

Hidden Markov model for discrete circular–linear wind data time series

Original

Hidden Markov model for discrete circular–linear wind data time series / Mastrantonio, Gianluca; Calise, Gianfranco. - In: JOURNAL OF STATISTICAL COMPUTATION AND SIMULATION. - ISSN 0094-9655. - 86:13(2016), pp. 2611-2624. [10.1080/00949655.2016.1142544]

Availability:

This version is available at: 11583/2664912 since: 2020-01-30T10:10:17Z

Publisher:

Taylor and Francis Ltd.

Published

DOI:10.1080/00949655.2016.1142544

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Hidden Markov model for discrete circular-linear wind data time series

Gianluca Mastrantonio¹ and Gianfranco Calise²

¹ Roma Tre University, Via Silvio D'Amico 77, Rome, 00145, Italy

²Department of Earth Science, University of Rome "Sapienza"

Abstract

In this work, we deal with a bivariate time series of wind speed and direction. Our observed data have peculiar features, such as informative missing values, non-reliable measures under a specific condition and interval-censored data, that we take into account in the model specification.

We analyze the time series with a non-parametric Bayesian hidden Markov model, introducing a new emission distribution based on the invariant wrapped Poisson, the Poisson and the hurdle density, suitable to model our data. The model is estimated on simulated datasets and on the real data example that motivated this work.

Keywords: Invariant wrapped Poisson; hurdle model; discrete circular variable; non-parametric Bayesian; Dirichlet process.

1 Introduction

The analysis of time series of meteorological data has an increasing interest in many fields. Their analysis is interesting in order to have a better understanding of the atmospheric phenomena, to determine the climate of a geographical data or to predict the occurrence of extreme events. In this work, motivated by our real data example, we are interested in the modelling of a bivariate time series of wind speed and direction.

The analysis of wind time series have been carried out by means of different approaches, for example mixture type models ([1]), harmonic analyses ([2]), ARMA-GARCH models ([3]) or stochastic differential equations ([4]). Among the others, the hidden Markov model (HMM), a class of mixture model for time series, proved to be well suited to model wind data, see for example [5], [6], [1], [7], [8] or [9]. In the HMM, the data-generative distribution is expressed as a mixture, with components that belong, generally, to an "easy tractable" distribution (called *emission distribution* or *regime-specific density*). The belonging to a component of the mixture, called also *regime* or *state*, depends on an underlying and unobservable discrete valued Markov process.

The wind direction is a circular variable, i.e. a variable that represents an angle or a point over the unit circle, while the speed is a linear (or inline) one. Hence, if a time series of wind direction and speed is modelled through an HMM, the emission distribution must be defined over a mixed circular-linear domain. Generally conditional independence between the circular and linear variables is assumed, but exceptions exist (see for example [7, 9]).

Although HMMs for wind data have been proposed in the literature, here, due to particular features of our data, a direct application of the models previously proposed is not possible. The measured linear and circular variables are interval-censored ([10]) and they are recorded as discrete variables; more precisely the wind speed is recorded as integer knots and the direction is measured on a discrete scale with 36 equally spaced values over the unit circle. Furthermore, the direction is strictly related

to the wind speed as when the latter is too small, the direction is not recorded and, due to the low instrument sensitivity, a recorded value of 0 or 1 wind speed is not reliable.

To take into account all these features, we introduce a new bivariate distribution to be used as the regime-specific density in the HMM. The model is estimated under a non-parametric Bayesian framework, using the hierarchical Dirichlet process-HMM (HDP-HMM) of [11] and the modification introduced by [12], namely the sticky HDP-HMM (sHPD-HMM). The sHPD-HMM allows us to estimate, as a model parameter, the number of regimes occupied by the time series. The model is applied to simulated and real data examples.

The paper is organized as follows. In Section 2 we describe the motivating example while in Section 3 we introduce the circular-linear distribution by first formalizing the marginal distribution of the linear component, Section 3.1, and then the one of the circular, conditioned to the linear one, Section 3.2. Section 3.3 is devoted to the model specification while Section 3.4 gives implementation details. Section 4 contains the simulated examples (Section 4.1) and the real data one (Section 4.2). The paper ends with a discussion (Section 5).

2 Motivating example

In last decades the Italian coast suffers an intensification of erosion, in particular in South Italy, along the coastline on the Tyrrhenian side, with manmade coastal structures. This is obviously a result of both natural and anthropogenic causes. The study of wind time series, recorded along the coast, is of a great importance since the wind generates waves, a significant factor of coastal geomorphology.

The data are recorded with an anemometer, located on the rocky cape of Capo Palinuro, in the town of Centola, province of Salerno, South Italy. The meteorological station of Capo Palinuro is one of the coastal stations managed by the Meteorological Service of the Military Italian Air Force. The instrument is placed away from obstacles, 10 meters above ground. The data are provided by the National Center of Aeronautical Meteorology and Climatology (C.N.M.C.A.), special office of the Meteorological Service of the Italian Air Force.

In this work we focus on the year 2010, that is a particular one. While the station, for decades, has recorded winds blowing between North and North-West or South and South-East octants, in 2010 the prevailing winds (the directions with the higher frequency) are the ones blowing from North-East and South-West, which is almost the opposite of the general trend. Then, it is important to understand and characterize in a greater details, i.e. through an HMM, the wind distribution of this year.


2.1 Data description

The linear variable is interval-censored, i.e. it is recorded with an error of 1 knot, and let Y^* be the recorded wind speed, if $Y^* = c$, with $c \in \mathbb{Z}^+$, the real (continuous) variable has value $\geq \min\{c-0.5, 0\}$ but $< c+0.5$. Further uncertainty is added to the values below 2 knots as the instrument sensitivity does not allow reliable recordings.


