

Ötting, Marius; Langrock, Roland; Maruotti, Antonello

**Article — Published Version**

## A copula-based multivariate hidden Markov model for modelling momentum in football

AStA Advances in Statistical Analysis

**Provided in Cooperation with:**

Springer Nature

*Suggested Citation:* Ötting, Marius; Langrock, Roland; Maruotti, Antonello (2021) : A copula-based multivariate hidden Markov model for modelling momentum in football, AStA Advances in Statistical Analysis, ISSN 1863-818X, Springer, Berlin, Heidelberg, Vol. 107, Iss. 1-2, pp. 9-27, <https://doi.org/10.1007/s10182-021-00395-8>

This Version is available at:

<https://hdl.handle.net/10419/286769>

**Standard-Nutzungsbedingungen:**

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

**Terms of use:**

*Documents in EconStor may be saved and copied for your personal and scholarly purposes.*

*You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.*

*If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.*



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



# A copula-based multivariate hidden Markov model for modelling momentum in football

Marius Ötting<sup>1</sup> · Roland Langrock<sup>1</sup> · Antonello Maruotti<sup>2,3</sup>

Received: 3 March 2020 / Accepted: 17 February 2021 / Published online: 19 March 2021  
© The Author(s) 2021

## Abstract

We investigate the potential occurrence of change points—commonly referred to as “momentum shifts”—in the dynamics of football matches. For that purpose, we model minute-by-minute in-game statistics of Bundesliga matches using hidden Markov models (HMMs). To allow for within-state dependence of the variables, we formulate multivariate state-dependent distributions using copulas. For the Bundesliga data considered, we find that the fitted HMMs comprise states which can be interpreted as a team showing different levels of control over a match. Our modeling framework enables inference related to causes of momentum shifts and team tactics, which is of much interest to managers, bookmakers, and sports fans.

## 1 Introduction

Vocabulary such as “momentum”, “momentum shift”, or related terms is commonly used to refer to change points in the dynamics of a sports match. Usage of such terms is typically associated with situations during a match where an event—such as a shot hitting the woodwork in an association football match—seems to change the dynamics of the match, e.g. in a sense that a team which prior to the event had been pinned back in its own half suddenly seems to dominate the match. A prominent example is the 2005 Champions League final between Milan and Liverpool, where Liverpool was trailing by three goals after the first half, but fought back after half time and eventually won by penalty shootout.

Despite the widespread belief in momentum shifts in sports, it is not always clear to what extent *perceived* shifts in the momentum are genuine. From the literature on the “hot hand”—i.e. research on serial correlation in human performances—it is

---

✉ Marius Ötting  
marius.oetting@uni-bielefeld.de

<sup>1</sup> Bielefeld University, Bielefeld, Germany

<sup>2</sup> Libera Università Maria Ss. Assunta, Rome, Italy

<sup>3</sup> University of Bergen, Bergen, Norway

well known that most people do not have a good intuition of randomness and in particular tend to overinterpret streaks of success and failure (see, for example, Thaler and Sunstein 2009, pp. 30–34, and Kahneman 2011, pp. 114–118). It is thus to be expected that many perceived momentum shifts are in fact cognitive illusions in the sense that the observed shift in a competition's dynamics is driven by chance only.

Momentum shifts have been investigated in qualitative psychological studies, e.g. by interviewing athletes, who reported momentum shifts during matches (see, for example, Richardson et al. 1988; Jones and Harwood 2008). Fuelled by the rapidly growing amount of freely available sports data, quantitative studies have investigated the drivers of ball possession in football (Lago-Peñas and Dellal 2010), the detection of main playing styles and tactics (Diquigiovanni and Scarpa 2018; Gonçalves et al. 2017) and the effects of momentum on risk-taking (Lehman and Hahn 2013). In some of the existing studies, e.g. in Lehman and Hahn (2013), momentum is not investigated in a purely data-driven way, but rather pre-defined as winning several matches in a row.

In this contribution, we analyse potential momentum shifts within football matches. Specifically, we investigate the potential occurrence of momentum shifts by analysing minute-by-minute bivariate summary statistics from the German Bundesliga using hidden Markov models (HMMs). The corresponding data are described in Sect. 2. Within the HMMs, we consider copulas to allow for within-state dependence of the variables considered. The corresponding methodology is presented in Sect. 3. Our results, which are presented in Sect. 4, suggest states which can be tied to different levels of control in a match. In addition, we investigate the causes of momentum shifts, e.g. the current score of the match. This type of insight could be of great interest to managers, bookmakers, and sports fans.

## 2 Data

We analyse minute-by-minute in-game statistics of Bundesliga matches, taken from [www.whoscored.com](http://www.whoscored.com), to investigate to what extent momentum shifts in a football match are genuine, and what kind of events lead to a shift. Since the quality and tactics differ between the teams, we do not pool data from multiple teams, but consider data from a single team. Throughout this paper, we consider data from Borussia Dortmund. In the Supplementary Material, we present the same analysis for Hannover 96.

As proxy measures for the current momentum within a football match, we consider the number of shots on goal and the number of ball touches, with both variables sampled on a minute-by-minute basis. For match  $m$ ,  $m = 1, \dots, 34$ , this results in a bivariate time series  $\{\mathbf{y}_{mt}\}_{t=1,2,\dots,T_m}$ , with  $\mathbf{y}_{mt} = (y_{mt1}, y_{mt2})$  the pair of variables observed at time  $t$  (out of  $T_m$  minutes played) during the match.

Due to injury times being added to the regular match length of 90 min, the lengths of the time series considered range from 91 to 100 min. The final data set then comprises 3214 bivariate observations from 34 matches of the season 2017/18. In addition, since the underlying dynamics of a match, from Borussia Dortmund's perspective, potentially depend on characteristics of the opponent (such as the strength of

the squad) as well as events in the match (such as goals), the following four covariates are considered:

- the market value of the opposing team (taken from [www.transfermarkt.com](http://www.transfermarkt.com));
- the goal difference in the current score;
- a dummy variable indicating whether the match is played at home or away;
- the current minute of the match.

The first covariate considered is a (crude) proxy for the quality of teams. Specifically, a team's market value is given by the sum of all players' market values at the beginning of the season and thus does not vary between matches or within matches, e.g. if players are substituted. The difference in the current score is calculated from Borussia Dortmund's point of view, i.e. positive values refer to a lead of Dortmund, whereas negative values represent that Dortmund is trailing. The dummy indicating whether the match is played at home is included since several studies provided evidence for a home field advantage, because of (e.g.) crowd effects and psychological advantage when playing at home (see, for example, Pollard 2008). Finally, to account for the potential state of exhaustion of players, the minute of the match is also included. The variables considered are summarised in Table 1.


