Learning hidden Markov models with persistent states by penalizing jumps

Nystrup, Peter; Lindström, Erik; Madsen, Henrik

Published in:
Expert Systems with Applications

Link to article, DOI:
[10.1016/j.eswa.2020.113307](https://doi.org/10.1016/j.eswa.2020.113307)

Publication date:
2020

Document Version
Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):
Nystrup, P., Lindström, E., & Madsen, H. (2020). Learning hidden Markov models with persistent states by penalizing jumps. *Expert Systems with Applications*, 150, [113307]. <https://doi.org/10.1016/j.eswa.2020.113307>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Learning hidden Markov models with persistent states by penalizing jumps

Peter Nystrup^{ab*}, Erik Lindström^a, and Henrik Madsen^b

February 11, 2020

^a*Centre for Mathematical Sciences, Lund University, 221 00 Lund, Sweden*

^b*Department of Applied Mathematics and Computer Science, Technical University of Denmark, 2800 Kgs. Lyngby, Denmark*

Abstract

Hidden Markov models are applied in many expert and intelligent systems to detect an underlying sequence of persistent states. When the model is misspecified or misestimated, however, it often leads to unrealistically rapid switching dynamics. To address this issue, we propose a novel estimation approach based on clustering temporal features while penalizing jumps. We compare the approach to spectral clustering and the standard approach of maximizing the likelihood function in an extensive simulation study and an application to financial data. The advantages of the proposed jump estimator include that it learns the hidden state sequence and model parameters simultaneously and faster while providing control over the transition rate, it is less sensitive to initialization, it performs better when the number of states increases, and it is robust to misspecified conditional distributions. The value of estimating the true persistence of the state process is illustrated through a simple trading strategy where improved estimates result in much lower transaction costs. Robustness is particularly critical when the model is part of a system used in production. Therefore, our proposed estimator significantly improves the potential for using hidden Markov models in practical applications.

Keywords: Time series analysis; clustering; unsupervised learning; regime switching; regularization; dynamic programming.

1 Introduction

A Markov process is a random process in which the future is independent of the past, given the present. When the state space is discrete, Markov processes are known as Markov chains. In a hidden Markov model (HMM), the probability distribution that generates an observation depends on the state of an unobserved Markov chain. The model is popular in many areas with applications

*Correspondence to: Peter Nystrup, Centre for Mathematical Sciences, Lund University, Box 118, 221 00 Lund, Sweden.

Tel: +45 20420498

Email: peter.nystrup@matstat.lu.se, erik.lindstrom@matstat.lu.se, hmad@dtu.dk

ranging from face recognition (Nefian and Hayes, 1998) and natural language processing (Gales and Young, 2008; Kang et al., 2018) to wind- and solar-power forecasting (Pinson et al., 2008; Bhardwaj et al., 2013), fraud detection (Robinson and Aria, 2018), risk and return modeling in finance and energy markets (Rydén et al., 1998; Dias and Ramos, 2014; Petropoulos et al., 2016), and genome annotation in computational biology (Choo et al., 2004).

“All models are wrong but some are useful” (Box, 1979).

In many applications the model is only successful if the state sequence has a certain level of persistence. If the model is misspecified, for example because of the Markov assumption or erroneous conditional distributions, or misestimated, for example due to limited data being available, high dimensionality, or nonstationarity, it typically leads to unstable and impersistent estimates of the underlying state sequence and creation of unnecessary extra states (Bulla, 2011; Johnson and Will-sky, 2013; Fiecas et al., 2017). When the state sequence contains far too many jumps, the model essentially converges to a mixture model with no temporal information, which is of limited use.

Several studies have shown the potential benefit from changing asset allocation depending on regimes inferred from financial returns (see, e.g., Guidolin and Timmermann, 2007; Bulla et al., 2011; Bae et al., 2014; Nystrup et al., 2015a, 2017a, 2018, 2019); however, if the inferred regime changes too often, this results in excessive trading costs and inferior performance. When forecasting up and down regulation in power balancing markets, in addition to significant costs, there are physical constraints on how fast production can be changed in response to a change in the forecasted direction (see, e.g., Nielsen et al., 2016; Blanco and Morales, 2017). Hence, it is of paramount importance to develop an estimation framework where expert knowledge can be utilized to estimate persistent models.

1.1 Related work

Given a long sequence of time series data, it is often desirable to partition it into segments, where each segment can be explained by as simple a model as possible. Variants of the problem have been studied in several contexts, including change-point detection (Gustafsson, 2000; Oh and Han, 2000; Nystrup et al., 2016), segmentation (Katz and Crammer, 2015; Hallac et al., 2019), trend filtering (Kim et al., 2009), and mixture models (Verbeek et al., 2003; Samé et al., 2011; Dias et al., 2015).

The task of segmenting a time series is different from clustering in general, because the time ordering of the data must be taken into account. The different methods make different assumptions about the data. The ergodic HMM assumes that the underlying states repeat themselves, with some structure to when transitions are likely to occur. It is possible, though, to impose additional constraints to ensure nonrepeatability of the states, in which case the model is referred to as a left-to-right HMM (Bakis, 1976).

The standard, two-state Gaussian HMM has been extended in numerous ways to overcome its shortcomings in specific applications. Through conditional distributions with heavier tails (Bulla, 2011) or time-varying parameters (Nystrup et al., 2017b) it is possible to increase persistence and improve the model’s ability to reproduce long memory. Other studies have relaxed the Markov assumption on sojourn-time distributions (Bulla and Bulla, 2006; Langrock and Zucchini, 2011; Johnson and Will-sky, 2013), considered higher-order Markov chains (Shamshad et al., 2005; Petropoulos et al., 2017), and increased the size of the hidden state space (Ghahramani and Jordan, 1997; Beal et al., 2002; Nystrup et al., 2015b; Sedoc et al., 2018). In a Bayesian framework the

level of persistence can be adjusted through a hyperparameter in the hierarchical Dirichlet prior distribution imposed on the transition probability matrix (Beal et al., 2002; Fox et al., 2011).

Traditional approaches to estimating HMMs include direct numerical maximization of the likelihood function, the Expectation–Maximization (EM) algorithm (Baum et al., 1970; Dempster et al., 1977), and Bayesian estimation using Markov chain Monte Carlo methods (Rydén, 2008). From a theoretical point of view the maximum likelihood estimate (MLE) has favorable properties; for example, it is consistent and asymptotically efficient. However, computationally it is not very tractable (Andersson et al., 2003). This led Hsu et al. (2012) to propose a spectral algorithm for learning discrete-output HMMs based on a singular value decomposition of a matrix of joint probabilities of past and future observations, closely related to subspace identification methods used in control theory. Song et al. (2010) proposed a kernel-based approach for learning HMMs based on Hilbert space embedding.

Rousseeuw et al. (2015) used the spectral clustering algorithm of Ng et al. (2002) to estimate the parameters of a discrete-output HMM. They were able to recover the hidden partition without prior knowledge of the specific model parameters. Subsequently, Zheng et al. (2019) used spectral clustering to learn a continuous-output HMM based on simple features constructed from the observed time series. Their results showed that spectral clustering is more robust and more accurate than maximum likelihood estimation when the state process is less persistent. Meanwhile, both methods had issues with the highly persistent processes that are characteristic for many financial time series (Rydén et al., 1998; Cont, 2001; Nystrup et al., 2017b).

1.2 Contribution

We propose a novel way of learning HMMs based on temporal features and penalizing jumps. Penalizing jumps is key to successfully applying simple clustering techniques to learn HMMs with persistent states from temporal features. We rely on the same set of features considered by Zheng et al. (2019) combined with a variant of the framework proposed by Bemporad et al. (2018) for fitting jump models. Our contribution is to combine these two approaches into a robust estimator with possibility for utilizing prior knowledge to estimate persistent models. Without regularization on jumps it nests the standard K -means clustering algorithm of Lloyd (1982). When penalizing jumps it is reminiscent of the Viterbi Path Counting algorithm proposed by Davis and Lovell (2004). By extending several other methods, it has a wide application within expert and intelligent systems.

We extend the simulation study of Zheng et al. (2019) by considering multiple simulation lengths, different levels of persistence in the state process, models with multiple states, and the impact of misspecified conditional distributions. Our proposal of learning HMMs by clustering temporal features while penalizing jumps compares favorably to both spectral clustering and the traditional approach of maximizing the likelihood function using the EM algorithm.

The jump penalty regularizes the estimation problem and provides control over the transition rate. Other advantages of the proposed estimator include that it learns the hidden state sequence and model parameters simultaneously, it needs much fewer iterations to converge, it is less sensitive to initialization, it performs better when the number of states increases, and it is robust to misspecified conditional distributions. Robustness is particularly critical when the model is part of an expert system used in production. Therefore, our proposed estimator significantly improves the potential for using hidden Markov models in practical applications.

We illustrate the value of estimating the true persistence of the state process through a

Algorithm 1 Spectral clustering for HMMs.

Input: Time series $y = \{y_1, \dots, y_T\}$ and number of latent states K .

