm$fit(silent = TRUE)
```

## 7.3. Results

We can get estimates of all model parameters using the methods described in Section 4.2. Here, we showcase plotting methods from the package, which can greatly help with model interpretation.

### *Interpreting the states*

The most likely state sequence is computed by `HMM$viterbi()`. It is often used to create a plot of the data, colored by state, and this is implemented in `HMM$plot_ts()` (Figure 1).

```
R> hmm$plot_ts("lon", "lat") +
+   coord_map("mercator") +
+   geom_point(size = 0.3) +
+   labs(x = "longitude", y = "latitude")
```

Another useful output is plots of the estimated state-dependent density functions for step length and turning angle. These plots are generated by the method `HMM$plot_dist()`, and shown in Figure 2.




Figure 1: Map of petrel tracks, colored by most likely state sequence.




Figure 2: Estimated distributions of step length (left) and turning angle (right) in the petrel analysis. The densities are weighted by the proportion of the states in the Viterbi sequence.

```
R> hmm$plot_dist("step") +
+   coord_cartesian(ylim = c(0, 0.45)) +
+   theme(legend.position = c(0.8, 0.8)) +
+   labs(x = "step (km)")
R> hmm$plot_dist("angle") +
+   coord_cartesian(ylim = c(0, 2.2)) +
+   theme(legend.position = "none") +
+   scale_x_continuous(breaks = seq(-pi, pi, by = pi/2),
+     labels = expression(-pi, -pi/2, 0, pi/2, pi))
```

Both Figures 1 and 2 suggest that the main difference between the states is in the speed



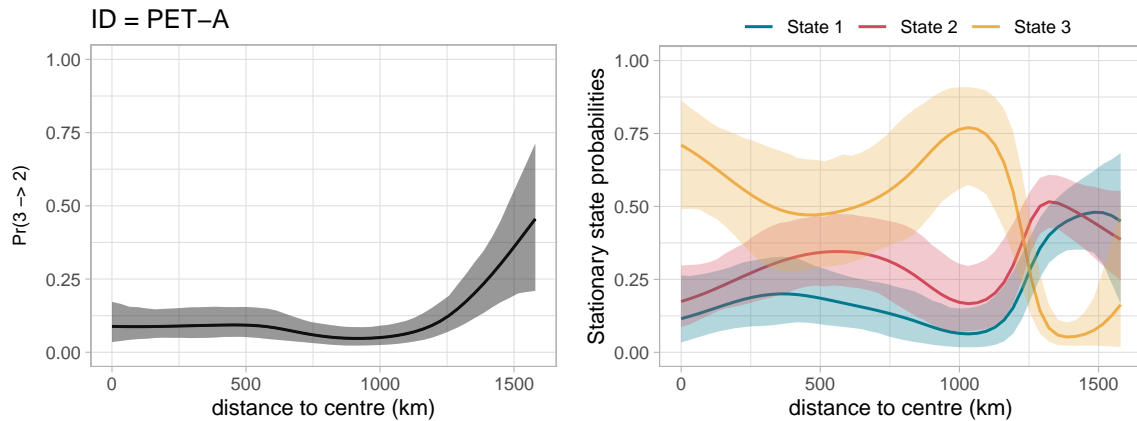


Figure 3: Transition probability  $\gamma_{32}$  (left) and stationary state probabilities (right) as functions of distance to central location (in km), with 95% pointwise confidence bands, for petrel analysis. The random intercept for individual PET-A is used.

of movement: state 1 captured slow movement, state 3 is fast movement, and state 2 is intermediate. All three states have high turning angle concentrations, i.e., strong directional persistence. In the following, we assume that these states correspond to three behaviors, “foraging” (state 1), “exploring” (state 2), and “traveling” (state 3), but we recognize that this likely ignores other important behaviors (e.g., resting on water).

#### *Effect of distance to center*

The main aim of this analysis was to investigate drivers of behavioral switching in petrels. We can use the method `HMM$plot()` to plot the transition probabilities (`what = "tpm"`), or the stationary state probabilities (`what = "delta"`), as functions of covariates. We use the arguments `i` and `j` to generate a plot of the transition probability  $\gamma_{32} = P(S_t = 2 \mid S_{t-1} = 3)$ , which is of particular interest here.

```
R> hmm$plot(what = "tpm", var = "d2c", i = 3, j = 2) +
+   labs(x = "distance to centre (km)")
R> hmm$plot(what = "delta", var = "d2c") +
+   theme(legend.position = "top", legend.margin = margin(c(0, 0, -10, 0))) +
+   labs(title = NULL, x = "distance to centre (km)")
```

Figure 3 indicates that the fast “traveling” state was most probable for distances to central location smaller than 1000km, but this probability decreases sharply above 1000 km. This suggests that petrels tend to travel at high speeds to reach good foraging grounds far from their central location, and slow down to forage (or search for food) when they reach those areas of interest.

#### *Inter-individual heterogeneity*

We can visualize the inter-individual heterogeneity in the state process using the `HMM$plot()` method, which shows the predicted values of the random effect with confidence intervals (Figure 4). This requires setting the other covariate(s) to a fixed value, and here we choose  $d_t = 1500$  km.




Figure 4: Predicted individual-specific random intercepts for petrel analysis, with 95% confidence intervals.

```
R> hmm$plot(what = "delta", var = "ID", covs = list(d2c = 1500))
```

In many applications, these random intercepts are not of direct interest, and we would instead report their variance. The standard deviation of the random effects, returned by `HMM$sd_re()`, quantifies inter-individual variability.