One example bivariate time series from the data set corresponding to the in-game statistics observed for Borussia Dortmund in the match against FC Schalke 04 played in November 2017 is shown in Fig. 1. In the media, this match was said to have a momentum shift, since Borussia Dortmund was in a 4:0 lead at half time, but Schalke 04 scored four goals in the second half so that the match resulted in a draw.

### 3 Modelling momentum

Figure 1 underlines that there are periods in the match where Borussia Dortmund's number of ball touches and the number of shots on goal are fairly low (e.g. around minute 75–90), as well as periods with relatively many ball touches and shots on goal (e.g. around minute 15–30). HMMs hence constitute a natural modelling approach for the minute-by-minute bivariate time series data, as they accommodate the idea of a match progressing through different phases, with potentially changing momentum. The states can be interpreted as the underlying momentum, i.e. as potentially different levels of control of the team considered. In the simplest model

**Table 1** Descriptive statistics of the variables analysed, “shots” and “ball touches”, as well as the covariates “market value” and “score difference”

	Mean	St. Dev.	Min.	Max.
shots	0.150	0.412	0	3
ball touches	6.101	5.036	0	28
market value (in $10^6$ Euro)	142.6	127.1	48.80	610.3
score difference	0.253	1.500	−6	5




**Fig. 1** Bivariate time series of the number of shots on goal (top) and the ball touches (bottom) of Borussia Dortmund for one example match from the data set (Borussia Dortmund vs. FC Schalke 04)

formulation with two states, the states could, for example, be interpreted as either the team considered or the opponent having a high level of control (i.e. dominating the match). In this section, the basic HMM model formulation will be introduced (Sect. 3.1) and extended to allow for within-state dependence using copulas (Sect. 3.2). The latter is desirable since the potential within-state dependence may lead to a more meaningful interpretation of the states regarding the underlying momentum. Finally, for the model formulation presented in Sect. 3.2, covariates will be included (Sect. 3.3).

### 3.1 A baseline model

HMMs involve two components: an unobserved Markov chain with  $N$  possible states and an observed state-dependent process, whose observations are assumed to be generated by one of  $N$  distributions as selected by the Markov chain. For the data considered in this paper, the observations and the state process are denoted by  $\mathbf{y}_{mt}$  and  $\{s_{mt}\}_{t=1,2,\dots,T_m}$ , respectively. Switches between the states are modelled by the transition probability matrix (t.p.m.)  $\mathbf{\Gamma} = (\gamma_{ij})$ , where  $\gamma_{ij} = \Pr(s_{mt} = j | s_{m,t-1} = i)$ ,  $i, j = 1, \dots, N$ . Figure 2 shows the model structure as directed graph. For the model formulation of an HMM to be completed, the number of states  $N$  and the class(es) of state-dependent distribution(s) have to be selected (see Zucchini et al. 2016, pp. 29–31). While choosing state-dependent distribution(s) is straightforward for univariate time series, it is generally not straightforward to define a multivariate distribution to allow for within-state dependence of the variables considered. This would be straightforward if the marginals are assumed to be normally distributed, as in that case a multivariate normal state-dependent distribution can be used (see, for example, Phillips et al. 2015). However, as this assumption would here clearly be inadequate given that



**Fig. 2** Dependence structure of the HMM considered: each pair of observations  $\mathbf{y}_{mt}$  is assumed to be generated by one of  $N$  (bivariate) distributions according to the state process  $s_{mt}$

we consider count data, for the vector of observations  $\mathbf{y}_{mt}$  in the baseline model formulation, we assume that the joint probability is obtained by the product of the marginal distributions,

$$f(\mathbf{y}_{mt} | s_{mt}) = \prod_{k=1}^K f(y_{mtk} | s_{mt}), \tag{1}$$

with  $K = 2$  here. This assumption, also known as *contemporaneous conditional independence*, is often used in practice (see, for example, Wall and Li 2009; DeRuiter et al. 2017; Punzo et al. 2018; van Beest et al. 2019). Taking the product of the marginal distributions is straightforward and allows a flexible choice of the marginals  $f(y_{mtk} | s_{mt})$ ,  $k = 1, \dots, K$ . In Eq. (1), each of these denotes a probability mass function (p.m.f.) since we deal with discrete data, but in principle  $f$  could also denote a density without any further changes in the baseline model formulation. The  $K = 2$  variables modelled here will still be unconditionally dependent when assuming contemporaneous conditional independence, as the underlying Markov chain induces both serial dependence and cross-dependence between them. The contemporaneous conditional independence assumption will not be made in the next subsection.

Since both the number of shots on goal and the number of ball touches are count data, the Poisson distribution would be a standard choice for either of the two variables. Here, to account for possible over- and underdispersion in the data, a Conway–Maxwell–Poisson (CMP) distribution is assumed both for the number of shots on goal and the number of ball touches, with p.m.f.

$$\Pr(X = x) = \frac{1}{Z(\lambda, \nu)} \frac{\lambda^x}{(x!)^\nu},$$

with  $Z(\lambda, \nu) = \sum_{k=0}^{\infty} \lambda^k / (k!)^\nu$ ,  $\lambda > 0$  and  $\nu \geq 0$  (Conway and Maxwell 1961). The CMP distribution contains some well-known discrete distributions:

- for  $\nu = 1$ ,  $Z(\lambda, \nu) = e^\lambda$ , and the CMP distribution simply reduces to the ordinary Poisson( $\lambda$ );

- for  $\nu \rightarrow \infty$ ,  $Z(\lambda, \nu) \rightarrow 1 + \lambda$ , and the CMP distribution approaches the Bernoulli with parameter  $\lambda(1 + \lambda)^{-1}$ ;
- for  $\nu = 0$  and  $0 < \lambda < 1$ ,  $Z(\lambda, \nu)$  is a geometric sum

$$Z(\lambda, \nu) = \sum_{j=0}^{\infty} \lambda^j = \frac{1}{1 - \lambda},$$

and, accordingly, the CMP distribution reduces to the geometric distribution  $p_x = \lambda^x(1 - \lambda)$ ;

- for  $\nu = 0$  and  $\lambda \geq 1$ ,  $Z(\lambda, \nu)$  does not converge, leading to an undefined distribution.

In general, the normalising constant  $Z(\lambda, \nu)$  does not reduce to such a simple closed-form expression. Asymptotic results are, however, available (Gillispie and Green 2015).