1. Construct a set of standardized features from the time series y .
2. Form an affinity matrix $A \in \mathbb{R}_{++}^{T \times T}$ from the standardized features.
3. Compute the random-walk normalized Laplacian $L_{\text{rw}} = I - D^{-1}A \in \mathbb{R}_{++}^{T \times T}$.
4. Find the first K eigenvectors of L_{rw} and collect them in a matrix $U \in \mathbb{R}^{T \times K}$.
5. Group the rows of U into K clusters using the K -means algorithm.
6. Compute the transition probabilities and distributional parameters for each cluster.

Output: HMM parameters and prediction of latent states.

long-short trading strategy using both simulated and real financial data. The ability to control the transition rate, and through that the turnover, results in much lower transaction costs.

We begin in Section 2 by revisiting the spectral clustering approach to HMM estimation. In Section 3, we outline our proposal of learning HMMs with persistent states by clustering temporal features while penalizing jumps. The simulation study is presented in Section 4. In Section 5, the estimated states form the basis of a trading strategy that illustrates the value of estimating the true persistence of the state process. Finally, Section 6 concludes.

2 Spectral clustering

Spectral clustering of affinity graphs is a popular tool in exploratory data analysis and has enjoyed much empirical success in machine learning (Shi and Malik, 2000; Ng et al., 2002; von Luxburg, 2007; Jia et al., 2013). The idea of spectral clustering is based on spectral graph theory. It treats the data clustering problem as a graph partitioning problem and constructs an undirected weighted graph with each point in the dataset being a vertex and the similarity value between any two points being the weight of the edge connecting the two vertices (Jia et al., 2013). This is commonly done by constructing a low-dimensional embedding of the affinity matrix, followed by K -means clustering in the low-dimensional space.

Algorithm 1 summarizes in six steps the spectral clustering algorithm for learning HMMs proposed by Zheng et al. (2019). The inputs are the time series $y = \{y_1, \dots, y_T\}$ and the number of latent states K . First, a set of features is constructed from the time series. Second, an affinity matrix $A \in \mathbb{R}_{++}^{T \times T}$ is constructed from the standardized features. Then the random-walk normalized Laplacian is computed, and its first K eigenvectors are collected in a matrix $U \in \mathbb{R}^{T \times K}$. Treating each row of U as a point in $\mathbb{R}^K$, they are grouped into K clusters using the K -means algorithm. Finally, the transition probabilities are computed, by counting the number of transitions between clusters, along with the distributional parameters for each cluster.

Normalizing the rows of U to have unit Euclidean norm before clustering, as proposed by Ng et al. (2002) and advocated by Weiss (1999), does not make a difference. Rather, it is important to standardize the features prior to forming the affinity matrix.

Spectral clustering estimates the parameters and the hidden states simultaneously. Each observation is assigned to a specific state, which is different from the traditional estimation methods that employ probabilistic assignment. Jin et al. (2006) proposed a probabilistic approach for op-

Algorithm 2 Features for HMM estimation.

Input: Time series $y = \{y_1, \dots, y_T\}$ and window length l .

1. Observation: y_t
2. Left absolute change: $|y_t - y_{t-1}|$
3. Right absolute change: $|y_{t+1} - y_t|$
4. Centered local mean: $\text{mean} \left[y_{t-\frac{l-1}{2}}, \dots, y_t, \dots, y_{t+\frac{l-1}{2}} \right]$
5. Centered local standard deviation: $\text{std} \left[y_{t-\frac{l-1}{2}}, \dots, y_t, \dots, y_{t+\frac{l-1}{2}} \right]$
6. Left local mean: $\text{mean} \left[y_{t-\frac{l-1}{2}}, \dots, y_t \right]$
7. Left local standard deviation: $\text{std} \left[y_{t-\frac{l-1}{2}}, \dots, y_t \right]$
8. Right local mean: $\text{mean} \left[y_t, \dots, y_{t+\frac{l-1}{2}} \right]$
9. Right local standard deviation: $\text{std} \left[y_t, \dots, y_{t+\frac{l-1}{2}} \right]$

Output: Feature set.

timizing spectral clustering and argued that soft, as opposed to hard, cluster membership meant that it was less likely to get trapped in a local minimum. However, that is not the impression the comparison in Section 4 with maximum likelihood estimation leaves.

2.1 Feature selection

The performance of spectral clustering is determined by the quality of the eigenvectors of the related graph Laplacian. Generally, the Laplacian is constructed using the full features, which will degrade the quality of the related eigenvectors when there are a large number of noisy or irrelevant features in datasets (Jiang and Ren, 2011). Without feature selection the performance of spectral clustering can degrade dramatically in the presence of irrelevant features (Bach and Jordan, 2004).

Algorithm 2 outlines the features proposed by Zheng et al. (2019) for learning HMMs. In addition to the observations and the left and right absolute changes, they include centered, left, and right local means and standard deviations. Based on a simulation experiment, Zheng et al. (2019) decided to use this feature set with two window lengths $l = 5$ and $l = 13$, i.e., 15 features in total.

The feature set in Algorithm 2 is simple, yet goes beyond the first two moments by including absolute differences. We will use this set in order to make our results directly comparable to those of Zheng et al. (2019). Additionally, we will use the same features as input to the jump estimator proposed in the next section to ensure that spectral clustering and jump estimation have the same starting point. It is left for future research to consider feature selection in depth.

2.2 Affinity matrix

The task of building a good affinity matrix has been shown to be largely responsible for the performance of spectral clustering algorithms (Bach and Jordan, 2004). We follow the most common approach by using the Gaussian kernel function defined by $A_{ij} = \exp \left(-\gamma \|x_i - x_j\|_2^2 \right)$, where x_i

and x_j are two feature vectors. It is simple to calculate and results in a positive definite affinity matrix. The scaling parameter $\gamma = 1/(2\sigma^2)$ controls how rapidly the affinity falls off with the distance between the features. It is a measure of when two feature vectors are considered similar and should be selected for the specific application, for example using cross-validation.

Successful recovery of the hidden partition requires the constructed affinity matrix to have a clear block structure. Fischer and Poland (2005) proposed different ways to amplify the block matrix structure for spectral clustering. Zhang et al. (2011) argued that it is desirable that with the same Euclidean distance, two points in the same cluster should have higher similarity than two points in different clusters, which is not the case for the standard Gaussian kernel. This led them to propose a local density adaptive similarity measure that uses the local density between two data points to scale the Gaussian kernel function. Some of the same smoothing effect is achieved through the use of rolling means and standard deviations in the feature calculation in Algorithm 2 as an alternative to considering the observations and their squared values.

2.3 Dimensionality reduction

After the affinity matrix is constructed, its random-walk normalized Laplacian $L_{\text{rw}} = I - D^{-1}A$ is computed (Meila and Shi, 2001). The degree matrix $D = \text{diag}(A1)$, where 1 is a column vector of ones, is a diagonal matrix whose i -th diagonal element is the sum of the elements in the i -th row of A .

Because the random-walk normalized Laplacian is positive semi-definite, its eigenvectors are real and nonnegative. In most of the literature on spectral clustering it is taken for granted that the eigenvalue problem is easy to solve (Bach and Jordan, 2004). In practice, it is necessary to use an algorithm that exploits the structure of a real symmetric matrix to avoid complex eigenvalues and -vectors as a result of numerical issues. Similarly, the first eigenvalue might be negative as a result of errors introduced by truncation and rounding errors. Several studies have demonstrated that by regularizing the Laplacian the ability of its eigenvectors to discriminate between the clusters can be improved (Chaudhuri et al., 2012; Amini et al., 2013; Qin and Rohe, 2013; Joseph and Yu, 2016).

It is a significant advantage of spectral clustering that the dimension of the clustering task only depends on the number of observations and the number of eigenvectors, and is independent of the number of features. Although the number of required eigenvectors is generally equal to the number of clusters, Jin et al. (2006) showed that it can be different when the clusters are not well separated. Zelnik-Manor and Perona (2005) suggested exploiting the structure of the eigenvectors to infer automatically the number of groups. Another possibility is to analyze the eigenvalues. We assume that the number of states K is known and select the same number of eigenvectors.

3 Penalizing jumps to increase persistence

Bemporad et al. (2018) proposed to fit jump models by minimizing the objective

$$\sum_{t=1}^{T-1} [\ell(y_t, \theta_{s_t}) + \lambda \mathbb{I}_{s_t \neq s_{t+1}}] + \ell(y_T, \theta_{s_T}), \quad (1)$$

where $\mathbb{I}$ is the indicator function, over the model parameters $\theta = \{\theta_1, \dots, \theta_K\}$ and the state sequence $s = \{s_1, \dots, s_T\}$. This can be viewed as a tradeoff between fitting the data and prior

Algorithm 3 Jump estimation of HMM.

Input: Time series $y = \{y_1, \dots, y_T\}$, number of latent states K , and initial state sequence $s^0 = \{s_1^0, \dots, s_T^0\}$.