```
R> hmm$sd_re()$hid
```

```

      [,1]
S1>S2.s(ID) 0.297697847
S1>S2.s(d2c) 0.012239009
S2>S1.s(ID) 0.111309940
S2>S1.s(d2c) 0.001021525
S2>S3.s(ID) 0.465517000
S2>S3.s(d2c) 0.123034525
S3>S2.s(ID) 0.414212750
S3>S2.s(d2c) 0.025433105

```

The function returns two standard deviations for each transition probability: one for the random intercept (interpreted as the standard deviation of the normal distribution), and one for the spline (inversely related to its smoothness). The distributions of random intercepts were estimated to  $\beta_{12}^{(k)} \sim N(0, 0.30^2)$ ,  $\beta_{21}^{(k)} \sim N(0, 0.11^2)$ ,  $\beta_{23}^{(k)} \sim N(0, 0.47^2)$ , and  $\beta_{32}^{(k)} \sim N(0, 0.41^2)$ .

## 8. Conclusion

HMMs are widely applied in many fields, and we anticipate that **hmmTMB** will be of general interest among statisticians and practitioners. The flexible model formulation allowed by the package, including covariate effects and parameter constraints, makes it possible to implement many general Markov-switching models, placing it at the cutting edge of HMM research.

The package is based on **TMB** to take advantage of its great flexibility and computational tractability, but this comes with some pitfalls. In particular, the Laplace approximation is used to derive the marginal likelihood of the model, and this may perform poorly in cases where the likelihood is multimodal or asymmetric. The performance of the Laplace approximation for HMMs is an open question, and we hope that this can be investigated by comparing outputs from `HMM$fit()` to outputs from `HMM$fit_stan()` (which samples random effects using **Stan**). It is also important to note that, regardless of the implementation method, complex model formulations (e.g., including multiple penalized splines) are susceptible to numerical instability and prohibitive computational cost. It is more challenging to estimate covariate effects in HMMs than in generalized linear models, for example, because parts of the model are latent and need to be inferred. Although **hmmTMB** makes it straightforward in principle to include many complex relationships in a model, we recommend parsimonious formulations.

The modular structure of **hmmTMB** makes it easy to extend, and we plan to include additional functionalities in the future. For example, we recognize that some important probability distributions are not currently included in the package, and we are open to request for additional distributions. Another important extension would be to implement continuous-time HMMs, which allow for observations at irregular time intervals, and where state transitions can occur at any time, making them popular in areas where sampling is irregular (Jackson, Sharples, Thompson, Duffy, and Couto 2003). It might also be possible to implement hierarchical HMMs within **hmmTMB** in the future, if such a model can be rewritten as an HMM over an expanded state space, similarly to the general dependence structures described in Section 5.4. Hierarchical HMMs are based on two nested state processes, and they have recently been used to study the patterns of financial markets (Oelschläger and Adam 2023) and the behavior of animals (Leos-Barajas, Gangloff, Adam, Langrock, Van Beest, Nabe-Nielsen, and Morales 2017a) over multiple time scales.

## Acknowledgments

I am very grateful to Richard Glennie, who provided key ideas and code in the early development of this package. I would also like to thank Carlina Feldmann, Sina Mews, Brett McClintock, Vinky Wang, and Jan-Ole Koslik for helpful feedback about the package, and Natasha Klappstein for providing comments about an earlier version of this paper.

## References

- Adam T, Mayr A, Kneib T (2022). “Gradient Boosting in Markov-Switching Generalized Additive Models for Location, Scale, and Shape.” *Econometrics and Statistics*, **22**, 3–16. doi:[10.1016/j.ecosta.2021.04.002](https://doi.org/10.1016/j.ecosta.2021.04.002).
- Altman RM (2007). “Mixed Hidden Markov Models: An Extension of the Hidden Markov Model to the Longitudinal Data Setting.” *Journal of the American Statistical Association*, **102**(477), 201–210. doi:[10.1198/016214506000001086](https://doi.org/10.1198/016214506000001086).
- Altman RM, Petkau AJ (2005). “Application of Hidden Markov Models to Multiple Sclerosis Lesion Count Data.” *Statistics in Medicine*, **24**(15), 2335–2344. doi:[10.1002/sim.2108](https://doi.org/10.1002/sim.2108).

- Bacri T, Berentsen GD, Bulla J, Hølleland S (2022). “A Gentle Tutorial on Accelerated Parameter and Confidence Interval Estimation for Hidden Markov Models Using Template Model Builder.” *Biometrical Journal*, **64**(7), 1260–1288. doi:10.1002/bimj.202100256.
- Bacri T, Berentsen GD, Bulla J, Støve B (2023). “Computational Issues in Parameter Estimation for Hidden Markov Models with Template Model Builder.” *Journal of Statistical Computation and Simulation*, **93**(18), 3421–3457. doi:10.1080/00949655.2023.2226788.
- Bates D, Eddelbuettel D (2013). “Fast and Elegant Numerical Linear Algebra Using the **RcppEigen** Package.” *Journal of Statistical Software*, **52**(5), 1–24. doi:10.18637/jss.v052.i05.
- Bulla J, Bulla I (2006). “Stylized Facts of Financial Time Series and Hidden Semi-Markov Models.” *Computational Statistics & Data Analysis*, **51**(4), 2192–2209. doi:10.1016/j.csda.2006.07.021.
- Carpenter B, Gelman A, Hoffman M, Lee D, Goodrich B, Betancourt M, Brubaker M, Guo J, Li P, Riddell A (2017). “Stan: A Probabilistic Programming Language.” *Journal of Statistical Software*, **76**(1), 1–32. doi:10.18637/jss.v076.i01.
- Chang W (2021). **R6: Encapsulated Classes with Reference Semantics**. doi:10.32614/CRAN.package.R6. R package version 2.5.1.
- Damiano L (2023). “**BayesHMM**: Full Bayesian Inference for Hidden Markov Models.” R package version 0.0.1, URL <https://github.com/luisdamiano/BayesHMM>.
- Descamps S, Tarroux A, Cherel Y, Delord K, Godø OR, Kato A, Krafft BA, Lorentsen SH, Ropert-Coudert Y, Skaret G, Varpe Ø (2016a). “At-Sea Distribution and Prey Selection of Antarctic Petrels and Commercial Krill Fisheries.” *PLOS One*, **11**(8), e0156968. doi:10.1371/journal.pone.0156968.
- Descamps S, Tarroux A, Cherel Y, Delord K, Godø OR, Kato A, Krafft BA, Lorentsen SH, Ropert-Coudert Y, Skaret G, Varpe Ø (2016b). “Data From: At-Sea Distribution and Prey Selection of Antarctic Petrels and Commercial Krill Fisheries.” doi:10.5441/001/1.q4gn4q56.
- Gabry J, Mahr T (2024). **bayesplot: Plotting for Bayesian Models**. doi:10.32614/CRAN.package.bayesplot. R package version 1.11.1.
- Gabry J, Simpson D, Vehtari A, Betancourt M, Gelman A (2019). “Visualization in Bayesian Workflow.” *Journal of the Royal Statistical Society A*, **182**, 389–402. doi:10.1111/rssa.12378.
- Hodges JS (2013). *Richly Parameterized Linear Models: Additive, Time Series, and Spatial Models Using Random Effects*. Chapman & Hall/CRC.
- Huang Q, Cohen D, Komarzynski S, Li XM, Innominato P, Lévi F, Finkenstädt B (2018). “Hidden Markov Models for Monitoring Circadian Rhythmicity in Telemetric Activity Data.” *Journal of The Royal Society Interface*, **15**(139), 20170885. doi:10.1098/rsif.2017.0885.

- Hughes JP, Guttorp P, Charles SP (1999). “A Non-Homogeneous Hidden Markov Model For Precipitation Occurrence.” *Journal of the Royal Statistical Society C*, **48**(1), 15–30. doi:10.1111/1467-9876.00136.
- Jackson C (2011). “Multi-State Models For Panel Data: The **msm** Package For R.” *Journal of Statistical Software*, **38**(8), 1–28. doi:10.18637/jss.v038.i08.
- Jackson CH, Sharples LD, Thompson SG, Duffy SW, Couto E (2003). “Multistate Markov Models For Disease Progression with Classification Error.” *Journal of the Royal Statistical Society D*, **52**(2), 193–209. doi:10.1111/1467-9884.00351.
- James G, Witten D, Hastie T, Tibshirani R (2013). *An Introduction to Statistical Learning*, volume 112. Springer-Verlag. doi:10.1007/978-1-4614-7138-7.