To formulate the likelihood for the baseline model, the  $i$ -th diagonal element of the  $N \times N$  diagonal matrix  $\mathbf{P}(\mathbf{y}_{m_i})$  consists of the joint probability of the observations  $y_{m_{i1}}$  and  $y_{m_{i2}}$  given state  $i$ , i.e.  $f(y_{m_{i1}} | s_{m_i} = i) \cdot f(y_{m_{i2}} | s_{m_i} = i)$ . Since the Conway–Maxwell–Poisson distribution contains an infinite sum in the normalising constant, the evaluation of the p.m.f. is not straightforward. Here, the R package `COMPOISSONREG` was used for this purpose (Sellers et al. 2018). Since stationarity cannot reasonably be assumed in our setting, we estimate the initial distribution  $\boldsymbol{\delta} = (\Pr(s_{m_1} = 1), \dots, \Pr(s_{m_1} = N))$ , regarding the parameters of  $\boldsymbol{\delta}$  as  $N - 1$  additional parameters to be estimated. The initial distribution is assumed to be constant across matches. With these quantities defined, the likelihood for a single match  $m$  is given by:

$$L = \boldsymbol{\delta} \mathbf{P}(\mathbf{y}_{m_1}) \boldsymbol{\Gamma} \mathbf{P}(\mathbf{y}_{m_2}) \dots \boldsymbol{\Gamma} \mathbf{P}(\mathbf{y}_{m_{T_m}}) \mathbf{1},$$

with column vector  $\mathbf{1} = (1, \dots, 1)' \in \mathbb{R}^N$  (see Zucchini et al. 2016, p. 37). Calculation of this matrix product expression amounts to the application of the forward algorithm, which is a powerful recursive technique for efficiently calculating the likelihood of an HMM at computational cost  $\mathcal{O}(TN^2)$  only (see Zucchini et al. 2016, p. 38). To obtain the likelihood for the full data set, we assume independence between the individual matches. The likelihood is thus given by the product of likelihoods for the individual matches:

$$L = \prod_{m=1}^{34} \boldsymbol{\delta} \mathbf{P}(\mathbf{y}_{m_1}) \boldsymbol{\Gamma} \mathbf{P}(\mathbf{y}_{m_2}) \dots \boldsymbol{\Gamma} \mathbf{P}(\mathbf{y}_{m_{T_m}}) \mathbf{1}. \tag{2}$$

The model formulation presented here could be extended to account for momentum carry-over effects across matches, but this is not investigated in the present work since there is usually a time difference of 5–7 days between matches. The model parameters are estimated by numerical maximum likelihood using the function `nlm()` in R (R Core Team 2017). To avoid local maxima, we selected starting values for the numerical maximisation by drawing random numbers from uniform

distributions 50 times and choosing the model with the best likelihood. An exploratory analysis guided the choice of what constitutes reasonable ranges for the parameter values for the state-dependent distributions. For a model with  $N = 2$  states, it took less than a minute to numerically maximise the likelihood on a standard desktop computer. In the Supplementary Material of this article, we provide data and code for all models presented.

### 3.2 Modelling within-state dependence using copulas

In the baseline model formulation, we assume contemporaneous conditional independence, i.e. that there is no within-state dependence between the two variables considered. However, when modelling momentum in football, it is of interest to explicitly model any within-state dependence to draw a comprehensive picture of the dynamics of a match. For example, high ball possession can be linked to both an attacking phase with lots of shots on goal, but also much less goal-oriented tactics, where the main aim is simply to control the match by keeping possession of the ball, without much pressure on goal. The between-variable correlation would likely be very different in those two scenarios. By estimating the within-state dependence between the two variables, we are better able to distinguish between such fairly subtle differences in a team's style of play.

To modify the contemporaneous conditional independence assumption, a multivariate distribution needs to be assumed to specify the dependence structure between the variables considered within states. Here, we allow for within-state dependence of our variables  $\mathbf{y}_{mt}$  by formulating a bivariate distribution as state-dependent distribution using a copula. A copula is a multivariate probability distribution with uniform margins. As introduced by Sklar (1959), the idea of a copula is to split a multivariate distribution into its univariate margins and the dependence structure, where the latter depends on the copula considered. Within the class of HMMs, copulas have previously been used by Härdle et al. (2015) to model within-state dependence in financial data, and by Brunel and Pieczynski (2005) and Lanchantin et al. (2011) for image analysis. For our modelling approach, we again consider the Conway–Maxwell–Poisson both for the number of shots on goal and the number of ball touches as marginal distribution. With  $F_1(y_{mt1} | s_{mt})$  and  $F_2(y_{mt2} | s_{mt})$  denoting the (state-dependent) cumulative distribution function of the marginals, the bivariate state-dependent distribution is given by

$$F(\mathbf{y}_{mt} | s_{mt}) = C(F_1(y_{mt1} | s_{mt}), F_2(y_{mt2} | s_{mt})),$$

where  $C(., .)$  is a bivariate copula. When deriving the corresponding p.m.f., differences are needed rather than derivatives, since the marginals are discrete (see, for example, Nikoloulopoulos 2013). Thus, the bivariate p.m.f. of  $\mathbf{y}_{mt}$  given state  $s_{mt}$  is given by



$$\begin{aligned}
 f(\mathbf{y}_{m_t} | s_{m_t}) &= C(F_1(y_{m_t1} | s_{m_t}), F_2(y_{m_t2} | s_{m_t})) \\
 &\quad - C(F_1(y_{m_t1} - 1 | s_{m_t}), F_2(y_{m_t2} | s_{m_t})) \\
 &\quad - C(F_1(y_{m_t1} | s_{m_t}), F_2(y_{m_t2} - 1 | s_{m_t})) \\
 &\quad + C(F_1(y_{m_t1} - 1 | s_{m_t}), F_2(y_{m_t2} - 1 | s_{m_t})).
 \end{aligned} \tag{3}$$

The copula  $C(\cdot, \cdot)$  needs to be selected from the large number of possible copula functions available in the literature. Here, we focus on copulas that can model positive and negative dependence. Archimedean copulas (see, for example, Nelsen 2006, pp. 116–118, for an overview) are convenient for this modelling purpose. We consider three different families of copulas, comparing their fit to the data in Sect. 4: first, the Frank copula, which is for two marginals  $u_1$  and  $u_2$  defined by

$$C(u_1, u_2) = -\frac{1}{\theta} \log \left( 1 + \frac{(\exp(-\theta u_1) - 1)(\exp(-\theta u_2) - 1)}{\exp(-\theta) - 1} \right), \theta \in \mathbb{R} \setminus \{0\},$$

second, the Clayton copula,

$$C(u_1, u_2) = \left( \max\{u_1^{-\theta} + u_2^{-\theta} - 1; 0\} \right)^{-1/\theta}, \theta \in [-1; \infty) \setminus \{0\},$$

and third, the Ali-Mikhail-Haq (AMH) copula,

$$C(u_1, u_2) = \frac{u_1 u_2}{1 - \theta(1 - u_1)(1 - u_2)}, \theta \in [-1, 1),$$

where for each copula considered the dependence parameter is denoted by  $\theta$ . As  $\theta \rightarrow 0$ , each of the three copulas above approaches the independence copula. For the Frank copula, as  $\theta \rightarrow \infty$ , the copula converges to the co-monotonicity copula corresponding to perfect positive dependence, while for  $\theta \rightarrow -\infty$ , it converges to the counter-monotonicity copula corresponding to perfect negative dependence. For the Clayton copula, as  $\theta \rightarrow -1$  ( $\theta \rightarrow \infty$ ), the copula converges to the counter-monotonicity (co-monotonicity) copula with perfect dependence. The AMH copula converges to neither the co-monotonicity nor the counter-monotonicity copula (see Nelsen 2006, pp. 116–118).