1. Construct a set of standardized features z from the time series y .
2. Iterate for $i = 1, \dots$
 - (a) $\theta^i = \operatorname{argmin}_{\theta} \sum_{t=1}^T \ell(z_t, \theta_{s_t^{i-1}})$ (model fitting).
 - (b) $s^i = \operatorname{argmin}_s \left\{ \sum_{t=1}^{T-1} [\ell(z_t, \theta_{s_t}^i) + \lambda \mathbb{I}_{s_t \neq s_{t+1}}] + \ell(z_T, \theta_{s_T}^i) \right\}$ (state-sequence fitting).
3. Until $s^i = s^{i-1}$.
4. Compute the transition probabilities and distributional parameters for each state.

Output: HMM parameters and prediction of latent states.

assumptions about the state sequence.

The parameter $\lambda \geq 0$ is the penalty for jumps. For λ large enough, the jump model becomes a single model that covers the whole dataset. When $\lambda = 0$ it reduces to automatically splitting the dataset in K clusters and fitting one model per cluster, which is a generalization of the K -means algorithm of Lloyd (1982). We will consider the loss function $\ell(y_t, \theta_{s_t}) = \|y_t - \theta_{s_t}\|_2^2$, which is equivalent to K -means clustering.

3.1 Model fitting

Algorithm 3 outlines the proposed jump model approach to estimating HMMs. The inputs are the time series $y = \{y_1, \dots, y_T\}$, the number of latent states K , and an initial state sequence $s^0 = \{s_1^0, \dots, s_T^0\}$. First, a set of standardized features is constructed from the time series. Then, it iterates between finding the parameters that minimize the loss function for a given state sequence—i.e., fitting the model—and finding the state sequence that minimizes the objective function (1) given the parameters—i.e., fitting the state sequence. This process is repeated until the state sequence does not change. In practice, we finish after a maximum of ten iterations or once the value of the objective changes by less than 10^{-6} from one iteration to the next, which usually takes less than five iterations.

Based on the final estimate of the state sequence s^i , transition probabilities are computed, by counting the number of transitions between states, along with the distributional parameters for each state. To improve the quality of the solution, we run Algorithm 3 from ten different initial state sequences generated randomly using K -means++ (Arthur and Vassilvitskii, 2007) and keep the model that achieves the lowest objective value.

In Algorithm 3, Item 2a, the parameters $\theta^i = \{\theta_1^i, \dots, \theta_K^i\}$ are estimated by computing analytically the minimum of the loss function for each of the K states. In Item 2b, the most likely sequence of states is found using dynamic programming (Bellman, 1957). Define

$$V(T, s) = \ell(z_T, \theta_s), \tag{2a}$$

$$V(t, i) = \ell(z_t, \theta_i) + \min_j [V(t+1, j) + \lambda \mathbb{I}_{i \neq j}], \quad t = T-1, \dots, 1. \tag{2b}$$

The most likely sequence of states is then given by

$$s_1 = \underset{j}{\operatorname{argmin}} V(1, j), \quad (3a)$$

$$s_t = \underset{j}{\operatorname{argmin}} [V(t, j) + \lambda \mathbb{I}_{s_{t-1} \neq j}], \quad t = 2, \dots, T. \quad (3b)$$

If the time order of operations is reversed, it becomes the Viterbi algorithm (Viterbi, 1967).

Finding the most likely sequence of states requires $O(TK^2)$ operations, which is the same as one forward and backward pass with the EM algorithm. With the same computational complexity per iteration, it is a significant advantage of the jump model that it takes much fewer iterations to converge. In our simulation study, five iterations are typically enough to fit a jump model, whereas the EM algorithm needs between 50 and 100 iterations to converge.

3.2 Selecting the penalty

The jump penalty λ is selected based on the specific application, for example using cross-validation. The penalty aims to regularize the problem (1) by increasing the difference between states, similar to the purpose of amplifying the block structure of the affinity matrix as described in Section 2.2.

The jump penalty summarizes prior knowledge or assumptions about the transition rate. Note that a Bayesian approach would be significantly more demanding computationally. A strong prior is particularly important when a limited number of observations is available. A challenge when interested in fitting highly persistent regime-switching models is that it requires a lot of observations to ensure that all latent states are observed. Hence, a prior or regularization penalty should be imposed.

In many applications of the HMM, asset allocation included, it is a definite advantage of the jump estimator that it is possible to control the transition rate by adjusting the penalty parameter. The penalty parameter can be viewed as a hyperparameter that can be adjusted to achieve the desired behavior. In the case of a trading strategy, tuning the jump penalty is an alternative or supplement to controlling turnover by imposing transaction costs (Boyd et al., 2017; García-Galicia et al., 2019; Nystrup et al., 2019). The penalty that yields the most profitable strategy is not necessarily the same value that yields the most accurate estimate of the hidden state sequence. Whereas it would be extremely challenging to select the parameters in an HMM in order to maximize the return of a trading strategy, this is easily done for the jump penalty using cross-validation.

3.3 Variations

It is straightforward to add a regularization term on the parameters θ to (1), but it comes at the cost of not being able to compute analytically the optimal parameters for a given state sequence. Another variation would be to implement a fading memory by adding a decay term to the objective values. Although it is outside the scope of this article, jump model fitting can be extended to streaming applications, as proposed by Bemporad et al. (2018).

When penalizing jumps it is no longer necessary to construct an affinity matrix and compute its Laplacian. Better results can be achieved by fitting jump models directly to the standardized temporal features in Algorithm 2. It is still possible, though, to consider the first few eigenvectors of the Laplacian if dimensionality reduction is needed. Jump fitting could even be applied as a

post-processing step after spectral clustering.

Bemporad et al. (2018) showed how the jump framework can be used for fitting HMMs by replacing the least-squares criterion by the likelihood function and selecting the jump penalty based on transition probabilities. After all, simple features constructed from the data cannot add information compared to using the correctly specified likelihood function. However, when using the likelihood objective the choice of the jump penalty becomes critical and can no longer be viewed as a regularization parameter. For λ large enough, the model never changes state and the accuracy converges to the reciprocal of the number of states, similar to when using (1). When λ is too small, it becomes a mixture model with no temporal information and the accuracy, once again, converges to the reciprocal of the number of states. This is crucially different from our proposal of using (1) based on temporal features, where too small of a value of λ means that the outcome essentially converges to the result of spectral clustering.

Learning HMMs by minimizing the least-squares criterion (1) is more robust than likelihood maximization, because it does not rely on distributional assumptions. Moreover, the method is expected to be more robust to initialization and an increasing number of states, as was the case with spectral clustering (Zheng et al., 2019). The potential for adding exogenous variables that could improve latent state prediction is perhaps the most exciting prospect. Similar to when using spectral clustering, it is straightforward to include additional features, because these do not have to be multivariate Gaussian.

4 Simulation study

We compare the accuracy of HMMs estimated using maximum likelihood, spectral, and jump estimation in a simulation study. The advantage of a simulation study is that the true state sequence is known, which makes it possible to evaluate the ability of each model to correctly identify the underlying hidden states.

Accuracy is defined as the maximum classification accuracy among all possible alignments of an inferred state sequence. It is necessary to check whether a different permutation of the hidden states yields higher accuracy due to the risk of label switching. It is the ratio $\text{tp}/(\text{tp} + \text{fn})$, where tp is the number of true positives and fn the number of false negatives. We use the balanced accuracy (BAC), which is the average of accuracy per observed state, to avoid inflated performance estimates on imbalanced datasets (Brodersen et al., 2010). If a classifier performs equally well on either state, BAC reduces to conventional accuracy. In contrast, if accuracy is above chance only because a classifier takes advantage of an imbalanced dataset, then balanced accuracy, as appropriate, drops to the reciprocal of the number of states.

4.1 Two-state HMM

Following Zheng et al. (2019), we simulate data from a two-state Gaussian HMM

$$y_t | s_t \sim N(\mu_{s_t}, \sigma_{s_t}^2), \quad (4)$$

where s_t is a first-order Markov chain, with parameters

$$\begin{aligned}\mu_1 &= .0123, & \mu_2 &= -.0157, \\ \sigma_1 &= .0347, & \sigma_2 &= .0778, \\ \Gamma &= \begin{pmatrix} .9629 & .0371 \\ .2101 & .7899 \end{pmatrix}.\end{aligned}\tag{5}$$

This is a model Hardy (2001) estimated from monthly returns on a stock index. As shown by Rydén et al. (1998), this model captures the stylized behavior of many financial time series including volatility clustering and leptokurtosis. The significant overlap between the conditional distributions makes it challenging to infer the unobserved state sequence.

Each series is simulated by first generating a sequence of states of desired length according to the transition probability matrix Γ starting from its stationary distribution. Secondly, observations are sampled from conditional Gaussian distributions with mean and volatility parameters μ and σ given by the simulated states.