- Kim CJ, Piger J, Startz R (2008). “Estimation of Markov Regime-Switching Regression Models with Endogenous Switching.” *Journal of Econometrics*, **143**(2), 263–273. doi:10.1016/j.jeconom.2007.10.002.
- Kristensen K, Nielsen A, Berg C, Skaug H, Bell B (2016). “**TMB**: Automatic Differentiation and Laplace Approximation.” *Journal of Statistical Software*, **70**(5), 1–21. doi:10.18637/jss.v070.i05.
- Langrock R, Kneib T, Glennie R, Michelot T (2017). “Markov-Switching Generalized Additive Models.” *Statistics and Computing*, **27**(1), 259–270. doi:10.1007/s11222-015-9620-3.
- Langrock R, Kneib T, Sohn A, DeRuiter SL (2015). “Nonparametric Inference in Hidden Markov Models Using P-Splines.” *Biometrics*, **71**(2), 520–528. doi:10.1111/biom.12282.
- Langrock R, Zucchini W (2011). “Hidden Markov Models with Arbitrary State Dwell-Time Distributions.” *Computational Statistics & Data Analysis*, **55**(1), 715–724. doi:10.1016/j.csda.2010.06.015.
- Lebedev S (2022). “**hmmlearn**.” GitHub Repository, URL <https://github.com/hmmlearn/hmmlearn>.
- Leos-Barajas V, Gangloff EJ, Adam T, Langrock R, Van Beest FM, Nabe-Nielsen J, Morales JM (2017a). “Multi-Scale Modeling of Animal Movement and General Behavior Data Using Hidden Markov Models with Hierarchical Structures.” *Journal of Agricultural, Biological and Environmental Statistics*, **22**, 232–248. doi:10.1007/s13253-017-0282-9.
- Leos-Barajas V, Photopoulou T, Langrock R, Patterson TA, Watanabe YY, Murgatroyd M, Papastamatiou YP (2017b). “Analysis of Animal Accelerometer Data Using Hidden Markov Models.” *Methods in Ecology and Evolution*, **8**(2), 161–173. doi:10.1111/2041-210x.12657.
- Maruotti A, Rydén T (2009). “A Semiparametric Approach to Hidden Markov Models under Longitudinal Observations.” *Statistics and Computing*, **19**(4), 381–393. doi:10.1007/s11222-008-9099-2.
- McClintock BT (2021). “Worth the Effort? A Practical Examination of Random Effects in Hidden Markov Models for Animal Telemetry Data.” *Methods in Ecology and Evolution*, **12**(8), 1475–1497. doi:10.1111/2041-210x.13619.

- McClintock BT, Langrock R, Gimenez O, Cam E, Borchers DL, Glennie R, Patterson TA (2020). “Uncovering Ecological State Dynamics with Hidden Markov Models.” *Ecology Letters*, **23**(12), 1878–1903. doi:10.1111/ele.13610.
- McClintock BT, Michelot T (2018). “**momentuHMM**: R Package for Generalized Hidden Markov Models of Animal Movement.” *Methods in Ecology and Evolution*, **9**(6), 1518–1530. doi:10.1111/2041-210x.12995.
- McKellar AE, Langrock R, Walters JR, Kesler DC (2015). “Using Mixed Hidden Markov Models to Examine Behavioral States in a Cooperatively Breeding Bird.” *Behavioral Ecology*, **26**(1), 148–157. doi:10.1093/beheco/aru171.
- Michelot T, Langrock R, Patterson TA (2016). “**moveHMM**: An R Package for the Statistical Modelling of Animal Movement Data Using Hidden Markov Models.” *Methods in Ecology and Evolution*, **7**(11), 1308–1315. doi:10.1111/2041-210x.12578.
- Miller DL (2025). “Bayesian Views of Generalized Additive Modelling.” *Methods in Ecology and Evolution*, **16**(3), 446–455. doi:10.1111/2041-210X.14498.
- Monnahan C, Kristensen K (2018). “No-U-Turn Sampling for Fast Bayesian Inference in **ADMB** and **TMB**: Introducing the **adnuts** and **tmbstan** R Packages.” *PLOS One*, **13**(5). doi:10.1371/journal.pone.0197954.
- Nash JC, Varadhan R (2011). “Unifying Optimization Algorithms to Aid Software System Users: **optimx** for R.” *Journal of Statistical Software*, **43**(9), 1–14. doi:10.18637/jss.v043.i09.
- Oelschläger L, Adam T (2023). “Detecting Bearish and Bullish Markets in Financial Time Series Using Hierarchical Hidden Markov Models.” *Statistical Modelling*, **23**(2), 107–126. doi:10.1177/1471082x211034048.
- Patterson TA, Basson M, Bravington MV, Gunn JS (2009). “Classifying Movement Behaviour in Relation to Environmental Conditions Using Hidden Markov Models.” *Journal of Animal Ecology*, **78**(6), 1113–1123. doi:10.1111/j.1365-2656.2009.01583.x.
- Pedersen EJ, Miller DL, Simpson GL, Ross N (2019). “Hierarchical Generalized Additive Models in Ecology: An Introduction with **mgcv**.” *PeerJ*, **7**, e6876. doi:10.7717/peerj.6876.
- Pohle J, Langrock R, Van Beest FM, Schmidt NM (2017). “Selecting the Number of States in Hidden Markov Models: Pragmatic Solutions Illustrated Using Animal Movement.” *Journal of Agricultural, Biological and Environmental Statistics*, **22**(3), 270–293. doi:10.1007/s13253-017-0283-8.
- Pohle J, Langrock R, Van der Schaar M, King R, Jensen FH (2021). “A Primer on Coupled State-Switching Models for Multiple Interacting Time Series.” *Statistical Modelling*, **21**(3), 264–285. doi:10.1177/1471082x20956423.
- R Core Team (2025). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. doi:10.32614/R.manuals. URL <https://www.R-project.org/>.



- Rigby RA, Stasinopoulos DM (2005). “Generalized Additive Models for Location, Scale and Shape.” *Journal of the Royal Statistical Society C*, **54**(3), 507–554. doi:[10.1111/j.1467-9876.2005.00510.x](https://doi.org/10.1111/j.1467-9876.2005.00510.x).
- Sanchez-Espigares JA, Lopez-Moreno A (2021). **MSwM**: *Fitting Markov Switching Models*. doi:[10.32614/CRAN.package.MSwM](https://doi.org/10.32614/CRAN.package.MSwM). R package version 1.5.
- Skaug HJ, Fournier DA (2006). “Automatic Approximation of the Marginal Likelihood in Non-Gaussian Hierarchical Models.” *Computational Statistics & Data Analysis*, **51**(2), 699–709. doi:[10.1016/j.csda.2006.03.005](https://doi.org/10.1016/j.csda.2006.03.005).
- Stan Development Team (2024). “**rstan**: The R Interface to Stan.” doi:[10.32614/CRAN.package.rstan](https://doi.org/10.32614/CRAN.package.rstan). R package version 2.32.6.
- Visser I, Speekenbrink M (2010). “**depmixS4**: An R Package for Hidden Markov Models.” *Journal of Statistical Software*, **36**(7), 1–21. doi:[10.18637/jss.v036.i07](https://doi.org/10.18637/jss.v036.i07).
- Viterbi A (1967). “Error Bounds for Convolutional Codes and an Asymptotically Optimum Decoding Algorithm.” *IEEE Transactions on Information Theory*, **13**(2), 260–269. doi:[10.1109/tit.1967.1054010](https://doi.org/10.1109/tit.1967.1054010).
- Wand MP (2003). “Smoothing and Mixed Models.” *Computational Statistics*, **18**, 223–249. doi:[10.1007/s001800300142](https://doi.org/10.1007/s001800300142).
- Wickham H (2016). **ggplot2**: *Elegant Graphics for Data Analysis*. Springer-Verlag. doi:[10.1007/978-0-387-98141-3](https://doi.org/10.1007/978-0-387-98141-3).
- Wickham H, Danenberg P, Csárdi G, Eugster M (2022). **roxygen2**: *In-Line Documentation for R*. doi:[10.32614/CRAN.package.roxygen2](https://doi.org/10.32614/CRAN.package.roxygen2). R package version 7.2.0.
- Wood SN (2017). *Generalized Additive Models: An Introduction with R*. 2nd edition. Chapman & Hall/CRC. doi:[10.1201/9781315370279](https://doi.org/10.1201/9781315370279).
- Zucchini W, MacDonald IL, Langrock R (2017). *Hidden Markov Models for Time Series: An Introduction Using R*. Chapman & Hall/CRC. doi:[10.1201/b20790](https://doi.org/10.1201/b20790).



## A. List of distributions

In **hmmTMB**, the observation process is modeled with state-dependent (parametric) distributions,

$$Z_t \mid \{S_t = j\} \sim \mathcal{D}(\omega_j),$$

where  $j \in \{1, 2, \dots, K\}$ . The choice of  $\mathcal{D}$  depends on the type of quantity modeled by  $Z_t$ ; e.g., a positive count might be modeled with Poisson distributions, a proportion between 0 and 1 might be modeled with beta distributions, and an unconstrained continuous variable might be modeled with normal distributions.

This is the list of distributions currently available in **hmmTMB**, with a list of parameters. It is relatively easy to add new distributions to the package, and users are invited to express interest in missing distributions, for future releases.