With the copulas defined as above, the diagonal matrix  $\mathbf{P}(\mathbf{y}_{m_t})$  in the HMM likelihood (see Eq. 2) changes slightly. The  $i$ -th diagonal entry is now equal to  $f(\mathbf{y}_{m_t} | s_{m_t} = i)$  as defined in Eq. (3) instead of the product of the marginals. The corresponding likelihood is then again numerically maximised using the function `nlm()` in R. For that purpose, we again carefully selected several starting values, as it was done for the baseline model introduced above.

### 3.3 A model including covariates

In the previous subsections, the transition probabilities  $\gamma_{ij}$  were assumed to be constant over time. To account for possible events which may lead to state-switching, and hence to possible momentum shifts, we modify this assumption by explicitly

allowing the transition probabilities  $\gamma_{ij}^{(t)}$  to depend on covariates at time  $t$ . This is done by linking  $\gamma_{ij}^{(t)}$  to covariates  $x_1^{(t)}, \dots, x_p^{(t)}$  using the multinomial logit link:

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

with

$$\eta_{ij}^{(t)} = \begin{cases} \beta_0^{(ij)} + \sum_{l=1}^p \beta_l^{(ij)} x_l^{(t)} & \text{if } i \neq j; \\ 0 & \text{otherwise.} \end{cases}$$

Since the transition probabilities depend on covariates, the t.p.m.  $\Gamma_t$  is not constant across time anymore, i.e. the Markov chain is non-homogeneous. However, the structure of the HMM likelihood as stated in Eq. (2) is unaffected, i.e. the likelihood can still be maximised numerically, again with several sets of starting values to avoid local maxima.

## 4 Results

In this section, the different models presented in Sect. 3 are fitted to data on the matches of Borussia Dortmund in the 2017/18 Bundesliga season. To further illustrate the methodology, in particular for lower-ranked teams, in the Supplementary Material we provide the results also for Hannover 96.

### 4.1 Baseline model

For the baseline model, we make the contemporaneous conditional independence assumption, cf. Eq. (1), initially focusing on the case of  $N = 2$  states. The corresponding parameter estimates associated with the number of shots on goal are  $\hat{\lambda}_{\text{shots}} = (0.125, 0.149)$ ,  $\hat{\nu}_{\text{shots}} = (0.206, 0.001)$ , while for the number of ball touches, they are  $\hat{\lambda}_{\text{touches}} = (0.971, 2.381)$ ,  $\hat{\nu}_{\text{touches}} = (0.102, 0.390)$ . It is not straightforward here to compute the means of the fitted distributions due to the infinite sum in the normalising constant. MacDonald and Bhamani (2018) discuss several approaches and calculate the mean by  $\frac{1}{Z(\lambda, \nu)} \sum_{k=0}^d k \lambda^k / (k!)^\nu$  using a very large  $d$  (say  $d = 100$ ). Following this approach, the means of the number of shots on goal are 0.138 and 0.175 for states 1 and 2. For the ball touches, the means are 4.080 (state 1) and 10.104 (state 2). Thus, state 2 can be interpreted as the team considered, Borussia Dortmund, being more dominant, i.e. having a higher level of control over the match, than when being in state 1. The t.p.m. is estimated as

$$\hat{\Gamma} = \begin{pmatrix} 0.867 & 0.133 \\ 0.280 & 0.720 \end{pmatrix},$$

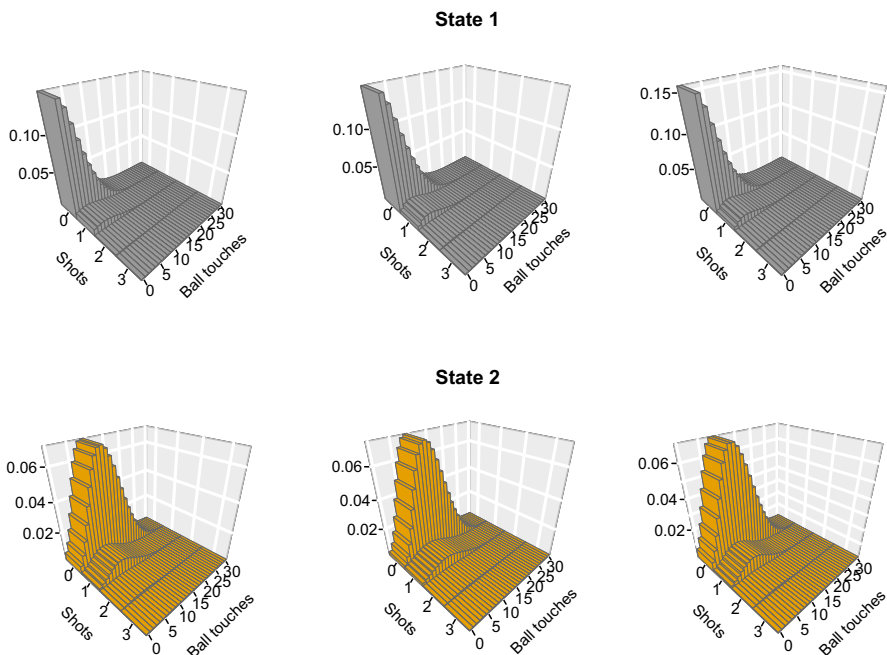
and the initial distribution as  $\hat{\delta} = (0.258, 0.742)$ . According to the t.p.m. of the fitted model, there is some persistence in both states. Although this is the simplest model

formulation considered here, the fitted model comprises interpretable states which refer to different levels of control over the match. The model can thus be regarded as a simple baseline model for capturing momentum shifts. We will now gradually increase its complexity to more fully capture the in-game dynamics.

#### 4.2 Copula-based HMM with $N = 2$

To capture possible within-state dependence of the variables, a multivariate distribution needs to be considered. For Poisson marginals, the bivariate Poisson as proposed by Karlis and Ntzoufras (2003) would be a possible candidate. However, as discussed in Sect. 3.1, this approach would have two limitations, namely the inability to capture overdispersion (or underdispersion) in the observations, and the restriction to positive between-variable correlation. Instead, we use more flexible CMP distributions for the marginals, stitching them together using a copula as described in Sect. 3.2.