Also the wind direction is interval-censored and it is measured in degrees on a discrete scale with 36 distinct values. If the recorded circular variable, X , assumes value c , with $c \in \{\frac{2\pi}{36}j\}_{j=0}^{35}$, the real direction has value $\geq c - \frac{1}{2}\frac{2\pi}{36}$ but $< c + \frac{1}{2}\frac{2\pi}{36}$. If the (continuous) wind speed is too low, i.e. close to 0, the instrument could not be able to record X and we have a missing recording. This missing gives information about the non-reliable linear recording, i.e. they are *informative* and *non-ignorable*, ([13]) and then the process that generates the missing observation must be taken into account in the model specification.

Due to the instrument malfunction, we have 2920 observations with, respectively, 3 and 2 non-informative missing in the circular and linear variables. y_t^* assumes value 0 494 times while x_t is equal to $\{\emptyset\}$ 328 times. Figure 1 shows barplots of the circular and linear data.

¹The direction recorded by the instrument assumes values in $(0, 10, 20, \dots, 350)$ but we rescale it to $[0, 2\pi)$.



(a) Wind speed



(b) Wind direction

Figure 1: Linear (a) and circular (b) barplot of the variable used in the real data example.

3 The model

Before the formalization of the HMM (Section 3.3), we introduce the circular-linear distribution we use in the model specification.

3.1 The linear distribution

We code the information carried by the recorded linear variable (Y^*) into two new variables. One representing the “true” discrete wind speed with support \mathbb{Z}^+ , Y , and a Bernoulli variable W , that assumes value 0 when $Y^* < 2$ and 1 otherwise. The marginal distribution of Y is setted to be Poisson with parameter λ_y : $Y \sim P(\lambda)$. Between Y and Y^* the following relation exists:

$$\begin{cases} Y = Y^* & \text{if } Y^* \geq 2, \\ Y \in \{0, 1\} & \text{if } Y^* < 2. \end{cases}$$

Let W be a binary random variable such that

$$\begin{aligned} W &= 1 \text{ if } Y^* \geq 2, \\ W &= 0 \text{ if } Y^* < 2, \end{aligned}$$

with

$$W \sim \text{Bern}(1 - e^{-\lambda_y}(1 + \lambda_y)), \quad (1)$$

where $\text{Bern}(\cdot)$ indicates the Bernoulli distribution. We define the distribution of $Y|W, \lambda_y$ as follows:

$$P(y|w = 1, \lambda_y) = \frac{\lambda_y^y e^{-\lambda_y}}{y!} \frac{1}{1 - e^{-\lambda_y}(1 + \lambda_y)} I(y \in \mathbb{Z}^+ \setminus \{0, 1\}) \quad (2)$$

and

$$P(y|w = 0, \lambda_y) = \left(\frac{\lambda_y e^{-\lambda_y}}{e^{-\lambda_y}(1 + \lambda_y)} \right)^y \left(1 - \frac{\lambda_y e^{-\lambda_y}}{e^{-\lambda_y}(1 + \lambda_y)} \right)^{1-y} I(y \in \{0, 1\}) \quad (3)$$

where $I(y \in A)$ is the indicator function.

3.2 The circular distribution

To model properly the non-ignorable missing value of X , we specify the domain of the circular variable as a mixed one, composed of a discrete circle $\mathbb{D} = \{\frac{2\pi}{36}j\}_{j=0}^{35}$, with 36 equally spaced values, and the empty set, $\{\emptyset\}$. The direction assumes value $\{\emptyset\}$ if the measure is a missing non-ignorable and then the distribution of the circular variable is an hurdle one. The hurdle distribution, first introduced by [14], assumes that the observation, in this case the wind direction, comes from two data-generative processes. With probability ν^* the observation belongs to the hurdle, i.e. the empty set, otherwise it belongs to the other portion of the mixed domain, i.e. \mathbb{D} , and a distribution over \mathbb{D} must be chosen.

Since the circular variable can assume value $\{\emptyset\}$ only if the discrete wind speed is 0, the probability ν^* must depend on Y . Indeed $(X = \{\emptyset\}) \Rightarrow (Y = 0)$ but it is not true that $(Y = 0) \Rightarrow (X = \{\emptyset\})$, as X may be measured when a light wind is blowing (below 1 knot). We specify the hurdle probability as $\nu^* = \nu I(Y = 0)$ and, as required, $\nu^* \neq 0$ only if $Y = 0$.

In the literature several distributions for discrete circular variables have been proposed, see for example the wrapped Binomial ([15]) or the wrapped Weibull ([16]). [17] proved that for most of them, the inference they allow strongly depends on the choice of the reference system origin and orientation, while a ‘‘proper’’ distribution must be independent on these choices, for details see [17]. As distribution over \mathbb{D} , we choose the invariant wrapped Poisson (IWP), proposed by [17], since it is a proper and flexible distribution for discrete circular data that is easy to implement in a Bayesian framework.

The density of X is

$$P(x|y, \lambda_x, \eta, \xi) = (\nu^*)^{I(x=\emptyset)} ((1 - \nu^*)P(x|x \in \mathbb{D}, \lambda_x, \eta, \xi))^{I(x \in \mathbb{D})}.$$

where $P(x|x \in \mathbb{D}, \lambda_x, \eta, \xi)$ is the pmf of the IWP:

$$P(x|x \in \mathbb{D}, \lambda_x, \eta, \xi) = \sum_{k=0}^{\infty} \frac{\lambda^{(\eta\theta - \xi) \bmod (2\pi)36/(2\pi) + k36} e^{-\lambda}}{((\eta\theta - \xi) \bmod (2\pi)36/(2\pi) + k36)!}, \quad \xi \in \mathbb{D}, \eta \in \{-1, 1\}. \quad (4)$$

(4) is obtained by wrapping the density of a linear transformation of a Poisson distributed variable. More in details, we first obtain the density of $Q^* = \eta(Q\frac{2\pi}{36} + \xi)$, where $Q \sim P(\lambda_x)$, that is

$$\frac{\lambda^{(\eta q^* - \xi)36/(2\pi)} e^{-\lambda}}{((\eta q^* - \xi)36/(2\pi))!}, \quad (5)$$

then we find the distribution of the associated wrapped variable $X = Q^* \bmod 2\pi$, that is (4). Note that the transformation $X = Q^* \bmod 2\pi$ wraps the linear variable Q^* , and its density, around the discrete circle \mathbb{D} . k , i.e. *winding number*, identify in (4) the k^{th} portion of the domain of Q^* that is wrapped around the circle.