```
"beta": beta(shape1, shape2)
"binom": binomial(size, prob)
"cat": categorical(...)
"dir": Dirichlet(alpha1, alpha2)
"exp": exponential(rate)
"foldednorm": folded normal(mean, sd)
"gamma": gamma(shape, scale)
"gamma2": gamma2(mean, sd)
"lnorm": log-normal(meanlog, sdlog)
"mvnorm": multivariate normal(mu1, mu2, sd1, sd2, corr12)
"nbinom": negative binomial(size, prob)
"norm": normal(mean, sd)
"pois": Poisson(rate)
"t": Student's t(mean, scale)
"truncnorm": truncated normal(mean, sd, min, max)
"tweedie": Tweedie(mean, p, phi)
"vm": von Mises(mu, kappa)
"weibull": Weibull(shape, scale)
"wrpcachy": wrapped Cauchy(mu, rho)
"zibinom": zero-inflated binomial(size, prob, z)
"zigamma": zero-inflated gamma(shape, scale, z)
"zigamma2": zero-inflated gamma2(mean, sd, z)
"zinbinom": zero-inflated negative binomial(size, prob, z)
"zipois": zero-inflated Poisson(rate, z)
"ztnbinom": zero-truncated negative binomial(size, prob)
"ztpois": zero-truncated Poisson(rate)
```

The link functions used for the parameters are the following:

- log for  $\omega > 0$ ,
- logit for  $\omega \in [0, 1]$ ,
- logit of scaled parameter for  $\omega \in [a, b]$  (e.g., mean of angular distribution  $\in [-\pi, \pi]$ ),
- identity for  $\omega \in \mathbb{R}$ .

## B. Energy price analysis

In this appendix, we illustrate the implementation of Markov-switching regression in package **hmmTMB**, with the inclusion of covariates in the observation parameters. We present the analysis of a data set of energy prices, included in the R package **MSwM** (Sanchez-Espigares and Lopez-Moreno 2021). The data set includes energy prices in Spain between 2002 and 2008, as well as some potential explanatory variables (e.g., raw material prices, financial indices). Refer to `?MSwM::energy` for more detail about the data set.

```
R> library("hmmTMB")
R> data("energy", package = "MSwM")
```

### B.1. Model formulation

We consider a 2-state model, where the energy price  $Z_t$  is modeled with a normal distribution in each state,

$$Z_t \mid S_t = j \sim N(\mu_j, \sigma_j)$$

We further make the assumption that, within each state, the parameters of the normal distribution depend on the exchange rate between Euro and US dollar,  $r_t$ . The mean is modeled with a cubic spline, and the standard deviation is modeled with a cubic polynomial.

$$\begin{aligned}\mu_j^{(t)} &= \alpha_0 + f_j(r_t), \\ \log(\sigma_j^{(t)}) &= \alpha_0 + \alpha_1 r_t + \alpha_2 r_t^2 + \alpha_3 r_t^3.\end{aligned}$$

This is an example of Markov-switching generalized additive model.

### B.2. Model fitting

The hidden state model does not include covariates in this analysis so, to specify the object of class ‘**MarkovChain**’, we only need to pass the data frame (needed to create design matrices, even when there are no covariates) and the number of states.

```
R> hid <- MarkovChain$new(data = energy, n_states = 2)
```

Defining the observation model involves a little more work, as it requires a list of observation distributions, a list of initial parameters, and a list of model formulas. We choose initial parameters based on inspection of a histogram of the `Price` variable. We use syntax from the R package **mgcv** to define non-parametric formula terms.

```
R> dists <- list(Price = "norm")
R> par0 <- list(Price = list(mean = c(3, 6), sd = c(1, 1)))
R> f <- list(Price = list(mean = ~ s(EurDol, k = 10, bs = "cs"),
+   sd = ~ poly(EurDol, 3)))
R> obs <- Observation$new(data = energy, n_states = 2, dists = dists,
+   par = par0, formulas = f)
```

We combine the ‘**MarkovChain**’ and ‘**Observation**’ objects to create the model, and we fit it. This takes about 50 sec on a laptop with an Intel i7-1065G7 CPU @1.30GHz with 16Gb RAM.

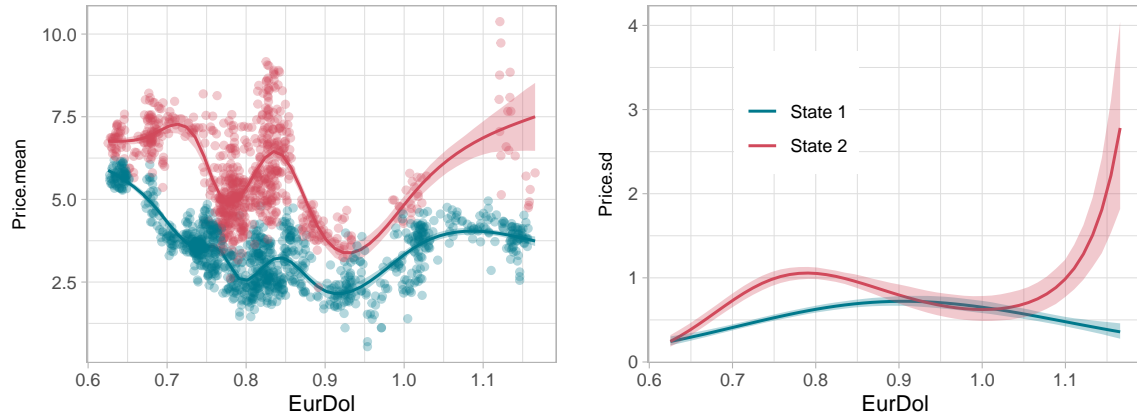


Figure 5: Results of energy price analysis. Left: relationship between mean price and euro-dollar exchange rate in each state, plotted with the data points colored by the most likely state sequence. Right: relationship between price standard deviation and euro-dollar exchange rate.