To simulate data from models with different levels of persistence in the state process, we assume the following relations between the model parameters at two time scales t_1 and t_2 :

$$\begin{aligned}\Gamma(t_1)^{1/t_1} &= \Gamma(t_2)^{1/t_2}, \\ \mu_s(t_1)/t_1 &= \mu_s(t_2)/t_2, \\ \sigma_s(t_1)/\sqrt{t_1} &= \sigma_s(t_2)/\sqrt{t_2},\end{aligned}$$

where $\Gamma(t)$, $\mu(t)$, and $\sigma(t)$ are the model parameters associated with time scale t . These relations are implied by assuming the instantaneous log return follows a Lévy process, i.e., a general continuous-time random walk with independent, stationary increments. The signal-to-noise ratio decreases with the time scale while persistence increases. In addition to the monthly base frequency $t = 20$ with parameters given in (5), we consider the weekly $t = 5$ and daily scale $t = 1$.

4.1.1 Selecting γ and λ

Figure 1 shows the BAC of spectral clustering as a function of the scaling parameter γ used in the Gaussian kernel function for different simulation lengths using the daily parameter values. The longer the simulated series, the higher the BAC. Across simulation lengths, any value $\gamma < 10^{-2}$ works. The corresponding figures for the weekly and monthly parameters look similar, with the only difference being the level of BAC. We use $\gamma = 10^{-3}$ throughout the simulation study, because the difference compared to using the optimal value for each time scale and simulation length is negligible. Using a common value across all simulation settings increases robustness of the results.

Figure 2 shows the BAC of jump estimation as a function of the penalty parameter λ for different simulation lengths using the daily parameter values. The main difference between simulation lengths is the level of BAC, with longer series leading to higher BAC. When the value of the penalty parameter is too large, the BAC approaches $1/2$. When the penalty parameter is too small, the BAC converges to that of the unpenalized case, which corresponds to applying K -means to the feature matrix. Across simulation lengths at the daily scale, the optimal value of the penalty parameter is around $\lambda = 10^2$. Without showing the corresponding figures for the weekly and monthly scales, we note that the optimal values are around $\lambda = 50$ and $\lambda = 1$, respectively. The less persistent the state process, the lower the optimal value of the penalty parameter and the

Figure 1: Balanced accuracy of spectral clustering as a function of the scaling parameter used in the Gaussian kernel function for different simulation lengths using the daily parameter values.

smaller the improvement in BAC that is achieved by penalizing jumps.

4.1.2 Daily scale

Table 1 shows the mean and standard deviation (in parentheses) of the estimated parameters based on 1000 series of different lengths simulated using the daily equivalent of the model (5). Bold entries identify the best estimate between maximum likelihood, spectral, and jump. The accuracy using the true parameters is based on the Viterbi path.

The reported MLEs are obtained using the regularized EM algorithm based on Huang et al. (2001, Chapter 9.6.1) that is implemented in the `hmmlearn` package in Python with the variance prior set to 10^{-4} . The MLE is reported for two different initializations: MLE_K is initialized using K -means with ten different centroid seeds generated using K -means++ (Arthur and Vassilvitskii, 2007), similar to how spectral clustering and jump estimation are initialized. MLE is initialized with a probability 1 of being in the first state, conditional means $\mu_1 = \mu_2 = 0$, standard deviations $\sigma_1 = \sigma_2 = 0.01$, and self-transition probabilities of 0.9.

We use the EM algorithm rather than direct numerical maximization of the likelihood function, because this is known to be more robust to initial values. Even so, at the shortest simulations lengths, the initialization makes a big difference in terms of the estimated transition probabilities and accuracy with MLE being much better than MLE_K . At the shortest simulation lengths, spectral clustering with $\gamma = 10^{-3}$ leads to much better estimates of the transition probabilities compared to MLE_K , but lower BACs. Jump estimation with $\lambda = 10^2$ yields more accurate parameter estimates compared to the other methods and a BAC inline with the best MLE.

It is arguably more fair to compare spectral clustering and jump estimation to MLE_K , which is initialized the same way and does not have the advantage of better initial values. The much higher standard deviation of the MLE_K parameter estimates is evidence of maximum likelihood estimation being less robust than spectral clustering and jump estimation.

The estimates of the daily parameters are plotted in Figure 3. It is clear from the figure that the MLE and the jump estimate are the fastest to converge to the true values as the simulation length

Figure 2: Balanced accuracy of jump estimation as a function of the penalty parameter for different simulation lengths using the daily parameter values.

	μ_1	μ_2	σ_1	σ_2	γ_{12}	γ_{21}	Accuracy 1	Accuracy 2	BAC
<i>250</i>									
True	.0006	-.0008	.0078	.0174	.0021	.0120	.9932 (.0404)	.8757 (.2469)	.9344 (.0587)
MLE _K	.0008 (.0022)	-.0001 (.0043)	.0075 (.0017)	.0119 (.0048)	.2889 (.2443)	.3694 (.2864)	.7689 (.2105)	.8101 (.2505)	.7895 (.0206)
MLE	.0007 (.0012)	.0000 (.0047)	.0077 (.0012)	.0123 (.0048)	.0453 (.0573)	.1180 (.1645)	.9250 (.1384)	.8686 (.2417)	.8968 (.0282)
Spectral	.0010 (.0024)	-.0006 (.0044)	.0073 (.0022)	.0121 (.0048)	.0933 (.0520)	.1605 (.0600)	.7384 (.1845)	.7405 (.1984)	.7394 (.0011)
Jump	.0007 (.0010)	.0000 (.0036)	.0079 (.0016)	.0126 (.0051)	.0056 (.0060)	.0278 (.0325)	.8766 (.1491)	.8772 (.2017)	.8769 (.0003)
<i>500</i>									
True	.0006	-.0008	.0078	.0174	.0021	.0120	.9968 (.0089)	.8662 (.2514)	.9315 (.0653)
MLE _K	.0007 (.0012)	-.0002 (.0022)	.0075 (.0009)	.0135 (.0047)	.2025 (.2334)	.2513 (.2619)	.8421 (.1992)	.8592 (.2076)	.8507 (.0086)
MLE	.0006 (.0006)	-.0004 (.0032)	.0077 (.0004)	.0138 (.0047)	.0303 (.0427)	.0762 (.1077)	.9591 (.0951)	.8770 (.2271)	.9181 (.0410)
Spectral	.0008 (.0018)	-.0008 (.0039)	.0072 (.0013)	.0136 (.0049)	.0770 (.0527)	.1518 (.0594)	.7874 (.1856)	.7734 (.1883)	.7804 (.0070)
Jump	.0006 (.0006)	-.0004 (.0029)	.0078 (.0004)	.0143 (.0048)	.0036 (.0029)	.0237 (.0249)	.9243 (.1214)	.8816 (.1903)	.9029 (.0214)
<i>1000</i>									
True	.0006	-.0008	.0078	.0174	.0021	.0120	.9962 (.0076)	.8993 (.1915)	.9478 (.0485)
MLE _K	.0006 (.0007)	-.0005 (.0017)	.0076 (.0004)	.0154 (.0039)	.1019 (.1881)	.1371 (.2165)	.9274 (.1475)	.8805 (.1817)	.9039 (.0235)
MLE	.0006 (.0004)	-.0005 (.0021)	.0077 (.0003)	.0159 (.0036)	.0157 (.0315)	.0395 (.0583)	.9820 (.0627)	.8921 (.2052)	.9370 (.0450)
Spectral	.0009 (.0014)	-.0014 (.0032)	.0073 (.0009)	.0152 (.0043)	.0565 (.0472)	.1434 (.0586)	.8467 (.1654)	.7774 (.1705)	.8120 (.0346)
Jump	.0006 (.0003)	-.0008 (.0025)	.0078 (.0003)	.0164 (.0035)	.0025 (.0015)	.0192 (.0176)	.9706 (.0772)	.8928 (.1623)	.9317 (.0389)
<i>2000</i>									
True	.0006	-.0008	.0078	.0174	.0021	.0120	.9965 (.0055)	.9264 (.1198)	.9614 (.0350)
MLE _K	.0006 (.0003)	-.0007 (.0011)	.0077 (.0003)	.0166 (.0025)	.0348 (.1077)	.0664 (.1480)	.9791 (.0759)	.9134 (.1273)	.9463 (.0329)
MLE	.0006 (.0002)	-.0007 (.0013)	.0078 (.0003)	.0169 (.0022)	.0071 (.0209)	.0229 (.0279)	.9946 (.0199)	.9036 (.1876)	.9491 (.0455)
Spectral	.0009 (.0007)	-.0017 (.0023)	.0074 (.0006)	.0164 (.0032)	.0387 (.0301)	.1421 (.0489)	.9021 (.1187)	.7918 (.1281)	.8470 (.0552)
Jump	.0006 (.0002)	-.0010 (.0017)	.0078 (.0002)	.0174 (.0018)	.0021 (.0010)	.0165 (.0123)	.9928 (.0246)	.9020 (.1161)	.9474 (.0454)

Table 1: Mean and standard deviation (in parentheses) of estimated parameters based on 1000 series of different lengths simulated using the daily parameter values. Bold entries identify the best estimate between MLE, spectral, and jump.

increases. The jump estimate particularly stands out when looking at the transition probabilities in the bottom row of the figure by being very close to the true values even for simulated series of only 250 observations. Even with 2000 observations per simulation, MLE_K and spectral clustering are far from estimating the true persistence of the state process.