We can compute in closed form the directional mean and concentration of the IWP ([17]), i.e. the circular counterparts of linear mean and concentration:

$$\mu = \eta\xi + \lambda \sin\left(\eta\frac{2\pi}{36}\right), \quad (6)$$

$$c = e^{-\lambda(1 - \cos(\frac{2\pi}{36}))}. \quad (7)$$

Equations (6) and (7), as the mean and concentration for the inline variables, are useful statistics to describe the circular variable behaviour.

It is not easy to work directly with equation (4), since it involves an infinite sum. When a wrapped distribution is used ([18, 19, 20]), a standard approach is to introduce the latent random variable K , and to work with the joint density of (X, K) that is the summand in (4) and does not require the infinite sum evaluation:

$$P(x, k|x \in \mathbb{D}, \lambda_x, \eta, \xi) = \frac{\lambda^{(\eta\theta - \xi) \bmod (2\pi)36/(2\pi) + kl} e^{-\lambda}}{((\eta\theta - \xi) \bmod (2\pi)36/(2\pi) + kl)!}.$$

The winding number is needed also to define an efficient sampling scheme for the IWP parameters, see [17] and Section 3.4.

The augmented hurdle type density is

$$P(x, k|y, \lambda_x, \eta, \xi) = (\nu^*)^{I(x=0)} ((1 - \nu^*)P(x, k|x \in \mathbb{D}, \lambda_x, \eta, \xi))^{I(x \in \mathbb{D})}. \quad (8)$$

We write $X, K, Y, W|\lambda_x, \eta, \xi, \nu, \lambda_y \sim HiwpP_{\mathbb{D}}(\lambda_x, \eta, \xi, \nu, \lambda_y)$ if $W|\lambda_y$ is distributed as (1), $Y|W, \lambda_y$ as in (2) and (3) and $X, K|Y, \lambda_x, \eta, \xi$ as in (8).

3.3 The HMM specification

Following the formalization introduced in Sections 3.1 and 3.2, we proceed by modeling the four-variate time series $\{\mathbf{x}, \mathbf{k}, \mathbf{y}, \mathbf{w}\}$, where $\mathbf{x} = \{x_t\}_{t=1}^T$, $\mathbf{k} = \{k_t\}_{t=1}^T$, $\mathbf{y} = \{y_t\}_{t=1}^T$ and $\mathbf{w} = \{w_t\}_{t=1}^T$. We time-cluster the data with a non-parametric Bayesian HMM, namely the sHDP-HMM, that allows us to group the time series in homogeneous regimes, as the standard HMM, but it does not need to assume known the number of regimes R that can be estimated along with the other model parameters.

The belonging to a regime is coded via a discrete random variable z_t ; if $z_t = r$, at time t the system is in regime r . In the sHDP-HMM is assumed that $z_t \in \mathbb{Z}^+ \setminus \{0\}$, i.e. the number of regimes is potentially infinite. Indeed since T is finite, the elements of the time series $\mathbf{z} = \{z_t \in \mathbb{Z}^+ \setminus \{0\}\}_{t=1}^T$ will assume only a finite number of states (R), i.e. the number of non-empty regimes.

The sHDP-HMM is a hierarchical model where, let $\boldsymbol{\psi}_r = \{\lambda_{x,r}, \eta_r, \xi_r, \nu_r, \lambda_{y,r}\}$ be the vector of parameters, the first stage is

$$P(\mathbf{x}, \mathbf{k}, \mathbf{y}, \mathbf{w}|\mathbf{z}, \{\boldsymbol{\psi}_r\}_{r=1}^{\infty}) = \prod_{t=1}^T \prod_{r=1}^{\infty} [P(x_t, k_t, y_t, w_t|\boldsymbol{\psi}_r)]^{I(r=z_t)}, \quad (9)$$

$$X_t, K_t, Y_t, W_t|\boldsymbol{\psi}_r \sim HiwpP_{\mathbb{D}}(\lambda_{x,r}, \eta_r, \xi_r, \nu_r, \lambda_{y,r}), \quad (10)$$

At the second level of the hierarchy, \mathbf{z} is assumed to follow a first order Markov process, i.e. $P(z_t|z_{t-1}, z_{t-2}, \dots, z_1) = P(z_t|z_{t-1})$, with

$$P(z_t|z_{t-1}, \boldsymbol{\pi}_{z_{t-1}}) = \pi_{z_{t-1}z_t}. \quad (11)$$

and $\boldsymbol{\pi}_j = \{\pi_{ji}\}_{i=1}^{\infty}$ is the j^{th} row of the transition matrix, i.e. the matrix that rules the probabilities to move from one state to another. The initial state z_0 , and the associated vector of probabilities, cannot be estimated consistently since we have no observations at time 0 (for details see [21]) and, without loss of generality, we set $z_0 = 1$. The equations (9), (10) and (11) define an HMM with an infinite number of states, i.e. $z_t \in \mathbb{Z}^+ \setminus \{0\}$, infinite number of parameters, $\{\boldsymbol{\psi}_r\}_{r=1}^{\infty}$, and vectors of probabilities of infinite length, $\boldsymbol{\pi}_j$ s.