```
R> hmm <- HMM$new(hid = hid, obs = obs)
R> hmm$fit(silent = TRUE)
```

### B.3. Visualize the results

The main object of interest in this analysis is the relationship between the observation parameters and the covariate. We can visualize this with the method `HMM_plot()`, with the arguments `what = "obspar"` and `var = "EurDol"` to indicate that we want to visualize the observation parameters as functions of the Euro–US dollar exchange rate covariate. We can also use the `i` argument to output only the plot of the mean or the plot of the standard deviation. Here, we want to overlay the relationship between mean price and Euro–US dollar exchange rate with the data. As `HMM_plot()` returns a ‘`ggplot`’ object, we can add the points with the usual **ggplot2** syntax. We color them by the most likely state sequence from the Viterbi algorithm, to visualize how observations were classified by the model. We use the same method to visualize the standard deviation of price (Figure 5).

```
R> energy$viterbi <- factor(paste0("State ", hmm$viterbi()))
R> hmm$plot(what = "obspar", var = "EurDol", i = "Price.mean") +
+   geom_point(aes(x = EurDol, y = Price, fill = viterbi, col = viterbi),
+     data = energy, alpha = 0.3) +
+   theme(legend.position = "none")
R> hmm$plot(what = "obspar", var = "EurDol", i = "Price.sd") +
+   theme(legend.position.inside = c(0.3, 0.7))
```

Overall, state 1 captured lower energy prices, and state 2 higher energy prices, although the mean price in each state varied greatly with the euro-dollar exchange rate. In both states, the relationship between mean price and exchange rate was highly non-linear. There also seemed to be a clear effect of the exchange rate on the standard deviation of price, i.e., its variability.

## B.4. Predict state-dependent distributions

In many studies, it is interesting to plot the state-dependent density functions, possibly on top of a histogram of the data. This can for example be helpful to interpret the states, to assess how much they overlap, or to determine whether the estimated distributions capture the distribution of the data. In this example, the density function in each state depends on the exchange rate variable, and so we can only create the plot for a chosen value of this covariate. Here, we do this for a few different values of the exchange rate, to visualize how the distributions change.

We first create a data frame of the Euro–US dollar exchange rate values for which the distributions should be computed; here, we choose six values that roughly span the range of the variable. We pass this data frame to `HMM$predict()` to get an array of the state-dependent observation parameters, calculated for those covariate values.

```
R> EurDol <- seq(0.65, 1.15, by = 0.1)
R> newdata <- data.frame(EurDol = EurDol)
R> par <- hmm$predict(what = "obspar", newdata = newdata)
```

In preparation to create the plots, we derive the proportion of time spent in each state (as determined by the Viterbi state sequence), which will be used as a weight for each state-dependent density. We also create a grid of values of energy prices, which will be shown on the  $x$  axis of the plots.

```
R> w <- table(hmm$viterbi())/nrow(energy)
R> grid <- seq(min(energy$Price), max(energy$Price), length = 100)
```

We compute the state-dependent probability densities using `dnorm()`, by plugging in the predict parameter values, and we format them into a data frame for plotting.

```
R> pdf_ls <- lapply(1:dim(par)[3], function(i) {
+   p <- par[, , i]
+   pdf1 <- w[1] * dnorm(grid, mean = p["Price.mean", "state 1"],
+     sd = p["Price.sd", "state 1"])
+   pdf2 <- w[2] * dnorm(grid, mean = p["Price.mean", "state 2"],
+     sd = p["Price.sd", "state 2"])
+   res <- data.frame(price = grid, pdf = c(pdf1, pdf2),
+     state = factor(rep(1:2, each = length(grid))),
+     eurdol = paste0("EurDol = ", EurDol[i]))
+   return(res)
+ })
R> pdf_df <- do.call(rbind, pdf_ls)
```

Finally, we plot histograms of the observed energy prices, and add the state-dependent normal density lines, for the different values of exchange rate. The distributions vary widely for different values of the covariate (Figure 6).

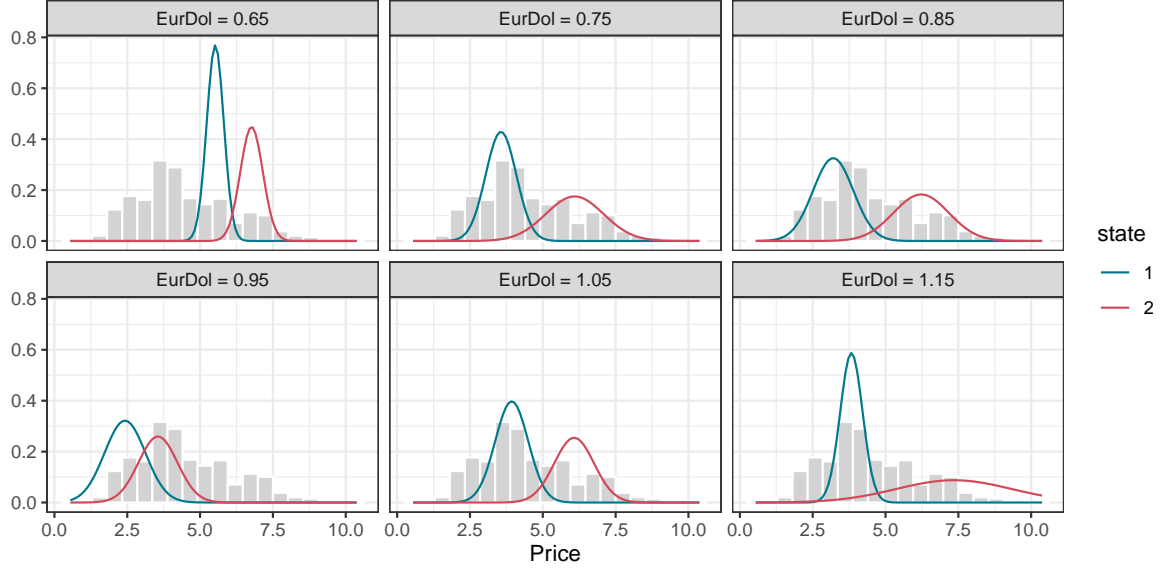


Figure 6: State-dependent distributions of energy prices, computed for various values of the euro-dollar exchange rate, shown over the histogram of observed values.