When simulating a series of 250 observations using the daily parameter values, the probability that it only includes observations from one state is 51%. When simulating a series of length 500, the probability is 30%, for 1000 observations it is 10%, and for 2000 observations it is 1%.¹ In their simulation study, Zheng et al. (2019) filtered the simulated series to exclude those that only included observations from one state. If one state is unobserved in a sample, the maximum likelihood problem is nonregular. However, excluding the most persistent samples inevitably introduces a bias in the transition probabilities.

In considering regime-switching models with high persistence, it is an issue that it requires a lot of observations to ensure that all latent states are featured. In practical applications, for example when using a rolling window for estimation, it might not be the case that all states are observed. It is important how an estimator performs in that situation and, therefore, those series should not be excluded. It reflects negatively on the average BAC when fitting a two-state model based on data from one state only, but it is important that the estimator does not break down. To that end, the regularization imbedded in the jump estimator and the variance prior in the chosen EM algorithm are pivotal. For a sufficiently large penalty on jumps, the estimator has no problem identifying when all observations originate from the same conditional distribution.

Comparing our results when simulating series of 500 observations to those of Zheng et al. (2019), the MLE has a higher accuracy while spectral clustering has a lower accuracy. The higher accuracy of MLE in our study can be explained by the choice of the more robust EM algorithm instead of direct numerical maximization of the likelihood function. This may also explain their choice of filtering the simulated series to exclude those that only include observations from one state. The results in Table 1 show that there is no need for filtering the simulated data, since the BACs of MLE and jump are close to those of the Viterbi path using the true parameters.

4.1.3 Weekly scale

Table 2 shows the mean and standard deviation (in parentheses) of the estimated parameters based on 1000 series of different lengths simulated using the weekly equivalent of the model (5). Given the lower persistence of the state process at the weekly scale, it is not longer likely that a sequence only includes observations from one state, even if it only consists of 250 observations. After all, 250 weeks correspond to a period of five years. Results are shown for $\lambda = 50$. Naturally, the optimal value of the jump penalty is lower, because the state process is less persistent.

The lower persistence leads to lower BACs, meaning the hidden states are harder to come by. This is also true for the Viterbi path based on the true parameter values. In terms of BAC, MLE_K has the highest value at the longest simulation lengths, but only by a narrow margin over the other estimators. Compared to at the daily scale, the different estimation methods are better at estimating the transition probabilities, with the jump estimates still being by far the best. The BAC of MLE and jump is close to what can be achieved using the true parameter values.

¹The probabilities are calculated analytically using the fact that sojourn times are geometrically distributed.

Figure 3: Estimates of daily parameters based on 1000 simulations as a function of the simulation length using different estimation approaches.

	μ_1	μ_2	σ_1	σ_2	γ_{12}	γ_{21}	Accuracy 1	Accuracy 2	BAC
<i>250</i>									
True	.0031	-.0039	.0174	.0389	.0103	.0582	.9858 (.0267)	.6890 (.2999)	.8374 (.1484)
MLE _K	.0033 (.0033)	-.0030 (.0104)	.0166 (.0018)	.0334 (.0106)	.1272 (.1975)	.2472 (.2682)	.8945 (.1596)	.6997 (.2542)	.7971 (.0974)
MLE	.0031 (.0017)	-.0031 (.0113)	.0171 (.0013)	.0342 (.0100)	.0296 (.0409)	.1234 (.1478)	.9517 (.0915)	.7064 (.2883)	.8291 (.1226)
Spectral	.0038 (.0032)	-.0043 (.0104)	.0165 (.0021)	.0339 (.0104)	.0547 (.0443)	.1484 (.0628)	.8483 (.1531)	.7649 (.2114)	.8066 (.0417)
Jump	.0031 (.0018)	-.0034 (.0101)	.0177 (.0016)	.0358 (.0101)	.0090 (.0056)	.0509 (.0317)	.9291 (.1113)	.7455 (.2502)	.8373 (.0918)
<i>500</i>									
True	.0031	-.0039	.0174	.0389	.0103	.0582	.9865 (.0176)	.7587 (.1884)	.8726 (.1139)
MLE _K	.0032 (.0015)	-.0037 (.0071)	.0170 (.0010)	.0367 (.0071)	.0524 (.1200)	.1348 (.1729)	.9566 (.0959)	.7554 (.1869)	.8560 (.1006)
MLE	.0031 (.0010)	-.0040 (.0079)	.0172 (.0008)	.0372 (.0063)	.0165 (.0195)	.0892 (.0956)	.9787 (.0386)	.7544 (.2077)	.8665 (.1122)
Spectral	.0039 (.0020)	-.0062 (.0078)	.0166 (.0015)	.0358 (.0080)	.0433 (.0318)	.1516 (.0515)	.8890 (.1181)	.7576 (.1608)	.8233 (.0657)
Jump	.0032 (.0011)	-.0054 (.0072)	.0178 (.0010)	.0385 (.0061)	.0074 (.0034)	.0508 (.0227)	.9703 (.0558)	.7584 (.1754)	.8643 (.1060)
<i>1000</i>									
True	.0031	-.0039	.0174	.0389	.0103	.0582	.9879 (.0113)	.7732 (.1165)	.8805 (.1073)
MLE _K	.0031 (.0007)	-.0036 (.0038)	.0172 (.0006)	.0380 (.0040)	.0196 (.0441)	.0884 (.0968)	.9845 (.0217)	.7714 (.1322)	.8780 (.1065)
MLE	.0031 (.0006)	-.0037 (.0038)	.0173 (.0007)	.0381 (.0038)	.0129 (.0120)	.0726 (.0457)	.9861 (.0147)	.7650 (.1576)	.8755 (.1105)
Spectral	.0039 (.0012)	-.0069 (.0056)	.0168 (.0010)	.0371 (.0055)	.0357 (.0199)	.1573 (.0412)	.9194 (.0825)	.7515 (.1065)	.8355 (.0840)
Jump	.0032 (.0007)	-.0057 (.0053)	.0178 (.0006)	.0394 (.0035)	.0067 (.0025)	.0486 (.0171)	.9818 (.0194)	.7561 (.1203)	.8689 (.1129)
<i>2000</i>									
True	.0031	-.0039	.0174	.0389	.0103	.0582	.9882 (.0079)	.7912 (.0757)	.8897 (.0985)
MLE _K	.0031 (.0004)	-.0041 (.0024)	.0173 (.0003)	.0387 (.0021)	.0111 (.0053)	.0648 (.0289)	.9877 (.0090)	.7896 (.0829)	.8887 (.0991)
MLE	.0031 (.0005)	-.0040 (.0026)	.0174 (.0007)	.0384 (.0031)	.0125 (.0123)	.0644 (.0213)	.9879 (.0090)	.7763 (.1327)	.8821 (.1058)
Spectral	.0039 (.0006)	-.0080 (.0036)	.0169 (.0007)	.0379 (.0038)	.0323 (.0119)	.1635 (.0314)	.9382 (.0460)	.7427 (.0761)	.8404 (.0977)
Jump	.0032 (.0005)	-.0065 (.0032)	.0179 (.0005)	.0398 (.0023)	.0067 (.0017)	.0478 (.0118)	.9846 (.0080)	.7565 (.0870)	.8706 (.1140)

Table 2: Mean and standard deviation (in parentheses) of estimated parameters based on 1000 series of different lengths simulated using the weekly parameter values. Bold entries identify the best estimate between MLE, spectral, and jump.

4.1.4 Monthly scale

Table 3 shows the mean and standard deviation (in parentheses) of the estimated parameters based on 1000 series of different lengths simulated using the monthly parameters (5). Results are shown for $\lambda = 1$.

The relatively low persistence at the monthly scale leads to lower BACs. Spectral clustering and jump estimation are much better than MLE both in terms of parameter estimates and BAC. It is remarkable that both estimators yield a higher BAC than the Viterbi path based on the true parameter values. Even though the motivation for penalizing jumps is to improve estimates of the less persistent process at the monthly scale. It is evident that MLE and the Viterbi algorithm have a hard time predicting the second, less persistent state with accuracy for this state not much above 1/2.