The sHDP-HMM specification is concluded assuming the following:

$$\boldsymbol{\pi}_r|\rho, \gamma, \{\beta_j\}_{j=1}^{\infty}, \boldsymbol{\psi}_r \sim DP \left(\gamma, \sum_{j=1}^{\infty} ((1 - \rho)\beta_j + \rho I(r = j)) \right), \quad \rho \in [0, 1], \gamma \in \mathbb{R}^+, \quad (12)$$

$$\beta_r = \beta_r^* \prod_{j=1}^{r-1} (1 - \beta_j^*)$$

$$\beta_r^*|\tau \sim B(1, \tau), \quad \tau \in \mathbb{R}^+$$

$$\boldsymbol{\psi}_r|H \sim H,$$

where $\beta_r > 0$, $r = 1, 2, \dots, \infty$ and $\sum_{r=1}^{\infty} \beta_r = 1$, i.e. $\{\beta_r\}_{r=1}^{\infty}$ is a vector of probabilities, $B(\cdot, \cdot)$ is the Beta distribution, $DP(\cdot)$ indicates the Dirichlet process while H is a distribution over $\boldsymbol{\psi}_r$ that acts as a prior for the model parameters. Without loss of generality, let suppose that the first R states are

the non-empty ones, and let $\boldsymbol{\pi}_r^* = (\pi_{r1}, \dots, \pi_{rR}, \sum_{j=R+1}^{\infty} \pi_{rj})$. The definition of the Dirichlet process (see for example [12]) implies that (12) can be written as

$$\boldsymbol{\pi}_r^* | \rho, \gamma, \{\beta_j\}_{j=1}^{\infty}, \boldsymbol{\psi}_r \sim \text{Dir} \left(\gamma((1-\rho)\beta_1 + \rho I(r=1)), \dots, \gamma((1-\rho)\beta_R + \rho I(r=R)), \gamma(1-\rho) \sum_{j=R+1}^{\infty} \beta_j \right), \quad (13)$$

where $\text{Dir}(\cdot)$ is the Dirichlet distribution.

From (13) it is clear that, for $r, j = 1, \dots, R$, $E(\pi_{rj}) = (1-\rho)\beta_r + \rho I(r=j)$, $\text{Var}(\pi_{rj}) = \frac{((1-\rho)\beta_j + \rho I(r=j))(1-(1-\rho)\beta_j - \rho I(r=j))}{\gamma+1}$. The vector $\{\beta_r\}_{r=1}^{\infty}$ and ρ rules the mean value of the $\boldsymbol{\pi}_r^*$ (and then also the one of $\boldsymbol{\pi}_r$) and ρ is an additional weight added to the auto-transition probability (or self-transition probability) π_{rr} . ρ is needed otherwise the HMM tends to create redundant states ([22]). γ is directly proportional to the precision (the inverse of the variance) of π_{rj} . The parameters ρ , γ and τ rule the number of non empty regimes, R , and as they increase, R decreases (see [12]).

For a more detailed explanation of the properties and interpretation of the sHDP-HMM we refer the reader to [12].


3.4 Prior distributions and implementation details

To specify the prior distribution over $\boldsymbol{\psi}_r$, H , we follow two standard advices. The first is to have prior distributions that make the model parameters easy to update in the Markov chain Monte Carlo (MCMC) algorithm, i.e. within a Gibbs step. The second advice is to use prior distributions that allow an easy learning from the data and a robust estimation of the posterior distribution and it is usually achieved using weakly informative priors.


In the simulated and real data examples (Section 4), we will use the same set of priors. Specifically $\lambda_{y,r} \sim G(a_y, b_y)I(0, c_y)$, where $G(\cdot, \cdot)I(0, \cdot)$ indicates the truncated gamma distribution expressed in terms of shape and rate that, under specific set of parameters, can resemble the uniform distribution over $[0, c_y]$. Due to the standard conjugacy between the Poisson and the gamma, the prior on $\lambda_{y,r}$ leads to a truncated gamma full conditional. We use a uniform distribution for ν_r that leads to a beta full conditional. For the IWP parameters we follow the work of [17]. They show that an efficient MCMC algorithm can be defined if we set an upper bound on the range of $\lambda_{x,r}$, i.e. $\lambda_{x,r} \in [0, \lambda_{x,max}]$, we use a truncated gamma distribution for $\lambda_{x,r}$ and we fix the maximum value that k_t s can assume. The value $\lambda_{x,max}$ is chosen such that the IWP evaluated with $\lambda_{x,r} = \lambda_{x,max}$ is indistinguishable from a circular discrete uniform. [17] propose to find the value $\lambda_{x,max}$ by using the wrapped normal approximation of the IWP and the truncation strategy of [19]. Once we have the maximum value of λ , we find the maximum value of the k_t s that is needed to obtain a reasonable approximation of the IWP, i.e. $k_{max} = \lceil \frac{3\sqrt{\lambda_{max}}}{l} + \frac{\lambda_{max}}{l} - \frac{1}{2} \rceil$, where $\lceil \cdot \rceil$ is the ceiling operator. Then following [17], regardless on the prior distributions for η_r and ξ_r , we can define an efficient Gibbs sampler. For η_r and ξ_r we use uniform distributions in their respective domains.

R depends on ρ , γ and τ and then we choose to treat them as random quantities, with non informative prior for each of them, so to be non informative on the distribution of R . Following [12], the following priors, $\rho \in U(0, 1)$, $\gamma \sim G(a_\rho, b_\rho)$ and $\tau \sim G(a_\tau, b_\tau)$, lead to full conditionals easy to simulate. The MCMC update of the regime indicator variables is done with the beam sampler ([23]).