First, we investigate the consequences of relaxing the contemporaneous conditional independence assumption. To this end, Fig. 3 displays the estimated state-dependent distributions of two-state copula-based HMM formulations, using the Frank, Clayton, and AMH copula. While visually there is no clear difference between the different copula functions considered, the application of the



**Fig. 3** Fitted state-dependent distributions for the two-state HMM for Borussia Dortmund. From left to right: Frank-, Clayton-, and AMH-copula

Clayton copula led to the highest likelihood of the fitted model. Compared to the baseline model, the copula-based model shows a clear improvement in the fit ( $\Delta AIC = 48; \Delta BIC = 35$ ). The fitted state-dependent distributions can again be interpreted as Borussia Dortmund exhibiting different levels of control, with state 1 corresponding to situations where the game is balanced, whereas state 2 refers to a high level of control. As for the baseline model, there is a fairly high persistence in the states, with the diagonal elements of the t.p.m. estimated as  $\hat{\gamma}_{11} = 0.852$  and  $\hat{\gamma}_{22} = 0.706$ .

### 4.3 Choosing the number of states

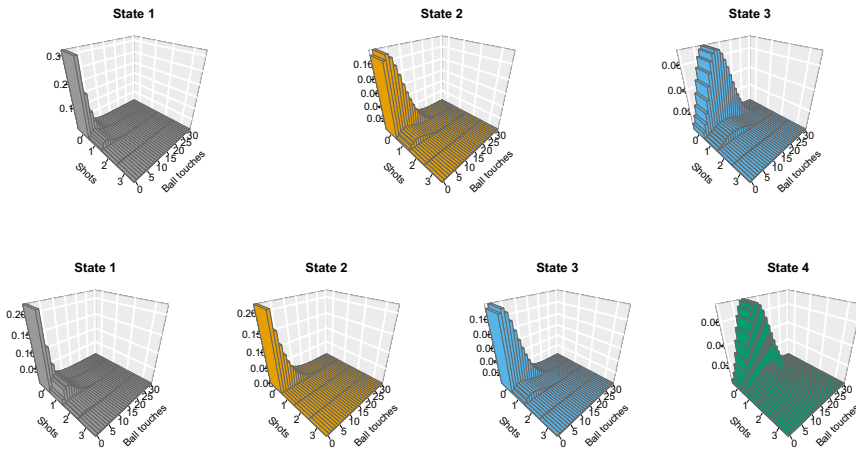
For the choice of the number of states, it is anything but clear how many states a given team may exhibit in a football match. To choose an appropriate number of states, and also a copula, we first consult the AIC and the BIC for the copula-based HMMs using different numbers of states and the three copulas considered above. The corresponding results are displayed in Table 2. Starting with the choice of the copula, the Clayton copula is preferred by both AIC and BIC. Hence, from now on, we use the Clayton copula. However, we note that when considering different marginal distributions, the fit of the copula also depends on the fit of the marginal distributions, which generally renders the choice of the copula a challenging task (see Mikosch 2006). For the number of states, the choice is not as conclusive: according to the AIC, the five-state model is preferred, whereas the BIC selects three states. As it is well known that the AIC tends to select too many states in a HMM (see Pohle et al. 2017), a choice of  $N = 3$  seems more appropriate based on these formal criteria. To make an informed choice based also on interpretability of the resulting model states, in Fig. 4 we further inspect the fitted models with three and four states, by means of their estimated state-dependent distributions. Figure 4 illustrates that the general patterns of the state-dependent distributions from the three-state model are also included in the four-state model, whereas the state-dependent distribution of state 2 in the four-state model seems to refer to an underlying level of control which is not included in the three-state model. However, at closer inspection of the distributional shapes in the four-state model, there is a substantial overlap between the state-dependent distributions of state 2 and state 3. Hence, given that the BIC points to the three-state model, and since we do not see meaningful additional information in a potential fourth state, from now on we focus exclusively on three-state models.

**Table 2** AIC and BIC for copula-based HMMs with different numbers of states. Values in bold indicate the number of states preferred by AIC and BIC, respectively

	Frank		Clayton		AMH	
	AIC	BIC	AIC	BIC	AIC	BIC
2 states	20,954	21,033	20,941	21,020	20,943	21,022
3 states	20,865	<b>21,005</b>	20,839	<b>20,979</b>	20,861	<b>21,001</b>
4 states	20,836	21,049	20,817	21,030	<b>20,831</b>	21,043
5 states	<b>20,814</b>	21,112	<b>20,801</b>	21,098	20,834	21,132

**Table 3** Parameter estimates for the state-dependent distributions of the Clayton-copula HMM with three states

Variable	State 1	State 2	State 3
Shots on goal	$\hat{\lambda} = 0.212, \hat{\nu} = 0.631$	$\hat{\lambda} = 0.117, \hat{\nu} \approx 0$	$\hat{\lambda} = 0.128, \hat{\nu} = 0.002$
Ball touches	$\hat{\lambda} = 0.670, \hat{\nu} \approx 0$	$\hat{\lambda} = 1.093, \hat{\nu} = 0.149$	$\hat{\lambda} = 2.145, \hat{\nu} = 0.352$
Dependence	$\hat{\theta} = 1.721$	$\hat{\theta} = 0.510$	$\hat{\theta} = -0.048$



**Fig. 4** State-dependent distributions for the three-state (top row) and four-state (bottom row) Clayton-copula HMM

### 4.4 Copula-based HMM with $N = 3$

For the Clayton-copula HMM with three states, Table 3 displays the estimated parameters of the marginal distributions as well as the dependence parameter of the copula. Deriving the corresponding means for the marginal distributions as described above yields means for the number of shots of 0.226, 0.132, and 0.147 for states 1, 2, and 3. For the number of ball touches, the corresponding means are 2.032 (state 1), 4.583 (state 2), and 9.732 (state 3). Based on the means and the corresponding distributional shapes (see top row in Fig. 4), the different states can be interpreted as Borussia Dortmund showing different levels of control over the match: low control in state 1, a fairly balanced match in state 2, and high control with lots of ball possession in state 3. State 1, with its relatively high mean number of shots on goal despite the fewer ball touches, likely includes several different styles of play with a low level of control, e.g. a defensive style of play, counter-attacks, and situations like (counter-)pressing. In state 3, the estimated negative dependence between the number of shots and ball touches may result from two different styles of high-control play: either Borussia Dortmund is

controlling and passing the ball without much pressure on goal, or they go effectively straight for goal, without much passing. In addition, the t.p.m. is estimated as

$$\hat{\Gamma} = \begin{pmatrix} 0.471 & 0.054 & 0.475 \\ 0.006 & 0.988 & 0.006 \\ 0.195 & \approx 0 & 0.805 \end{pmatrix}.$$

Here, with  $\hat{\gamma}_{22} = 0.988$  and  $\hat{\gamma}_{33} = 0.805$ , there is very high persistence in state 2 (balanced state) and moderately high persistence in state 3 (high-control state). Staying in state 1 (low control and quick counter-attacks) is relatively unlikely ( $\hat{\gamma}_{11} = 0.471$ ), and switching to the high-control state when being in state 1 is most likely. Up next we will present the results for the model including covariates in the state process.