4.2 Three-state HMM

It is well known that HMMs become considerably harder to estimate as the number of hidden states increases. To compare the different estimation methods when estimating a model with three hidden states, we follow Zheng et al. (2019) by augmenting the two-state model (5) with an intermediate third state:

$$\begin{aligned}
\mu_1 &= .0123, & \mu_2 &= .0000, & \mu_3 &= -.0157, \\
\sigma_1 &= .0347, & \sigma_2 &= .0500, & \sigma_3 &= .0778, \\
\Gamma &= \begin{pmatrix} .9629 & .0185 & .0186 \\ .0618 & .8764 & .0618 \\ .1051 & .1050 & .7899 \end{pmatrix}.
\end{aligned} \tag{6}$$

	μ_1	μ_2	σ_1	σ_2	γ_{12}	γ_{21}	Accuracy 1	Accuracy 2	BAC
250									
True	.0123	-.0157	.0347	.0778	.0371	.2101	.9845 (.0194)	.4914 (.1882)	.7380 (.2465)
MLE _K	.0127 (.0045)	-.0194 (.0303)	.0338 (.0032)	.0708 (.0176)	.0942 (.1257)	.3389 (.2338)	.9382 (.1037)	.5234 (.2089)	.7308 (.2074)
MLE	.0124 (.0032)	-.0187 (.0313)	.0344 (.0033)	.0711 (.0178)	.0513 (.0460)	.2651 (.1931)	.9544 (.0775)	.5157 (.2320)	.7351 (.2193)
Spectral	.0146 (.0049)	-.0158 (.0184)	.0332 (.0034)	.0676 (.0152)	.0539 (.0344)	.1735 (.0556)	.8489 (.1150)	.7029 (.1737)	.7759 (.0730)
Jump	.0143 (.0050)	-.0165 (.0176)	.0333 (.0031)	.0688 (.0154)	.0425 (.0233)	.1494 (.0471)	.8695 (.1072)	.7163 (.1568)	.7929 (.0766)
500									
True	.0123	-.0157	.0347	.0778	.0371	.2101	.9833 (.0136)	.5230 (.1230)	.7532 (.2301)
MLE _K	.0125 (.0021)	-.0171 (.0148)	.0343 (.0018)	.0758 (.0101)	.0495 (.0487)	.2541 (.1323)	.9729 (.0389)	.5266 (.1535)	.7498 (.2231)
MLE	.0124 (.0020)	-.0167 (.0147)	.0344 (.0019)	.0756 (.0102)	.0423 (.0222)	.2332 (.1091)	.9732 (.0358)	.5283 (.1587)	.7507 (.2225)
Spectral	.0145 (.0031)	-.0183 (.0129)	.0335 (.0024)	.0694 (.0107)	.0480 (.0225)	.1802 (.0431)	.8758 (.0811)	.6921 (.1240)	.7839 (.0919)
Jump	.0146 (.0029)	-.0193 (.0114)	.0335 (.0020)	.0703 (.0101)	.0390 (.0151)	.1583 (.0345)	.8946 (.0681)	.7197 (.1095)	.8071 (.0875)
1000									
True	.0123	-.0157	.0347	.0778	.0371	.2101	.9838 (.0095)	.5302 (.0873)	.7570 (.2268)
MLE _K	.0124 (.0014)	-.0160 (.0088)	.0345 (.0012)	.0769 (.0065)	.0409 (.0214)	.2262 (.0753)	.9794 (.0218)	.5337 (.1131)	.7566 (.2229)
MLE	.0123 (.0015)	-.0153 (.0094)	.0347 (.0019)	.0762 (.0081)	.0408 (.0160)	.2171 (.0691)	.9794 (.0177)	.5244 (.1373)	.7519 (.2275)
Spectral	.0146 (.0019)	-.0193 (.0080)	.0336 (.0016)	.0699 (.0076)	.0463 (.0149)	.1883 (.0318)	.8898 (.0534)	.6911 (.0870)	.7904 (.0994)
Jump	.0145 (.0017)	-.0206 (.0073)	.0336 (.0014)	.0710 (.0071)	.0374 (.0101)	.1645 (.0261)	.9082 (.0419)	.7147 (.0756)	.8114 (.0968)
2000									
True	.0123	-.0157	.0347	.0778	.0371	.2101	.9842 (.0065)	.5291 (.0648)	.7567 (.2275)
MLE _K	.0124 (.0009)	-.0161 (.0059)	.0346 (.0008)	.0772 (.0043)	.0379 (.0092)	.2172 (.0465)	.9827 (.0100)	.5305 (.0826)	.7566 (.2261)
MLE	.0122 (.0012)	-.0151 (.0075)	.0350 (.0022)	.0759 (.0076)	.0401 (.0151)	.2105 (.0499)	.9832 (.0103)	.5099 (.1316)	.7465 (.2367)
Spectral	.0146 (.0012)	-.0199 (.0056)	.0337 (.0011)	.0700 (.0051)	.0448 (.0103)	.1924 (.0222)	.8982 (.0343)	.6899 (.0607)	.7941 (.1041)
Jump	.0145 (.0012)	-.0213 (.0051)	.0337 (.0010)	.0710 (.0049)	.0367 (.0075)	.1683 (.0192)	.9137 (.0291)	.7129 (.0532)	.8133 (.1004)

Table 3: Mean and standard deviation (in parentheses) of estimated parameters based on 1000 series of different lengths simulated using the monthly parameter values. Bold entries identify the best estimate between MLE, spectral, and jump.

Models with more than two states are typically preferred when considering long series of returns (Dias et al., 2015; Nystrup et al., 2015b).

The figures corresponding to Figure 1 and Figure 2 look similar when simulating from a model with three states, except for the level of BAC. Therefore, we continue using $\gamma = 10^{-3}$ and $\lambda = 10^2$. Table 4 shows the mean and standard deviation (in parentheses) of the estimated parameters based on 1000 series of different lengths simulated using the daily equivalent of the three-state model (6). We do not show results for series of 250 observations, as this is too few to estimate a three-state model with this level of persistence.

As expected, adding an intermediate third state has a very negative effect on the BAC. With three states the difference between the two initializations of MLE is decisive. When MLE is initialized using K -means—similar to how spectral clustering and the jump estimator are initialized—the results are terrible. MLE_K works poorly because the transition probability matrix never converges to something meaningful, which is why the order of the states appears to be wrong despite the presented results being for the most accurate permutation. Interestingly, the BACs of MLE_K and MLE decrease when the simulation length increases.

MLE requires much better initial values to contend with the two other estimators. That is despite having the advantage of knowing the true underlying model and even if 2000 observations are available for estimation. The challenges are most evident when looking at the transition probabilities where, once again, the jump estimates are by far the best. Even for the jump estimator, the BAC is much lower than what is achieved using the true parameter values.

4.3 Misspecified conditional distributions

In order to compare the different estimation methods when conditional distributions are misspecified, we repeat the two-state simulation study, only this time we sample the observations from Student’s t rather than Gaussian distributions. The conditional distributions have the same mean and standard deviation, but are t -distributions with five degrees of freedom, which have heavier