With regard to the time series \mathbf{y} , we can have two types of missing observations; when $y_t^* = 0, 1$, i.e. an informative missing since we know that the value of y_t is either 0 or 1, and when y_t^* is missing due to the malfunction of the anemometer, i.e. a non informative missing. In both cases the missing observations are estimated during the model fitting. If y_t is missing and $y_t^* = 0$ then $w_t = 0$ and the full conditional of y_t is $\text{Bern} \left(\frac{\lambda_{y,r} e^{-\lambda_{y,r}}}{e^{-\lambda_{y,r}}(1+\lambda_{y,r})} \right)$ if $x_t \in \mathbb{D}$, see equation (5), while $y_t = 0$ if $x = \{\emptyset\}$. If both y_t and y_t^* are missing, then w_t is missing as well and, if $x = \{\emptyset\}$ we have that $y_t = 0$ and $w_t = 0$, otherwise if $x \in \mathbb{D}$, we can simulate from the joint full conditional of (y_t, w_t) by first simulate y_t from




(a) Linear densities




(b) Circular densities

Figure 2: Simulate example: densities used to generate the example 1. The solid line is the first regime, the dashed the second and the dotted the third




(a) Linear densities




(b) Circular densities

Figure 3: Simulate example: densities used to generate the example 2. The solid line is the first regime, the dashed the second and the dotted the third

a Poisson with parameter $\lambda_{y,r}$, and then we set $w_t = 1$ if $y_t \geq 2$ and 0 otherwise. The non informative missing observation in the time series \mathbf{x} , due to a malfunction of the station, can be easily simulated, i.e. if $y_t \neq 0$ we simulate a value from the IWP, while if $y_t = 0$ we simulate w_t from a $Bern(\nu_r)$ and, if $w_t = 1$ then $x_t = \{\emptyset\}$ while if $w_t = 0$ we simulate a value from the IWP.




(a) Linear densities




(b) Circular densities

Figure 4: Simulate example: densities used to generate the example 3. The solid line is the first regime, the dashed the second and the dotted the third



(a) Linear densities



(b) Circular densities

Figure 5: Simulate example: densities used to generate the example 4. The solid line is the first regime, the dashed the second and the dotted the third

4 Application

In this Section, the HMM is applied to simulated and real data examples.

	Regime 1	Regime 2	Regime 3
$\hat{\lambda}_{y,r}$	1.056	10.059	30.103
CI	(0.981 1.138)	(9.860 10.266)	(29.771 30.467)
$\hat{\lambda}_{x,r}$	4.887	1.024	5.055
CI	(4.740 5.049)	(0.961 1.093)	(4.917 5.190)
$\hat{\nu}_r$	0.102	0.001	0.002
CI	(0.078 0.124)	(0.000 0.004)	(0.000 0.005)

Table 1: Simulated example: posterior estimates ($\hat{\cdot}$) and 95% credible intervals (CI) for $\lambda_{y,r}$, $\lambda_{x,r}$ and ν_r : example 1.

	Regime 1	Regime 2	Regime 3
$\hat{\lambda}_{y,r}$	1.010	5.130	10.249
CI	(0.930 1.093)	(4.954 5.307)	(9.993 10.518)
$\hat{\lambda}_{x,r}$	4.931	2.075	4.941
CI	(4.765 5.100)	(1.978 2.173)	(4.767 5.106)
$\hat{\nu}_r$	0.094	0.001	0.001
CI	(0.083 0 0.106)	(0.000 0.004)	(0.000 0.004)

Table 2: Simulated example: posterior estimates ($\hat{\cdot}$) and 95% credible intervals (CI) for $\lambda_{y,r}$, $\lambda_{x,r}$ and ν_r : example 2.

	Regime 1	Regime 2	Regime 3
$\hat{\lambda}_{y,r}$	1.015	10.029	30.419
CI	(0.940 1.089)	(9.817 10.244)	(30.075 30.776)
$\hat{\lambda}_{x,r}$	243.205	1.016	5.041
CI	(170.534 303.098)	(0.949 1.083)	(4.896 5.189)
$\hat{\nu}_r$	0.094	0.001	0.001
CI	(0.083 0.105)	(0.000 0.004)	(0.000 0.004)

Table 3: Simulated example: posterior estimates ($\hat{\cdot}$) and 95% credible intervals (CI) for $\lambda_{y,r}$, $\lambda_{x,r}$ and ν_r : example 3.

	Regime 1	Regime 2	Regime 3
$\hat{\lambda}_{y,r}$	0.954	4.907	10.117
CI	(0.881 1.030)	(4.701 5.111)	(9.868 10.379)
$\hat{\lambda}_{x,r}$	256.186	2.927	4.919
CI	(183.894 301.445)	(2.811 3.050)	(4.759 5.080)
$\hat{\nu}_r$	0.094	0.001	0.001
CI	(0.083 0.107)	(0.000 0.004)	(0.000 0.004)

Table 4: Simulated example: posterior estimates ($\hat{\cdot}$) and 95% credible intervals (CI) for $\lambda_{y,r}$, $\lambda_{x,r}$ and ν_r : example 4.

4.1 Simulated examples

We simulate 4 datasets with 3 regimes ($R = 3$), from the model described in Section 3.3. The parameters are chosen so that in two examples (examples 1 and 3) the marginal circular distributions are slightly overlapping (Figures 2 (b) and 4 (b)) as well as the linear ones (Figures 2 (a) and 4 (a)), and in two examples (examples 2 and 4) the overlapping involves larger portions of the distributions (Figures 3 and 5). In the examples 3 and 4, one of the marginal circular distribution is a discrete

		Destination		
		Regime 1	Regime 2	Regime 3
Origin	Regime 1	0.771 (0.743 0.812)	0.101 (0.083 0.121)	0.128 (0.098 0.150)
	Regime 2	0.108 (0.090 0.128)	0.788 (0.763 0.813)	0.104 (0.085 0.122)
	Regime 3	0.095 (0.079 0.113)	0.107 (0.090 0.126)	0.798 (0.774 0.820)

Table 5: Simulated example: posterior mean estimates and 95 % credible intervals for the transition probability matrix: example 1.

		Destination		
		Regime 1	Regime 2	Regime 3
Origin	Regime 1	0.772 (0.742 0.809)	0.121 (0.097 0.147)	0.107 (0.086 0.132)
	Regime 2	0.099 (0.081 0.119)	0.790 (0.762 0.817)	0.110 (0.090 0.135)
	Regime 3	0.097 (0.077 0.120)	0.132 (0.094 0.164)	0.791 (0.746 0.820)

Table 6: Simulated example: posterior mean estimates and 95 % credible intervals for the transition probability matrix: example 2.