```
R> ggplot(pdf_df, aes(price, pdf)) +
+   facet_wrap("eurdol") +
+   geom_histogram(aes(x = Price, y=..density..), bins = 20,
+     col = "white", bg = "lightgrey", data = energy) +
+   geom_line(aes(col = state)) +
+   theme_bw() +
+   labs(x = "Price", y = NULL) +
+   scale_color_manual(values = hmmTMB:::hmmTMB_cols)
```

### C. Bayesian analysis of human activity data

We analyze a subset of the data described by Huang, Cohen, Komarzynski, Li, Innominato, Lévi, and Finkenstädt (2018), accessible on GitHub at [github.com/huang1010/hmms](https://github.com/huang1010/hmms). In that study, accelerometers were attached to human subjects to measure their activity through time. Raw acceleration was processed to obtain continuous variable summarizing physical activity at a 5-min resolution. We consider data from two subjects, named “S9” and “S20”, comprising a total of 2317 observations over around 8 days.

The aim of the study was to investigate circadian rhythms, i.e., daily patterns in activity. Similarly to Huang *et al.* (2018), we included time of day as a covariate on the transition probabilities but, unlike them, we used a non-parametric approach. Specifically, we modeled the relationship between transition probabilities and time of day using a cyclic spline (“cc” in **mgcv**), with period set to 24 hours.

We follow a Bayesian approach for this analysis, i.e., we perform posterior inference using `HMM$fit_stan()`.




Figure 7: Activity data from two patients, taken from [Huang \*et al.\* \(2018\)](#).

### C.1. Prepare data

We load the data from GitHub, save them in the format expected by **hmmTMB**, and add a column for time of day (numeric between 0 and 24). Plotting the time series of activity measurements suggests that there are some clearly distinct phases of low and high activity (Figure 7). We will use a 2-state HMM to try to identify those phases.

```
R> path <- "https://raw.githubusercontent.com/huang1010/HMMS/master/"
R> data1 <- read.csv(url(paste0(path, "S9.csv")))[-3]
R> data2 <- read.csv(url(paste0(path, "S20.csv")))[-3]
R> data <- rbind(cbind(ID = "S9", data1),
+   cbind(ID = "S20", data2))
R> data$time <- as.POSIXct(data$time)
R> data$tod <- as.numeric(format(data$time, "%H")) +
+   as.numeric(format(data$time, "%M"))/60
R> theme_set(theme_bw())
R> ggplot(data, aes(time, activity)) +
+   facet_wrap("ID", scales = "free_x") +
+   geom_line()
```

### C.2. Model specification

We use a cyclic spline for time of day to capture its effect on the transition probabilities. We use `initial_state = "stationary"` to indicate that the initial distribution of the Markov chain should be fixed to the stationary distribution of the first transition probability matrix, rather than estimated. This is to avoid mixing problems arising from poor identifiability of the initial distribution.

For the observation model, we use the "zigamma2" distribution, i.e., the zero-inflated gamma distribution specified in terms of mean and standard deviation (rather than shape and scale). We need the zero-inflated distribution because a small number of activity values are exactly zero, and those cannot be modeled with a gamma distribution. As a result, an additional parameter `z` is required, to measure the proportion of values that are exactly zero in each

state. We choose initial values based on the expectation that state 1 will capture active behavior, and state 2 will capture less active periods.

In this example, we fix the zero inflation parameter for state 1 to zero, i.e., we assume that state 2 will capture all observations that are exactly zero. We do this because initial experiments revealed that the zero inflation parameter for state 1 was estimated to be virtually zero, and mixing was poor.

```
R> hid <- MarkovChain$new(data = data, n_states = 2,
+   initial_state = "stationary", formula = ~ s(tod, k = 5, bs = "cc"))
R> dists <- list(activity = "zigamma2")
R> par0 <- list(activity = list(mean = c(20, 150), sd = c(20, 40),
+   z = c(0.1, 0)))
R> obs <- Observation$new(data = data, dists = dists,
+   n_states = 2, par = par0)
R> fixpar <- list(obs = c("activity.z.state2.(Intercept)" = NA))
R> hmm <- HMM$new(obs = obs, hid = hid, fixpar = fixpar)
```

We also specify prior distributions for some of the parameters using `HMM$set_priors()`. We pass a list with elements `coeff_fe_obs` and `coeff_fe_hid`, for the fixed effect parameters of the observation model and the hidden state model, respectively. For the observation model, we provide a matrix with six rows, corresponding to intercepts for the mean in state 1, mean in state 2, standard deviation in state 1, standard deviation in state 2, zero mass in state 1, and zero mass in state 2. The matrix has one column for the mean of the normal prior distribution, and one column for its standard deviation. We use diffuse priors centered on our initial guesses (with standard deviation 10), except for the zero mass in state 2, which is fixed to zero and not estimated. For the hidden state process, we specify a matrix with one row each for the intercepts of  $P(S_{t+1} = 2 \mid S_t = i)$  and  $P(S_{t+1} = 1 \mid S_t = 2)$ , and we choose these to represent our expectation that the off-diagonal elements of the transition probability matrix will be small. (This is common because the state process tends to display persistence.) We could also specify priors for the smoothness parameters of the penalized splines, but we leave them as NA here (i.e., improper flat priors are used).

```
R> prior_obs <- matrix(c(
+   log(20), 10,
+   log(150), 10,
+   log(20), 10,
+   log(40), 10,
+   qlogis(0.1), 10,
+   NA, NA),
+   ncol = 2, byrow = TRUE)
R> prior_hid <- matrix(c(
+   qlogis(0.05), 10,
+   qlogis(0.05), 10),
+   ncol = 2, byrow = TRUE)
R> hmm$set_priors(list(coeff_fe_obs = prior_obs,
+   coeff_fe_hid = prior_hid))
R> hmm$priors()
```



```

$coeff_fe_obs
              mean sd
activity.mean.state1.(Intercept)  2.995732 10
activity.mean.state2.(Intercept)  5.010635 10
activity.sd.state1.(Intercept)    2.995732 10
activity.sd.state2.(Intercept)    3.688879 10
activity.z.state1.(Intercept)     -2.197225 10
activity.z.state2.(Intercept)           NA NA

$coeff_fe_hid
              mean sd
S1>S2.(Intercept) -2.944439 10
S2>S1.(Intercept) -2.944439 10

$log_lambda_obs
              mean sd

$log_lambda_hid
              mean sd
S1>S2.s(tod)    NA NA
S2>S1.s(tod)    NA NA

```

### C.3. Model fitting using Stan

We run Hamiltonian Monte Carlo using **Stan**, with the function `HMM$fit_stan()`. We need to specify the number of chains and the number of iterations of each chain. Here, we run 1000 iterations for one chain; in practice, this choice should be made based on criteria such as effective sample sizes. This takes around 5 min on a laptop with an Intel i7-1065G7 CPU @1.30GHz with 16Gb RAM.

```

R> n_chain <- 1
R> n_iter <- 1000
R> hmm$fit_stan(chains = n_chain, iter = n_iter, seed = 9745)

```

After model fitting, `HMM$iters()` returns a data frame of posterior samples, which can then be used to create trace plots or density plots, for example (Figure 8). By default, it outputs posterior samples for the HMM parameters (transition probabilities and observation parameters), but posterior samples for the linear predictor parameters (fixed effects, random effects, smoothness parameters) can be obtained using the argument `type = "raw"`. These data frames can also directly be passed to functions from the **bayesplot** package for visualization (Gabry, Simpson, Vehtari, Betancourt, and Gelman 2019; Gabry and Mahr 2024), e.g., `bayesplot::mcmc_pairs()` or `bayesplot::mcmc_areas()`.

```

R> par_df <- as.data.frame.table(hmm$iters()[,c(1:4)])
R> colnames(par_df) <- c("iter", "par", "value")

```




Figure 8: Histograms of posterior samples for observation parameters in activity data analysis.