	μ_1	μ_2	μ_3	σ_1	σ_2	σ_3	γ_{12}	γ_{13}
<i>500</i>								
True	.0006	.0000	-.0008	.0078	.0112	.0174	.0009	.0010
MLE _K	.0003 (.0030)	.0003 (.0035)	.0001 (.0035)	.0094 (.0038)	.0093 (.0030)	.0104 (.0039)	.2256 (.1843)	.2115 (.1795)
MLE	.0005 (.0009)	.0001 (.0027)	-.0002 (.0022)	.0087 (.0027)	.0114 (.0037)	.0126 (.0044)	.0219 (.0302)	.0243 (.0233)
Spectral	.0007 (.0038)	-. .0001 (.0063)	-.0001 (.0067)	.0068 (.0026)	.0096 (.0033)	.0123 (.0048)	.0938 (.0483)	.0712 (.0465)
Jump	.0005 (.0026)	-. .0001 (.0048)	-. .0004 (.0043)	.0088 (.0031)	.0115 (.0041)	.0134 (.0049)	.0040 (.0089)	.0050 (.0114)
<i>1000</i>								
True	.0006	.0000	-.0008	.0078	.0112	.0174	.0009	.0010
MLE _K	.0001 (.0018)	.0003 (.0019)	.0004 (.0020)	.0115 (.0042)	.0091 (.0029)	.0093 (.0030)	.1290 (.1628)	.1498 (.1790)
MLE	.0005 (.0006)	-.0001 (.0016)	-.0002 (.0022)	.0084 (.0021)	.0125 (.0036)	.0134 (.0041)	.0156 (.0222)	.0181 (.0229)
Spectral	.0006 (.0031)	.0006 (.0066)	-. .0008 (.0068)	.0065 (.0017)	.0098 (.0030)	.0131 (.0047)	.0947 (.0484)	.0574 (.0423)
Jump	.0006 (.0018)	.0000 (.0049)	-.0004 (.0041)	.0083 (.0021)	.0124 (.0036)	.0150 (.0044)	.0024 (.0051)	.0028 (.0060)
<i>2000</i>								
True	.0006	.0000	-.0008	.0078	.0112	.0174	.0009	.0010
MLE _K	-.0003 (.0012)	.0005 (.0011)	.0005 (.0012)	.0133 (.0034)	.0087 (.0023)	.0085 (.0022)	.0518 (.1082)	.0796 (.1487)
MLE	.0006 (.0003)	-.0002 (.0014)	-.0003 (.0011)	.0081 (.0014)	.0132 (.0030)	.0139 (.0034)	.0108 (.0188)	.0122 (.0203)
Spectral	.0006 (.0028)	.0004 (.0069)	-. .0009 (.0069)	.0065 (.0008)	.0096 (.0026)	.0141 (.0038)	.1005 (.0482)	.0445 (.0351)
Jump	.0006 (.0010)	-. .0001 (.0046)	-.0010 (.0039)	.0082 (.0014)	.0131 (.0030)	.0165 (.0030)	.0016 (.0026)	.0016 (.0031)
	γ_{21}	γ_{23}	γ_{31}	γ_{32}	Accuracy 1	Accuracy 2	Accuracy 3	BAC
<i>500</i>								
True	.0032	.0037	.0058	.0062	.9534 (.1737)	.7417 (.3857)	.7734 (.3464)	.8228 (.0932)
MLE _K	.2594 (.2129)	.2784 (.1870)	.2478 (.1990)	.2844 (.1830)	.6250 (.2832)	.3879 (.3587)	.3415 (.4017)	.4515 (.1242)
MLE	.0370 (.0507)	.0444 (.0553)	.0346 (.0340)	.0381 (.0312)	.8726 (.2790)	.5957 (.4203)	.6886 (.3996)	.7190 (.1151)
Spectral	.1318 (.0493)	.0999 (.0547)	.1232 (.0519)	.1176 (.0475)	.5912 (.2264)	.4583 (.1926)	.5634 (.1963)	.5377 (.0572)
Jump	.0123 (.0205)	.0158 (.0257)	.0142 (.0210)	.0154 (.0231)	.7839 (.2811)	.6035 (.3631)	.6773 (.2635)	.6883 (.0741)
<i>1000</i>								
True	.0032	.0037	.0058	.0062	.9727 (.1021)	.7545 (.3489)	.8205 (.2829)	.8492 (.0913)
MLE _K	.1491 (.1969)	.3399 (.1895)	.1644 (.1972)	.3475 (.1919)	.7457 (.3094)	.2527 (.3732)	.1628 (.3379)	.3871 (.2562)
MLE	.0268 (.0403)	.0381 (.0322)	.0292 (.0482)	.0387 (.0362)	.9007 (.2446)	.5177 (.4253)	.6385 (.4185)	.6856 (.1599)
Spectral	.1491 (.0432)	.0942 (.0525)	.1193 (.0439)	.1207 (.0465)	.6290 (.2002)	.4218 (.1721)	.5644 (.1979)	.5384 (.0866)
Jump	.0112 (.0153)	.0154 (.0209)	.0126 (.0157)	.0157 (.0169)	.8777 (.2299)	.5467 (.3771)	.6612 (.2478)	.6952 (.1373)
<i>2000</i>								
True	.0032	.0037	.0058	.0062	.9834 (.0355)	.8101 (.2732)	.8453 (.2234)	.8796 (.0748)
MLE _K	.0613 (.1422)	.4085 (.1686)	.0889 (.1731)	.4132 (.1721)	.8662 (.2863)	.1360 (.3157)	.0681 (.2313)	.3568 (.3613)
MLE	.0172 (.0266)	.0356 (.0229)	.0190 (.0206)	.0371 (.0253)	.9304 (.1794)	.4930 (.3787)	.6171 (.3983)	.6802 (.1840)
Spectral	.1627 (.0359)	.0822 (.0500)	.1170 (.0395)	.1233 (.0374)	.6608 (.1671)	.3969 (.1381)	.5685 (.1805)	.5421 (.1094)
Jump	.0095 (.0112)	.0143 (.0170)	.0108 (.0104)	.0164 (.0137)	.9484 (.1628)	.5324 (.3752)	.6590 (.2301)	.7132 (.1741)

Table 4: Mean and standard deviation (in parentheses) of estimated parameters based on 1000 series of different lengths simulated using the daily parameter values. Bold entries identify the best estimate between MLE, spectral, and jump.

	μ_1	μ_2	σ_1	σ_2	γ_{12}	γ_{21}	Accuracy 1	Accuracy 2	BAC
250									
True	.0006	-.0008	.0078	.0174	.0021	.0120	.9802 (.0750)	.7968 (.3200)	.8885 (.0917)
MLE _K	.0005 (.0019)	-.0005 (.0121)	.0071 (.0046)	.0153 (.0056)	.2041 (.2018)	.5197 (.3428)	.8600 (.1949)	.6690 (.2947)	.7645 (.0955)
MLE	.0006 (.0014)	-.0003 (.0065)	.0072 (.0044)	.0155 (.0066)	.1137 (.1447)	.4090 (.3566)	.9020 (.1619)	.7350 (.2840)	.8185 (.0835)
Spectral	.0007 (.0024)	-.0006 (.0066)	.0075 (.0044)	.0140 (.0080)	.0715 (.0480)	.1635 (.0710)	.7653 (.1893)	.6568 (.2524)	.7111 (.0542)
Jump	.0005 (.0024)	-.0004 (.0057)	.0084 (.0050)	.0149 (.0077)	.0082 (.0191)	.0421 (.0492)	.8778 (.1840)	.7530 (.3130)	.8154 (.0624)
500									
True	.0006	-.0008	.0078	.0174	.0021	.0120	.9873 (.0217)	.8361 (.2590)	.9117 (.0756)
MLE _K	.0006 (.0007)	.0003 (.0093)	.0067 (.0027)	.0165 (.0046)	.1238 (.1469)	.4156 (.3462)	.9199 (.1307)	.7080 (.2788)	.8140 (.1060)
MLE	.0006 (.0006)	-.0001 (.0033)	.0068 (.0024)	.0167 (.0044)	.0830 (.1014)	.3621 (.3428)	.9395 (.1083)	.7465 (.2687)	.8430 (.0965)
Spectral	.0007 (.0016)	-.0004 (.0051)	.0073 (.0030)	.0153 (.0071)	.0563 (.0379)	.1580 (.0705)	.8104 (.1586)	.6750 (.2445)	.7427 (.0677)
Jump	.0006 (.0016)	-.0004 (.0049)	.0079 (.0033)	.0171 (.0066)	.0043 (.0087)	.0395 (.0430)	.9466 (.1154)	.7869 (.2608)	.8668 (.0799)
1000									
True	.0006	-.0008	.0078	.0174	.0021	.0120	.9875 (.0154)	.8651 (.2033)	.9263 (.0612)
MLE _K	.0006 (.0003)	-.0009 (.0181)	.0066 (.0008)	.0172 (.0034)	.0749 (.0935)	.3104 (.3083)	.9524 (.0605)	.7196 (.2619)	.8360 (.1164)
MLE	.0006 (.0003)	-.0004 (.0018)	.0067 (.0007)	.0174 (.0039)	.0608 (.0702)	.2827 (.2976)	.9605 (.0393)	.7293 (.2684)	.8449 (.1156)
Spectral	.0007 (.0009)	-.0010 (.0047)	.0072 (.0014)	.0160 (.0081)	.0462 (.0294)	.1469 (.0695)	.8268 (.1397)	.6894 (.2289)	.7581 (.0687)
Jump	.0006 (.0007)	-.0012 (.0058)	.0078 (.0012)	.0188 (.0097)	.0030 (.0033)	.0332 (.0354)	.9769 (.0585)	.7900 (.2421)	.8835 (.0934)
2000									
True	.0006	-.0008	.0078	.0174	.0021	.0120	.9892 (.0091)	.8958 (.1264)	.9425 (.0467)
MLE _K	.0006 (.0002)	.0000 (.0117)	.0067 (.0004)	.0177 (.0024)	.0466 (.0525)	.2325 (.2404)	.9682 (.0250)	.7453 (.2290)	.8567 (.1114)
MLE	.0006 (.0002)	-.0004 (.0012)	.0068 (.0004)	.0178 (.0026)	.0406 (.0410)	.2146 (.2266)	.9702 (.0181)	.7590 (.2241)	.8646 (.1056)
Spectral	.0008 (.0005)	-.0013 (.0174)	.0071 (.0011)	.0163 (.0074)	.0390 (.0204)	.1418 (.0788)	.8382 (.1380)	.7010 (.1993)	.7696 (.0686)
Jump	.0007 (.0002)	-.0010 (.0045)	.0077 (.0003)	.0188 (.0088)	.0025 (.0010)	.0249 (.0250)	.9885 (.0132)	.8234 (.1712)	.9060 (.0826)

Table 5: Mean and standard deviation (in parentheses) of estimated parameters based on 1000 series of different lengths simulated from conditional t -distributions with five degrees of freedom using the daily parameter values. Bold entries identify the best estimate between MLE, spectral, and jump.

tails than the Gaussian distribution. Several studies have found that this is in fact a better fit to the empirical distribution of financial returns (Cont, 2001; Bulla, 2011; Nystrup et al., 2017b).

Table 5 shows the mean and standard deviation (in parentheses) of the estimated parameters based on 1000 series of different lengths simulated from conditional t -distributions with five degrees of freedom using the daily parameter values. The accuracy using the true parameter values is based on the Viterbi path under the erroneous assumption that conditional distributions are Gaussian.