		Destination		
		Regime 1	Regime 2	Regime 3
Origin	Regime 1	0.777 (0.752 0.802)	0.108 (0.090 0.128)	0.114 (0.095 0.134)
	Regime 2	0.122 (0.092 0.142)	0.774 (0.749 0.800)	0.104 (0.085 0.123)
	Regime 3	0.107 (0.090 0.126)	0.102 (0.085 0.121)	0.791 (0.766 0.813)

Table 7: Simulated example: posterior mean estimates and 95 % credible intervals for the transition probability matrix: example 3.

		Destination		
		Regime 1	Regime 2	Regime 3
Origin	Regime 1	0.786 (0.760 0.812)	0.112 (0.091 0.136)	0.101 (0.082 0.123)
	Regime 2	0.112 (0.091 0.133)	0.799 (0.770 0.826)	0.090 (0.069 0.113)
	Regime 3	0.104 (0.084 0.126)	0.095 (0.073 0.121)	0.800 (0.771 0.828)

Table 8: Simulated example: posterior mean estimates and 95 % credible intervals for the transition probability matrix: example 4.

uniform. In all the examples we set $T = 3000$ and $\mathbb{D} = \left\{ \frac{2\pi j}{36} \right\}_{j=0}^{35}$. The transition matrix has diagonal elements equal to 0.8 while the off-diagonal ones are 0.1.

We use the following set of parameters:

- Example 1:

$$\lambda_y = \begin{bmatrix} 1 \\ 10 \\ 30 \end{bmatrix}, \lambda_x = \begin{bmatrix} 5 \\ 1 \\ 5 \end{bmatrix}, \eta = \begin{bmatrix} -1 \\ 1 \\ 1 \end{bmatrix}, \xi = \begin{bmatrix} 5\frac{2\pi}{36} \\ 15\frac{2\pi}{36} \\ 0 \end{bmatrix}, \nu = \begin{bmatrix} 0.1 \\ 0 \\ 0 \end{bmatrix}.$$

- Example 2:

$$\lambda_y = \begin{bmatrix} 1 \\ 5 \\ 10 \end{bmatrix}, \lambda_x = \begin{bmatrix} 5 \\ 1 \\ 5 \end{bmatrix}, \eta = \begin{bmatrix} -1 \\ 1 \\ 1 \end{bmatrix}, \xi = \begin{bmatrix} 10\frac{2\pi}{36} \\ 15\frac{2\pi}{36} \\ 10\frac{2\pi}{36} \end{bmatrix}, \nu = \begin{bmatrix} 0.1 \\ 0 \\ 0 \end{bmatrix}.$$

- Example 3:

$$\lambda_y = \begin{bmatrix} 1 \\ 10 \\ 30 \end{bmatrix}, \lambda_x = \begin{bmatrix} 300 \\ 1 \\ 5 \end{bmatrix}, \eta = \begin{bmatrix} -1 \\ 1 \\ 1 \end{bmatrix}, \xi = \begin{bmatrix} 5\frac{2\pi}{36} \\ 15\frac{2\pi}{36} \\ 0 \end{bmatrix}, \nu = \begin{bmatrix} 0.1 \\ 0 \\ 0 \end{bmatrix}.$$

- Example 4:

$$\lambda_y = \begin{bmatrix} 1 \\ 5 \\ 10 \end{bmatrix}, \lambda_x = \begin{bmatrix} 300 \\ 1 \\ 5 \end{bmatrix}, \eta = \begin{bmatrix} -1 \\ 1 \\ 1 \end{bmatrix}, \xi = \begin{bmatrix} 10\frac{2\pi}{36} \\ 15\frac{2\pi}{36} \\ 10\frac{2\pi}{36} \end{bmatrix}, \nu = \begin{bmatrix} 0.1 \\ 0 \\ 0 \end{bmatrix}.$$

The example 4 is particularly challenging since there is a strong overlap between the circular distributions (see Figure 5).

Here, and in the real data example of Section 4.2, we run models estimations using a MCMCs with 100000 iterations, a burnin of 50000 and keep for inference one observation every 10 samples, i.e. for posterior estimates we use 5000 samples. Since in the simulated examples and in the real data one, we never observed a value of the linear variable greater than 50, then as prior for $\lambda_{y,r}$, the mean of the linear variable, we use $G(1, 0.00005)I(0, 50)$. Following [17], as priors for $\lambda_{x,r}$ we use $G(1, 0.00005)I(0, 500)$. To conclude the priors specification we choose $\gamma \sim G(1, 0.1)$, $\tau \sim G(1, 0.1)$, that are standard weak informative distributions, and discrete uniform for ν_k and ξ_k .

In all four examples the posterior distribution of R is concentrated over 3 (regimes). The parameters $\lambda_{y,r,s}$, ν_r s and the transition matrices are correctly estimated², see Tables from 1 to 8, and the marginal posterior distributions of the η_r s are concentrated on the ‘‘true’’ values used to simulate each dataset. In the examples 1 and 3 both $\lambda_{x,r}$ and ξ_r are always correctly estimated, except in the first regime of the third example where the posterior distribution of ξ_1 has non-zero probability on $22\frac{2\pi}{36}$, $23\frac{2\pi}{36}$ and $24\frac{2\pi}{36}$. Note that when the density of the discrete circular variable is really close to the uniform, the parameters of the IWP becomes weakly identifiable and then we are not surprised that in the third example, the CI of $\lambda_{x,1}$ has length ≈ 133 and the true value of ξ_1 is not inside the associated CI.

In the examples 2 and 4, the $\lambda_{x,r}$ s are right estimated in the first and third regimes while in the second regime is overestimated. The posterior distributions of the ξ_r s are concentrated over the values used to simulate the data in the first and third regime of the example 2 and the first regime of the example 4. In the second regimes of the example 2 and 4, the posteriors of ξ_k are concentrated over $14\frac{2\pi}{36}$ and $13\frac{2\pi}{36}$ respectively, while in the third regime, example 4, is concentrated over 0.