```
R> ggplot(par_df, aes(x = value)) +
+   geom_histogram(bins = 30, col = "grey", fill = "lightgrey") +
+   facet_wrap("par", scales = "free", ncol = 2) +
+   labs(x = NULL)
```

We can also access the `stanfit` object using `HMM$out_stan()`, which can be printed, or passed to functions such as `rstan::stan_trace()` or `rstan::stan_dens()`. Printing the output gives posterior quantiles, as well as information about effective sample size and convergence. Here, `Rhat` is 1 for all parameters, indicating convergence of the sampler to the posterior distribution (see `rstan` and `Stan` documentation).

```
R> hmm$out_stan()
```

Inference for Stan model: `hmmTMB`.

1 chains, each with iter=1000; warmup=500; thin=1;

post-warmup draws per chain=500, total post-warmup draws=500.

	mean	se_mean	sd	2.5%	25%
<code>coeff_fe_obs[1]</code>	2.99	0.00	0.05	2.89	2.96
<code>coeff_fe_obs[2]</code>	4.89	0.00	0.01	4.88	4.89
<code>coeff_fe_obs[3]</code>	3.14	0.00	0.05	3.04	3.11
<code>coeff_fe_obs[4]</code>	3.49	0.00	0.03	3.44	3.47
<code>coeff_fe_obs[5]</code>	-2.30	0.00	0.10	-2.50	-2.37
<code>coeff_fe_hid[1]</code>	-3.19	0.01	0.19	-3.58	-3.32
<code>coeff_fe_hid[2]</code>	-3.46	0.01	0.27	-4.04	-3.61
<code>log_lambda_hid[1]</code>	0.62	0.05	1.15	-2.01	0.04
<code>log_lambda_hid[2]</code>	0.75	0.06	1.16	-1.59	0.05
<code>coeff_re_hid[1]</code>	-1.28	0.02	0.44	-2.22	-1.57

coeff_re_hid[2]	0.68	0.02	0.42	-0.13	0.39
coeff_re_hid[3]	1.25	0.01	0.34	0.62	1.03
coeff_re_hid[4]	-0.58	0.03	0.58	-1.97	-0.87
coeff_re_hid[5]	-1.42	0.02	0.43	-2.24	-1.71
coeff_re_hid[6]	-0.55	0.02	0.42	-1.38	-0.84
lp__	-10379.95	0.20	3.00	-10387.63	-10381.65
	50%	75%	97.5%	n_eff	Rhat
coeff_fe_obs[1]	2.99	3.02	3.07	289	1.00
coeff_fe_obs[2]	4.89	4.90	4.91	1059	1.00
coeff_fe_obs[3]	3.14	3.18	3.23	260	1.00
coeff_fe_obs[4]	3.49	3.51	3.55	1096	1.00
coeff_fe_obs[5]	-2.30	-2.23	-2.10	772	1.00
coeff_fe_hid[1]	-3.19	-3.05	-2.84	646	1.00
coeff_fe_hid[2]	-3.44	-3.29	-2.99	468	1.00
log_lambda_hid[1]	0.76	1.34	2.44	469	1.01
log_lambda_hid[2]	0.78	1.58	2.88	359	1.00
coeff_re_hid[1]	-1.26	-0.97	-0.41	551	1.00
coeff_re_hid[2]	0.67	0.98	1.54	771	1.00
coeff_re_hid[3]	1.24	1.49	1.92	616	1.00
coeff_re_hid[4]	-0.54	-0.20	0.40	399	1.00
coeff_re_hid[5]	-1.40	-1.12	-0.68	483	1.00
coeff_re_hid[6]	-0.52	-0.24	0.21	555	1.00
lp__	-10379.68	-10377.81	-10375.22	222	1.00

Samples were drawn using NUTS(diag\_e) at Thu May 1 10:42:07 2025.

For each parameter, n\_eff is a crude measure of effective sample size, and Rhat is the potential scale reduction factor on split chains (at convergence, Rhat=1).

#### C.4. Observation distributions

It is often useful to visualize the state-dependent observation distributions, possibly combined with a histogram of the data, to help with state interpretation and assess goodness-of-fit. We can generate posterior samples of the distributions, and plot those to also visualize the associated uncertainty (Figure 9). The plot shows that the two states were clearly distinct: one of them captures high activity, and the other low activity. There seems to be some lack of fit for high activity values (i.e., the density lines don't match the distribution of observed data), and a 3-state model might be able to provide the additional flexibility needed here.

```
R> w <- table(hmm$viterbi())/nrow(data)
R> ind_post <- sort(sample(1:(n_iter/2), size = 100))
R> dens_df <- data.frame()
R> activity <- seq(0.4, 200, length = 100)
R> for(i in ind_post) {
+   par <- hmm$iters()[i,1:6]
+   dens1 <- obs$dists()$activity$pdf(x = activity, mean = par[1],
+   sd = par[3], z = par[5])
```




Figure 9: State-dependent distributions and histogram of activity data. The lines are 100 posterior samples.

```
+ dens2 <- obs$dists()$activity$pdf()(x = activity, mean = par[2],
+   sd = par[4], z = par[6])
+ dens <- data.frame(state = paste0("state ", rep(1:2, each = 100)),
+   dens = c(w[1] * dens1, w[2] * dens2))
+ dens$group <- paste0("iter ", i, " - ", dens$state)
+ dens_df <- rbind(dens_df, dens)
+ }
R> dens_df$activity <- activity
R> ggplot(dens_df, aes(activity, dens)) +
+   geom_histogram(aes(y = ..density..), data = data, fill = "lightgrey",
+   col = "grey", breaks = seq(0, 200, by = 10)) +
+   geom_line(aes(col = state, group = group), size = 0.1, alpha = 0.5) +
+   scale_color_manual(values = hmmTMB::hmmTMB_cols, name = NULL) +
+   guides(color = guide_legend(override.aes = list(size = 0.5, alpha = 1)))
```

### C.5. Covariate effects

The aim of the study was to estimate daily variations in activity. We can use the function `HMM$predict()` to compute the stationary state probabilities over a grid of values of time of day. The stationary state probabilities are the stationary distribution of the corresponding transition probabilities, and they are meant to roughly measure the proportion of time spent in each state, for a given covariate value. Like above, we generate curves from 100 randomly selected posterior samples, so that the plot captures the uncertainty in the estimated relationship. To do this, we use `hmm$update_par(iter = i)` to update the parameter values stored in the model (and used in `HMM$predict()`) to the  $i$ -th posterior sample, and we loop over  $i$ . Although there is high variation between the posterior samples, they almost all indicate that it was much more likely to be in the high activity state during the day, and in the low activity state at night (Figure 10).

```
R> ind_post <- sort(sample(1:(n_iter/2), size = 100))
```




Figure 10: Estimated stationary state probabilities as functions of time of day, in the human activity analysis. The lines are 100 posterior samples.