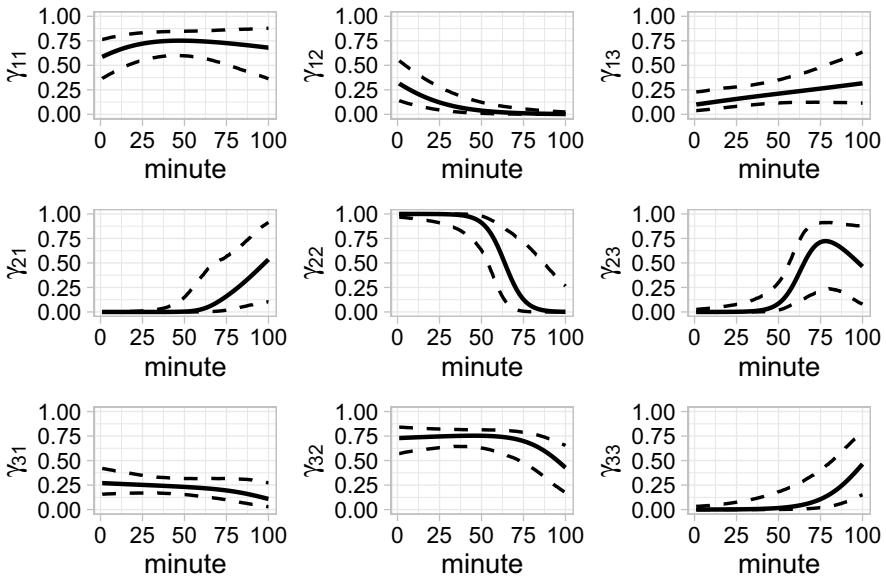
#### 4.5 A model including covariates

The models presented so far already provide interesting insights into the dynamics of football matches, since the state-dependent distributions can be tied to different levels of control of the team considered. To gain further insights, we incorporate covariates to investigate potential drivers of momentum shifts. According to the AIC, the model including all covariates considered is preferred over the model without covariates ( $\Delta\text{AIC} = 51$ ); we do not conduct variable selection as we regard this step of the analysis as explanatory (rather than an attempt to find the best model).

For ease of interpretation, we visualise the estimated transition probabilities as functions of covariates and present the theoretical stationary distributions of the Markov state process when fixing the covariate values at certain levels. The theoretical stationary distributions indicate how state occupancy, i.e. how much time is spent in a state, varies across different values of the covariate considered (Patterson et al. 2009). To illustrate these two approaches, we present (i) the transition probabilities as functions of the covariate minute and (ii) the stationary distributions with respect to the score difference. In Table 5 in the Supplementary Material, the estimated  $\beta_0^{(ij)}, \dots, \beta_p^{(ij)}$  and their 95% CIs are displayed.

For (i), as displayed in Fig. 5, the values of the score difference and the market value of the opponent are set to 0 and 200, corresponding to situations where the score is even and the opponent's strength is about average. In addition, we focus on home matches only, since the corresponding dummy variable in the linear predictor does not affect the overall pattern regarding the direction of the effect. The confidence intervals (indicated by the dashed lines) are obtained based on Monte Carlo simulation from the approximate multivariate normal distribution of the estimator. According to the estimated effects, switching from state 1 (low control and quick counter-attacks) and state 2 (balanced state) to state 3 (high-control state) becomes more likely at the end of matches. In addition, staying in state 3 also becomes more likely at the end of matches.

The stationary distributions for the score difference are shown in Table 4. The values of the minute and the market value of the opponent are fixed at 80 and 200, corresponding to situations in the final stage of a match with the opponent's



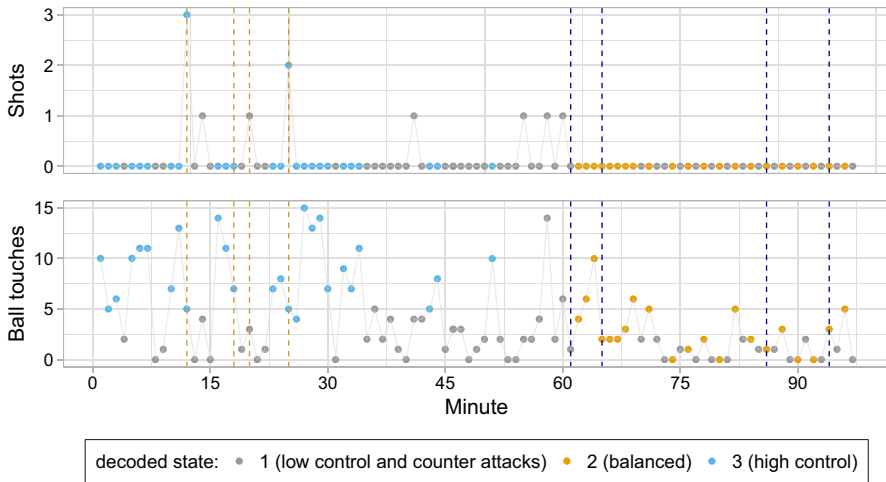
**Fig. 5** Transition probabilities as functions of the covariate minute. The dashed lines indicate confidence intervals (based on Monte Carlo simulation). The values of the score difference and the market value of the opponent are set to 0 and 200. Table 5 in the Supplementary Material displays the coefficients of the multinomial logistic regression underlying this Fig.

strength being about average. The stationary distributions indicate that there is a high probability for Borussia Dortmund to be in state 3 (high-control state) either if they have a clear lead or if they are trailing. In contrast, if they hold only a slender lead, then the probability of being in state 1 (low control and quick counter-attacks) is highest.