The model we propose is able to recover the parameters used to simulate the data in all identifiable situations.

4.2 Real data application

After the model fitting we observed that the posterior distribution of R is concentrated over 3. The three estimated regimes are numbered in increasing order, based on the value of mean wind speeds,

² A parameter is considered correctly estimated if the 95% credible interval (CI) contains the value used to simulate the data.

	Regime 1	Regime 2	Regime 3
$\hat{\lambda}_{y,r}$	3.908	8.646	13.112
CI	(3.794 4.017)	(8.222 9.061)	(12.633 13.590)
$\hat{\lambda}_{x,r}$	93.374	16.486	59.344
CI	(92.571 94.145)	(13.304 20.676)	(52.494 60.536)
$\hat{\nu}_r$	0.156	0.002	0.002
CI	(0.141 0.172)	(0.000 0.009)	(0.000 0.011)
$\hat{\mu}_r$	5.951	0.983	4.161
CI	(5.817 6.091)	(0.885 1.077)	(3.971 4.347)
\hat{c}_r	0.242	0.778	0.406
CI	(0.239 0.245)	(0.730 0.817)	(0.399 0.450)

Table 9: Real data example: posterior estimates ($\hat{\cdot}$) and 95% credible intervals (CI) for $\lambda_{y,r}$, $\lambda_{x,r}$ and ν_r .

		Destination		
		Regime 1	Regime 2	Regime 3
Origin	Regime 1	0.946 (0.934 0.957)	0.037 (0.028 0.047)	0.017 (0.011 0.025)
	Regime 2	0.127 (0.091 0.167)	0.808 (0.763 0.850)	0.064 (0.039 0.095)
	Regime 3	0.155 (0.115 0.200)	0.020 (0.005 0.043)	0.824 (0.780 0.867)

Table 10: Real data example: Posterior mean estimates and 95 % credible intervals for the transition probability matrix.

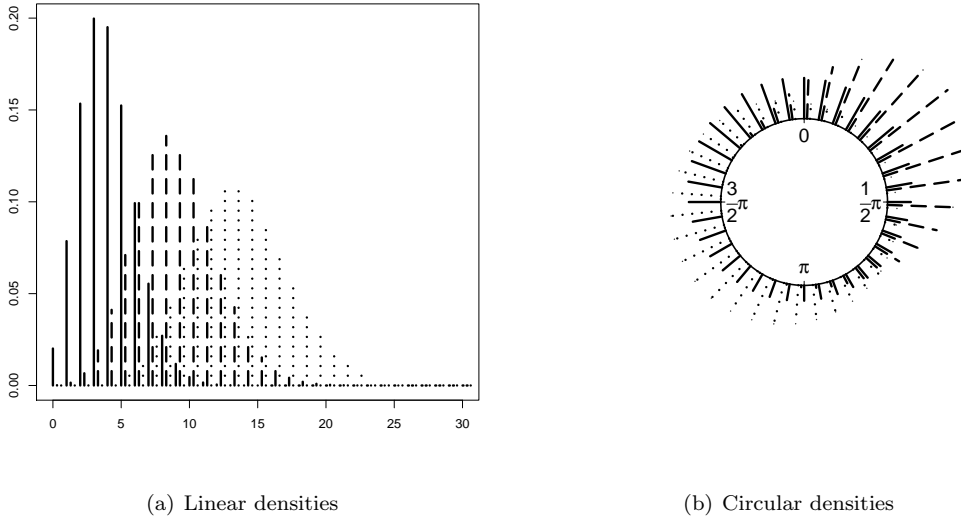


Figure 6: Real data example: predictive densities. The solid line is the first regime, the dashed the second and the dotted the third

$\lambda_{y,r}$, that can be seen in Table 9 along with the posterior estimates of $\lambda_{x,r}$, ν_r , and the circular mean and concentration. The predictive posterior densities are depicted in Figure 6.

With probability 1 the parameter η_r is -1 in the first regime and 1 in the second, while in the third is equal to -1 with probability 0.989 and 1 with probability 0.011. The posterior distribution of ξ_r is concentrated over $18\frac{2\pi}{36}$ in the first regime, in the second regime it has probability greater than 0 between the values $19\frac{2\pi}{36}$ and $31\frac{2\pi}{36}$ and its mode, $27\frac{2\pi}{36}$, has probability 0.223. In the third regime the posterior distribution of ξ_r assumes positive values between $26\frac{2\pi}{36}$ and $6\frac{2\pi}{36}^3$ and the modal value, $26\frac{2\pi}{36}$, has probability 0.845.

In the first regime we have a mean wind speed of 3.908, while the mean wind direction is 5.951, that corresponds to a direction between North and North-West. The circular concentration is 0.242 and the distribution is close to the discrete uniform, see Figure 6. The posterior mean value of ν_r , the hurdle probability, is 0.156. According to the empirical Beaufort scale ([24]), the first regime represents the *light breeze* state, where there are ripples without crests or small wavelets and then there is not cost erosion.

In the second regime, that can be considered as a transition state between the calm (first regime) and the storm (third) regime, the mean wind speed is 8.646 and the mean wind direction is 0.983, roughly North-East, while the circular concentration is 0.778, i.e. the second regime has a directional distribution more concentrated. The hurdle probability is 0.002, really close to zero.

In the third regime the mean wind speed is 13.112 and the mean direction is about South-West, $\hat{\mu}_3 = 4.161$. The circular concentration is 0.406 and again the mean hurdle probability is 0.002. In this regime, the distribution of the wind speed is fully concentrated between 5 and 23 knots, resulting in an extreme wave height of almost 4 meters in open water. It is interesting to note that, in this year, the winds with the higher speed are the ones blowing from the sea, more precisely from South West quadrant, resulting in a year with waves with more energy, intensification of erosion and changes in the longshore drift.