```
R> probs_df <- data.frame()
R> newdata <- data.frame(tod = seq(0, 24, length = 100))
R> for(i in ind_post) {
+   hmm$update_par(iter = i)
+   probs <- data.frame(state = rep(paste0("state ", 1:2), each = 100),
+     prob = as.vector(hmm$predict(what = "delta", newdata = newdata)))
+   probs$group <- paste0("iter ", i, " - ", probs$state)
+   probs_df <- rbind(probs_df, probs)
+ }
R> probs_df$tod <- newdata$tod
R> ggplot(probs_df, aes(tod, prob, group = group, col = state)) +
+   geom_line(size = 0.1, alpha = 0.5) +
+   scale_x_continuous(breaks = seq(0, 24, by = 4)) +
+   labs(x = "time of day", y = "stationary state probabilities",
+     col = NULL) +
+   scale_color_manual(values = hmmTMB::hmmTMB_cols) +
+   guides(color = guide_legend(override.aes = list(size = 0.5, alpha = 1)))
```

## D. Simulation study

In this appendix, we investigate the ability of **hmmTMB** to recover relationships between parameters and covariates in three scenarios: (1) non-homogeneous hidden Markov model with non-linear covariate effects on transition probabilities, (2) Markov-switching generalized additive model with non-linear covariate effects on observation parameters, and (3) mixed hidden Markov model with random effect on transition probabilities.

### D.1. Non-homogeneous hidden Markov model

We first considered a 2-state HMM where the transition probabilities are functions of a time-

varying continuous covariate  $x_{1t} \in [-1, 1]$ . The true relationship used to simulate data was

$$\begin{aligned}\text{logit}(\gamma_{12}^{(t)}) &= -3 + 3x_{1t}^2 \\ \text{logit}(\gamma_{21}^{(t)}) &= -2 + \sin(\pi x_{1t}).\end{aligned}$$

The observations were simulated from state-dependent normal distributions,

$$Z_t \mid \{S_t = j\} \sim N(\mu_j, \sigma_j^2)$$

where  $\mu_1 = -5$ ,  $\mu_2 = 5$ , and  $\sigma_1 = \sigma_2 = 1$ .

We fitted a 2-state HMM where the transition probabilities were modeled with non-parametric splines. We used the "cs" basis (cubic spline) for  $\gamma_{12}^{(t)}$  and the "cc" basis (cyclical cubic spline) for  $\gamma_{21}^{(t)}$ . We assumed normal observation distributions, for which the parameters were estimated during model fitting.

We repeated the procedure 200 times; at each iteration,

1. Simulate  $n = 5000$  values of the covariate using a reflected Gaussian random walk on  $(-1, 1)$ .
2. Derive one transition probability matrix for each time step  $t = 1, \dots, n$ , based on the assumed relationships.
3. Simulate state sequence based on transition matrices, and then simulate observations  $z_1, \dots, z_n$  from state-dependent distributions.
4. Fit HMM to the observations  $z_1, \dots, z_n$ .
5. Predict the transition probabilities over a grid of values of  $x_1$ .

The estimated relationships between the transition probabilities and the covariate are shown in Figure 11, together with the true relationships. For each transition probability, all estimates captured the shape of the true function well.

## D.2. Markov-switching regression

Next, we considered a 2-state HMM with covariate effects on the observation parameters. Specifically, we modeled observations with state-dependent Poisson distributions,

$$Z_t \mid \{S_t = j\} \sim \text{Poisson}(\lambda_j^{(t)}).$$

The rate parameter was chosen to be constant in state 2 (i.e.,  $\lambda_2^{(t)} = \lambda_2$ ), and to be a function of a continuous time-varying covariate  $x_{1t} \in [-1, 1]$  in state 1:

$$\lambda_1^{(t)} = \begin{cases} \exp(1 + \sin(2\pi x_{1t})) & \text{if } -0.5 \leq x_{1t} \leq 0.5, \\ \exp(1) & \text{otherwise.} \end{cases}$$

We tried to recover this relationship using a Markov-switching generalized additive model, i.e., by modeling the (log) rate with a cubic spline ("cs" in **mgcv** syntax).

We repeated the procedure 200 times; at each iteration,




Figure 11: Results of non-homogeneous HMM simulation. The red lines are the true relationships between transition probabilities and the covariate  $x_1$  used for simulating, and the black lines are estimates.

1. Simulate  $n = 2000$  values of the covariate using a reflected Gaussian random walk on  $(-1, 1)$ .
2. Derive the time-varying Poisson rates based on the function described above.
3. Simulate state sequence based on the transition probabilities  $\gamma_{12} = \gamma_{21} = 0.1$ , and then simulate observations  $z_1, \dots, z_n$  from state-dependent distributions with time-varying parameters.
4. Fit HMM to the observations  $z_1, \dots, z_n$ .
5. Predict the Poisson rate in state 1 over a grid of values of  $x_1$ .

The estimated parameters  $\lambda_1$  are shown as functions of  $x_1$  in Figure 12. The true non-linear function was captured well by the estimated splines.

### D.3. Mixed hidden Markov model

We simulated data from a 2-state HMM where the transition probabilities were themselves random, to check whether their distribution could be estimated using **hmmTMB**. We consider individual-specific transition probabilities (where the individual could be a subject in the study, or an animal, or any relevant grouping) where, for individual  $m \in \{1, \dots, M\}$ ,

$$\begin{aligned} \text{logit}(\gamma_{12}^{(m)}) &= -2.5 + \beta_{12}^{(m)}, & \beta_{12}^{(m)} &\sim N(0, 1^2) \\ \text{logit}(\gamma_{21}^{(m)}) &= -2.5 + \beta_{21}^{(m)}, & \beta_{21}^{(m)} &\sim N(0, 0.5^2). \end{aligned}$$

In this scenario, the aim was to estimate the variance of the distributions of the random effects  $\beta_{12}^{(m)}$  and  $\beta_{21}^{(m)}$ .






Figure 12: Results of Markov-switching regression simulation. The red line is the true relationship between the rate parameter  $\lambda_1$  and the covariate  $x_1$  used in the simulations, and the black lines are estimates.

Observations were simulated from state-dependent gamma distributions,

$$Z_t \mid \{S_t = j\} \sim \text{gamma}(\mu_j, \sigma_j^2),$$

with means  $\mu_1 = 3$  and  $\mu_2 = 15$ , and standard deviations  $\sigma_1 = 2$  and  $\sigma_2 = 5$ .

We fitted an HMM where the transition probabilities included an i.i.d. normal random effect for the individual, using `s(ID, bs = "re")`.

We repeated the procedure 200 times; at each iteration,

1. Randomly generate 20 transition probability matrices (one for each of 20 individuals), based on the true distribution described above.
2. For each individual, simulate a state sequence of length 500 based on the individual-specific transition probability matrix, and then simulate observations from the state-dependent gamma distributions.
3. Fit HMM to the observations from all individuals.
4. Save the estimated standard deviation of the random effects.

The distributions of the estimated random effect standard deviations are shown in Figure 13. The true standard deviations were recovered successfully in most analyses.




Figure 13: Results of mixed HMM simulation. The vertical dotted lines show the true values of the standard deviation of random effects used in the simulations, and the histograms show the distribution of 200 estimates.

### Affiliation:

Théo Michelot  
 Department of Mathematics and Statistics  
 Dalhousie University  
 Halifax NS, B3H 1Z9, Canada  
 E-mail: [theo.michelot@dal.ca](mailto:theo.michelot@dal.ca)  
 URL: <https://theomichelot.github.io/>