To further investigate typical patterns of momentum shifts according to the state process  $\{s_{mt}\}$ , we calculate the most likely trajectory of the states for match  $m$ . Specifically, for a given match  $m$ , we seek

$$(s_{m1}^*, \dots, s_{mT_m}^*) = \underset{s_{m1}, \dots, s_{mT_m}}{\operatorname{argmax}} \Pr(s_{m1}, \dots, s_{mT_m} \mid \mathbf{y}_{m1}, \dots, \mathbf{y}_{mT_m}),$$

i.e. the most likely state sequence, given the observations. Maximising this probability is equivalent to finding the optimal of  $N^{T_m}$  possible state sequences. This can be achieved at computational cost  $\mathcal{O}(T_m N^2)$  using the Viterbi algorithm (see Zucchini et al. 2016, pp. 88–92). Figure 6 displays the decoded sequences for the match Borussia Dortmund against Schalke 04 which is already shown in Fig. 1. We see that Borussia Dortmund started the match in the high-control state with occasional switches to the low control state with quick counter-attacks. According to the decoded state sequence, Borussia Dortmund was in the high-control state for most of the first half, and scored three of their four goals while in that state. After the half-time break, Borussia Dortmund was primarily in the low-control state with quick counter-attacks for about 15 min and subsequently alternated between this and



**Fig. 6** Decoded most likely state sequence of the match Borussia Dortmund against Schalke 04 according to the three-state Clayton-copula HMM including covariates. The vertical dashed lines denote goals scored by Borussia Dortmund (yellow lines) and Schalke 04 (blue lines). The goal leading to the intermediate score 2–0 was an own goal by Schalke 04

the balanced state. In the entire second half, Borussia Dortmund only once was in the high-control state.

At this point, it is worth emphasising that our fitted HMM cannot be expected to fully represent all structure and dynamics related to momentum shifts. First, when applied in an unsupervised setting as was done here, an HMM's model states will generally only be proxies for genuine states (Leos-Barajas et al. 2017). Second, while discrete states are conceptually appealing and mathematically convenient, it is not necessarily clear that different levels of control and hence momentum shifts are adequately represented by only finitely many states (cf. Ötting et al. 2020). Thus, the actual sequence of control levels may of course differ from the decoded sequence as shown in Fig. 6, and not every inferred state switch refers to a genuine switch in the actual momentum. However, as Borussia Dortmund was occupying the high-control state for most of the first half, but only once in the second half, the decoded state sequence is in agreement with the momentum shift around halftime as suggested by the media.

## 5 Discussion

There is wide interest in the dynamics of football matches, and specifically in potential momentum shifts, in particular by fans and the media. From a managerial perspective, it is important to understand the causes of such shifts, and hence also how to potentially exert an influence on the match outcome. With data sets on in-game summary statistics becoming freely available, we now have the opportunity to statistically investigate the corresponding processes. To that end, here we provide a



modelling framework—copula-based multivariate HMMs—which naturally accommodates potential changes in the dynamics of a match by relating the observed in-game match statistics to latent states. A key strength of the proposed approach is that we not only partition a given match into different phases but also allow for the investigation into drivers of how a match unfolds dynamically over time. Such in-game modelling could also be useful for bookmakers to obtain more precise estimations of betting odds. For instance, when modelling the time until the next goal during a football match, bookmakers could take into account the latent dynamics of a match as modelled here.

In our exploratory case study, we tested the feasibility of our approach by analysing minute-by-minute data on matches of one particular team, namely Borussia Dortmund. The three underlying states of the fitted model correspond to match phases where Borussia Dortmund exhibits a low level of control with quick counter-attacks, to phases where the match is balanced, and to those with high level of control. In addition, the estimated effects of the covariates shed some light on what kind of events may lead to switches between those states. Specifically, we found that Borussia Dortmund has the highest probability of being in the high-control state when having a clear lead or when trailing.

Although the states of the fitted models are tied to different levels of control, it remains unclear to what extent these can be attributed to shifts in the underlying momentum. Inference into the existence of potential momentum shifts is generally challenging given the absence of any formal definition of what constitutes momentum in sports. Without a clear definition, especially the relation between tactical changes and momentum shifts remains unclear—depending on the definition, it may be necessary to clearly differentiate between these, or alternatively tactical changes may at least need to be taken into account when investigating momentum shifts. In our case study, some of the reported effects may clearly arise from tactical considerations rather than momentum shifts. For example, for one-goal leads, switching to the low control and quick counter-attacks state may of course be a tactical consideration rather than a shift in the underlying momentum. The data considered here do not allow us to disentangle these two possible causes, rendering it impossible to arrive at a definitive conclusion whether the switches between the states are momentum shifts or tactical considerations. However, with the states and effects of the covariates considered (cf. Figure 5 and Table 4) being easy to interpret, they still provide interesting insights to dynamics of football matches.

**Table 4** Stationary distributions when fixing the score difference at certain levels

	-6	-5	-4	-3	-2	-1	0	1	2	3	4	5
state 1	0.073	0.100	0.134	0.175	0.222	0.280	0.523	0.732	0.705	0.642	0.560	0.475
state 2	0.391	0.364	0.334	0.301	0.267	0.234	0.206	0.175	0.147	0.122	0.098	0.076
state 3	0.535	0.535	0.532	0.524	0.511	0.486	0.271	0.094	0.148	0.236	0.342	0.450

Probabilities were calculated for each value of the score difference, with the market value of the opponent and the minute of the match fixed at 200 and 80, corresponding to situations in the final stage of a match against an opposing team of average quality

A clear limitation of the approach as presented here is that we focus on the in-game dynamics of only one of the two teams involved in a match, when in fact it is clear that the dynamics of a match result from the combination of both teams' actions. One way to achieve this would be to consider or even construct variables for the state-dependent process which reflect the actions of both teams, e.g. both teams' ball touches, or one team's proportion of ball touches in any given minute. It then seems conceptually appealing to jointly model both teams' underlying latent state corresponding to their exertion of control over the match, which could be achieved using a bivariate Markov chain, resulting in  $N^2$  combinations of states (see, for example, Sherlock et al. 2013; Pohle et al. 2020). In these model formulations, both teams' underlying state variables are allowed to interact. To further improve the realism of these models, it would be beneficial to also include tracking data, e.g. by considering the distances run per minute as covariate information.

The modelling framework used in the present contribution, i.e. copula-based HMMs for modelling football minute-by-minute data, can easily be transferred to other sports for further investigations and possible characteristics of momentum shifts. These sports include, e.g. basketball, where the variables to be modelled comprise, for example, the number of points/shots, the number of rebounds, and the number of blocks/steals. More generally, sports with two individuals or teams competing against each other and multiple variables measured on a fine-grained scale are best suitable for analysing momentum shifts using the modelling framework provided here.

**Supplementary Information** The online version supplementary material available at <https://doi.org/10.1007/s10182-021-00395-8>.