The posterior distribution of the transition probability matrix, Table 10, shows a strong self transition, i.e. the left side interval of the CIs of the self transition are always higher than 0.75.

5 Discussion

Motivated by our real data example, we introduced a new HMM for discrete circular-linear variables. Our data have some peculiar features: i) the linear and circular observations are interval-censored, ii) measurements of wind speed equal to 0 and 1 are not reliable, iii) some of the missing observations of the circular variables are informative on the values of the non reliable wind speed measurements. All these features was taken into account when the regime-specific density of the HMM was specified. We introduced a new circular-linear distribution that is suited to model our data. We estimated the model in a non-parametric Bayesian framework and we have shown how specific choices of prior distributions lead to a MCMC algorithm based only on Gibbs steps. We estimated the model on 4 simulated examples and then on the real one.

Future work will find us enriching the model in at least two directions. First we want to introduce a time-dependent transition matrix. Second, since the Poisson, due to the unit variance-to-mean ratio, can not model over-dispersed data, we are going to change the regime-specific marginal linear density to increase flexibility.

References

- [1] Ailliot P, Monbet V. Markov-switching autoregressive models for wind time series. *Environmental Modelling and Software*. 2012;30:92 – 101.
- [2] Belu R, Koracin D. Statistical and spectral analysis of wind characteristics relevant to wind energy assessment using tower measurements in complex terrain. *Journal of Wind Energy*. 2013;2013.

³ ξ is a discrete circular variable and then the left end of an interval can be greater than the right one

- [3] Lojowska A, Kurowicka D, Papaefthymiou G, van der Sluis L. Advantages of arma-garch wind speed time series modeling. In: Probabilistic Methods Applied to Power Systems (PMAPS), 2010 IEEE 11th International Conference on; June; 2010. p. 83–88.
- [4] Zárata-Miñano R, Anghel M, Milano F. Continuous wind speed models based on stochastic differential equations. *Applied Energy*. 2013 4;104:42–49.
- [5] Martin M, Cremades L, Santabarbara J. Analysis and modelling of time series of surface wind speed and direction. *International Journal of Climatology*. 1999;19(2):197–209.
- [6] Holzmann H, Munk A, Suster M, Zucchini W. Hidden Markov models for circular and linear-circular time series. *Environmental and Ecological Statistics*. 2006;13(3):325–347.
- [7] Bulla J, Lagona F, Maruotti A, Picone M. A multivariate hidden Markov model for the identification of sea regimes from incomplete skewed and circular time series. *Journal of Agricultural, Biological, and Environmental Statistics*. 2012;17(4):544–567.
- [8] Lagona F, Picone M, Maruotti A. A hidden markov model for the analysis of cylindrical time series. *Environmetrics*. 2015;.
- [9] Mastrantonio G, Maruotti A, Jona Lasinio G. Bayesian hidden Markov modelling using circular-linear general projected normal distribution. *Environmetrics*. 2015;26:145–158.
- [10] Lindsey JC, Ryan LM. Methods for interval-censored data. *Statistics in Medicine*. 1998;17(2):219–238.
- [11] Teh YW, Jordan MI, Beal MJ, Blei DM. Hierarchical dirichlet processes. *Journal of the American Statistical Association*. 2006;101(476):1566–1581.
- [12] Fox EB, Sudderth EB, Jordan MI, Willsky AS. A sticky hdp-hmm with application to speaker diarization. *The Annals of Applied Statistics*. 2011 06;5(2A):1020–1056.
- [13] Rubin DB. Inference and missing data. *Biometrika*. 1976;63(3):581–592.
- [14] Mullahy J. Specification and testing of some modified count data models. *Journal of Econometrics*. 1986;33(3):341 – 365.
- [15] Girija SVS, Rao AVD, Srihari GVLN. On wrapped binomial model characteristics. *Mathematics and Statistics*. 2014;2(7):231 – 234.
- [16] Sarma R, Rao AVD, Girija SV. On characteristic functions of the wrapped lognormal and the wrapped weibull distributions. *Journal of Statistical Computation and Simulation*. 2011; 81(5):579–589.
- [17] Mastrantonio G, Jona Lasinio G, Maruotti A, Calise G. On initial direction, orientation and discreteness in the analysis of circular variables. *ArXiv e-prints*. 2015 Sep;.
- [18] Coles S. Inference for Circular Distributions and Processes. *Statistics and Computing*. 1998; 8(2):105–113.
- [19] Jona Lasinio G, Gelfand A, Jona Lasinio M. Spatial analysis of wave direction data using wrapped Gaussian processes. *Annals of Applied Statistics*. 2012;6(4):1478–1498.
- [20] Mastrantonio G, Gelfand AE, Jona Lasinio G. The wrapped skew Gaussian process for analyzing spatio-temporal data. *Stochastic Environmental Research and Risk Assessment*. 2015;To appear.
- [21] Cappé O, Moulines E, Ryden T. Inference in hidden markov models. *Springer Series in Statistics*; Springer; 2005; Available from: http://books.google.it/books?id=-3_A3_11yssc.

- [22] Teh YW, Jordan MI. Hierarchical bayesian nonparametric models with applications. In: Hjort N, Holmes C, Müller P, Walker S, editors. Bayesian nonparametrics: Principles and practice. Cambridge University Press; 2010.
- [23] Van Gael J, Saatchi Y, Teh YW, Ghahramani Z. Beam sampling for the infinite hidden Markov model. In: Proceedings of the 25th International Conference on Machine Learning; Helsinki, Finland; ICML '08. New York, NY, USA: ACM; 2008. p. 1088–1095.
- [24] Scot H. Defining the wind: The beaufort scale, and how a 19th century admiral turned science into poetry. *The Professional Geographer*. 2005;57(3):474–475; Available from: http://dx.doi.org/10.1111/j.0033-0124.2005.493_1.x.