**Acknowledgements** This research was supported by Deutsche Forschungsgemeinschaft (Grant 431536450).

**Funding** Open Access funding enabled and organized by Projekt DEAL.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Brunel, N., Pieczynski, W.: Unsupervised signal restoration using hidden Markov chains with copulas. *Sig. Process.* **85**(12), 2304–2315 (2005)
- Conway, R.W., Maxwell, W.L.: A queuing model with state dependent service rates. *J. Ind. Eng.* **12**(2), 132–136 (1961)
- DeRuiter, S.L., Langrock, R., Skirbutas, T., Goldbogen, J.A., Calambokidis, J., Friedlaender, A.S., Southall, B.L., et al.: A multivariate mixed hidden Markov model for blue whale behaviour and responses to sound exposure. *Anna. Appl. Stat.* **11**(1), 362–392 (2017)

- Diquigiovanni, J., Scarpa, B.: Analysis of association football playing styles: an innovative method to cluster networks. *Stat. Model.* **19**(1), 1–27 (2018)
- Gillispie, S.B., Green, C.G.: Approximating the Conway-Maxwell-Poisson distribution normalization constant. *Statistics* **49**(5), 1062–1073 (2015)
- Gonçalves, B., Coutinho, D., Santos, S., Lago-Penas, C., Jiménez, S., Sampaio, J.: Exploring team passing networks and player movement dynamics in youth association football. *PLoS One* **12**(1), e0171156 (2017)
- Härdle, W.K., Okhrin, O., Wang, W.: Hidden Markov structures for dynamic copulae. *Econ. Theory* **31**(5), 981–1015 (2015)
- Jones, M.I., Harwood, C.: Psychological momentum within competitive soccer: players' perspectives. *J. Appl. Sport Psychol.* **20**(1), 57–72 (2008)
- Kahneman, D.: *Thinking, Fast and Slow*. Farrar, Straus and Giroux, New York (2011)
- Karlis, D., Ntzoufras, I.: Analysis of sports data by using bivariate Poisson models. *J. Roy. Stat. Soc. (Series D)* **52**(3), 381–393 (2003)
- Lago-Peñas, C., Dellal, A.: Ball possession strategies in elite soccer according to the evolution of the match-score: the influence of situational variables. *J. Hum. Kinetics* **25**, 93–100 (2010)
- Lanchantin, P., Lapuyade-Lahorgue, J., Pieczynski, W.: Unsupervised segmentation of randomly switching data hidden with non-Gaussian correlated noise. *Sig. Process.* **91**(2), 163–175 (2011)
- Lehman, D.W., Hahn, J.: Momentum and organizational risk taking: evidence from the National Football League. *Manage. Sci.* **59**(4), 852–868 (2013)
- Leos-Barajas, V., Photopoulou, T., Langrock, R., Patterson, T.A., Watanabe, Y.Y., Murgatroyd, M., Papastamatiou, Y.P.: Analysis of animal accelerometer data using hidden Markov models. *Methods Ecol. Evol.* **8**(2), 161–173 (2017)
- MacDonald, I.L., Bhamani, F.: A time-series model for underdispersed or overdispersed counts. *Am. Stat.* (2018). <https://doi.org/10.1080/00031305.2018.1505656>
- Mikosch, T.: Copulas: tales and facts. *Extremes* **9**(1), 3–20 (2006)
- Nelsen, R.B.: *An Introduction to Copulas*. Springer Science & Business Media, New York (2006)
- Nikoloulopoulos, A.K.: Copula-based models for multivariate discrete response data. In: Jaworski, P., Durante, F., Härdle, W.K. (eds.) *Copulae in Mathematical and Quantitative Finance*, pp. 231–249. Springer, Berlin (2013)
- Ötting, M., Langrock, R., Deutscher, C., Leos-Barajas, V.: The hot hand in professional darts. *J. Roy. Stat. Soc.: Ser. A (Statistics in Society)* **183**(2), 565–580 (2020)
- Patterson, T.A., Basson, M., Bravington, M.V., Gunn, J.S.: Classifying movement behaviour in relation to environmental conditions using hidden Markov models. *J. Anim. Ecol.* **78**(6), 1113–1123 (2009)
- Phillips, J.S., Patterson, T.A., Leroy, B., Pilling, G.M., Nicol, S.J.: Objective classification of latent behavioural states in bio-logging data using multivariate-normal hidden Markov models. *Ecol. Appl.* **25**(5), 1244–1258 (2015)
- Pohle, J., Langrock, R., van der Schaar, M., King, R., Jensen, F. H.: A primer on coupled state-switching models for multiple interacting time series. (2020). arXiv preprint [arXiv:2004.14700](https://arxiv.org/abs/2004.14700)
- Pohle, J., Langrock, R., van Beest, F.M., Schmidt, N.M.: Selecting the number of states in hidden Markov models: pragmatic solutions illustrated using animal movement. *J. Agric. Biol. Environ. Stat.* **22**(3), 270–293 (2017)
- Pollard, R.: Home advantage in football: a current review of an unsolved puzzle. *Open Sports Sci. J.* **1**(1), 12–14 (2008)
- Punzo, A., Ingrassia, S., Maruotti, A.: Multivariate generalized hidden Markov regression models with random covariates: physical exercise in an elderly population. *Stat. Med.* **37**(19), 2797–2808 (2018)
- R Core Team: *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria (2017)
- Richardson, P.A., Adler, W., Hankes, D.: Game, set, match: psychological momentum in tennis. *Sport Psychol.* **2**(1), 69–76 (1988)
- Sellers, K., Lotze, T., Raim, A.: *COMpoissonReg: Conway-Maxwell Poisson (COM-Poisson) Regression*. R package version (6), 1 (2018)
- Sherlock, C., Xifara, T., Telfer, S., Begon, M.: A coupled hidden Markov model for disease interactions. *J. Roy. Stat. Soc. (Series C)* **62**(4), 609–627 (2013)
- Sklar, M.: Fonctions de repartition an dimensions et leurs marges. *Publ. Inst. Stat. Univ. Paris* **8**, 229–231 (1959)
- Thaler, R.H., Sunstein, C.R.: *Nudge: Improving Decisions About Health, Wealth, and Happiness*. Penguin, London (2009)

- van Beest, F.M., Mews, S., Elkenkamp, S., Schuhmann, P., Tsolak, D., Wobbe, T., Bartolino, V., Bastardie, F., Dietz, R., von Dorrien, C., Galatius, A., Karlsson, O., McConnell, B., Nabe-Nielsen, J., Tange Olsen, M., Teilmann, J., Langrock, R.: Classifying grey seal behaviour in relation to environmental variability and commercial fishing activity—a multivariate hidden Markov model. *Sci. Rep.* **9**(1), 5642 (2019)
- Wall, M.M., Li, R.: Multiple indicator hidden Markov model with an application to medical utilization data. *Stat. Med.* **28**(2), 293–310 (2009)
- Zucchini, W., MacDonald, I.L., Langrock, R.: *Hidden Markov Models for Time Series: An Introduction Using R*. Chapman & Hall/CRC, Boca Raton (2016)

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.