Compared to Table 1, the BACs are around five percentage points lower as a result of the misspecified conditionals. When simulating series of 250 observations, the BAC of MLE with good initial values is similar to that of the jump estimator; but when simulating longer series the jump estimator has the highest BAC. MLE_K yields significantly worse estimates than MLE at the shortest simulation lengths, once again emphasizing the method’s lack of robustness.

The jump estimates of the transition probabilities stand out compared to the other methods. This is even more evident when looking at the plots of the estimated parameters in Figure 4. Misspecification of the conditional distributions causes a severe upward bias in the estimated transition probabilities for the other methods, leading to unrealistically rapid switching dynamics. Meanwhile, the jump estimator has no difficulty in estimating the transition probabilities. As a result of the jump estimator being very close to the true values, its BAC is relatively close to what can be achieved using the true parameter values. Spectral clustering has a lower BAC than the other methods, but provides better estimates of the transition probabilities compared to MLE, regardless of the initialization.

5 Trading strategy

Several studies have shown the potential benefit from changing asset allocation depending on regimes inferred from financial returns (see, e.g., Guidolin and Timmermann, 2007; Bulla et al.,

Figure 4: Estimates of daily parameters based on 1000 simulations from conditional t -distributions with five degrees of freedom as a function of the simulation length using different estimation approaches.

	Return	Risk	Sharpe ratio	Turnover	Break-even
MLE	.1120	.1767	.6338	21.52	.0025
Spectral	.0899	.1878	.4784	17.21	.0025
Jump	.1271	.1836	.6923	4.55	.0131

Table 6: Summary of long–short strategy based on 1000 simulations using the last 250 out of 1000 observations simulated from conditional t -distributions with five degrees of freedom using the daily parameters. Break-even is the cost per transaction that would offset the return.

2011; Bae et al., 2014; Nystrup et al., 2015a, 2017a, 2018, 2019). The HMM is a popular choice for inferring the hidden state of financial markets, because it can match the tendency of financial markets to change their behavior abruptly and the phenomenon that the new behavior often persists for several periods after a change (Ang and Timmermann, 2012). To illustrate the value of estimating the true persistence of the underlying state process, we consider a simple long–short trading strategy based on the inferred states.

5.1 Simulated data

We first consider a simulation example based on the misspecified two-state HMM from the previous section. We reuse the 1000 series of 1000 observations that were simulated using the daily parameter values, but sampled from conditional t -distributions with five degrees of freedom. Long and short positions are implemented for the last 250 observations based on the inferred states. In the state with positive mean and low volatility the strategy is long, and in the other state with negative return and higher volatility the strategy is short the underlying stock index.

Allocating for one year based on three years of data is a realistic setup. We fit the model to the first 750 observations using the different methods and then apply the Viterbi algorithm sequentially to the next 250 observations for state decoding. Thus, differences in the inferred state sequences are entirely due to the differences in the estimated parameter values. An important advantage of spectral clustering and jump estimation is the simultaneous estimation of parameters and state sequence; however, it is outside the scope of this article to extend these methods to streaming applications.

Table 6 shows the results when taking long and short positions based on the states inferred using the different parameter estimates. Results are only shown for MLE as the two initializations yield relatively similar results in this case. Spectral clustering leads to a lower return with similar risk as the other methods and, consequently, a lower Sharpe ratio—i.e., ratio between return and risk. Risk is measured as the standard deviation of the 1000 annualized returns. The turnover using the spectral parameter estimates is lower than that using the MLE. The two methods have the same break-even transaction cost, which is the cost per transaction that would offset the return.

The long–short strategy based on the states inferred using the jump parameter estimates yields a slightly higher return, risk, and Sharpe ratio than that using the MLE, while incurring a substantially lower turnover. This amounts to a break-even transaction cost that is five times higher (131 basis points), clearly illustrating the value of estimating the true persistence of the underlying state process. Had the simulation been based on series that were shorter than 1000 observations, the difference would have been even bigger.

Figure 5: S&P 500 index values and log returns from 2009 through 2018.

	μ_1	μ_2	σ_1	σ_2	γ_{12}	γ_{21}
MLE	.0011	-.0005	.0054	.0154	.0270	.0430
Spectral	.0011	-.0025	.0069	.0200	.0265	.1281
Jump	.0008	-.0015	.0092	.0268	.0033	.0182

Table 7: Estimated parameters for a two-state Gaussian HMM fitted to the log returns of the S&P 500 index from 2009 through 2018.

5.2 Application to S&P 500

Finally, we consider an application to real data using daily values of the S&P 500 total return index from 2009 through 2018. Figure 5 shows the index and its log returns over the 20-year period. This period after the financial crisis in 2008 is special in that it does not include a major drawdown, which is in sharp contrast to the ten-year period that preceded. As a consequence, it is harder for a regime-based strategy to outperform the index, because it is unable to benefit from prolonged periods of drawdown. This makes it an interesting period to study.

Table 7 shows the estimated parameters for a two-state Gaussian HMM fitted to the log returns of the S&P 500 index. Once again we use $\gamma = 10^{-3}$ and $\lambda = 10^2$, though we could have chosen the values that yield the highest return or Sharpe ratio rather than the highest BAC. Results are only shown for MLE as the two initializations yield almost similar results. The estimated parameters are fairly similar to those used in the simulation study in Section 4. The higher values for the transition probabilities of MLE and spectral compared to jump could indicate that the conditional distributions have heavier tails than the assumed Gaussian distribution.

Table 8 shows the result when taking long and short positions, respectively, in the S&P 500 index based on the states inferred using the different methods. The purpose is to compare the results using the different estimation methods and not to conclude whether a trading strategy is profitable or not. The latter would require out-of-sample testing. The results confirm the conclusions from the simulation study. Although the turnover of the MLE is lower than in Table 6, it is still above that of the jump estimate, which has by far the highest break-even transaction cost. Over the 20-year period, the Sharpe ratio of the jump estimator compares favorably with that of the S&P 500 index, which is even far above its long-term average.

	Return	Risk	Sharpe ratio	Turnover	Break-even
S&P 500	.1114	.1657	.6726		
MLE	.1529	.1656	.9234	4.175	.0169
Spectral	.2181	.1653	1.319	11.03	.0089
Jump	.2370	.1652	1.434	1.392	.0736

Table 8: Summary of long-short strategy based on the S&P 500 from 2009 through 2018. Break-even is the cost per transaction that would offset the return.

6 Conclusion

From a theoretical point of view, the MLE has favorable properties; however, it required accurate initial values and correctly specified conditional distributions to obtain those properties in practice. The spectral clustering approach to estimating HMMs had some merit in leading to improved estimates of transition probabilities while being less sensitive to initialization compared to the MLE, but it generally led to lower accuracy when the underlying states were highly persistent. The MLE required considerably better initial values in order to compete with the proposed jump estimate based on temporal features and penalizing jumps. That alone is a strong argument in favor of the jump estimator, since the required prior knowledge is often not available in practical applications.

It was convincing that jump estimation, which required much fewer iterations, could compete with maximizing the correctly specified likelihood function. Even when initialized with sufficiently accurate values, maximizing the likelihood at best yielded the same accuracy as the jump estimator for highly persistent processes. Meanwhile, the jump estimates of transition probabilities were more precise, especially when the samples available for estimation were short. When the hidden state process was less persistent, the jump estimator, as well as spectral clustering, yielded significantly higher accuracy compared to the MLE and, in some cases, even higher than the Viterbi path based on the true parameters.

As the number of latent states increased, the initialization became even more critical to the performance of the MLE. Only if initial values were sufficiently close to the true parameter values, was it able to contend with the jump estimator; but it still led to worse estimates of transition probabilities and oftentimes lower accuracy. When simulating from conditional distributions with heavier tails than the assumed Gaussian distribution, the jump estimator yielded much better results than the other methods considered. The misspecification caused a severe bias in the estimated transition probabilities for the other methods, which led to unrealistically rapid switching dynamics, whereas the jump estimator was not susceptible.

The value of estimating the true persistence of the hidden state process was illustrated through applications to long-short trading strategies. The improved estimates of the latent state process resulted in a substantially lower turnover and higher break-even transaction costs.

Compared to maximization of the likelihood function using the EM algorithm, the advantages of the proposed jump estimator have been shown to include that it learns the hidden state sequence and model parameters simultaneously while providing control over the transition rate, it takes much fewer iterations to converge, it is less sensitive to initial values, it performs better when the number of hidden states increases, and it is more robust to heavy-tailed conditional distributions because it does not assume Gaussianity.

It is left for future research to consider feature selection in depth. The main limitation of the

spectral clustering and jump approaches is the forward-looking nature of some of the temporal features used for estimation. This needs to be addressed in future research. At the same time, the possibility of extending the feature space to include additional features that can improve state inference is perhaps the most exciting prospect for future work. This can be easily done because the features do not have to multivariate Gaussian. Moreover, the computational advantage of the jump estimator could be useful in considering data with a frequency higher than daily.

Replacing the forward-looking features is a prerequisite to developing a jump estimator that can be applied in a streaming setting. Developing a streaming implementation is particularly important for the model’s application in real-life expert and intelligent systems. Unlike spectral clustering, which happens in the frequency domain, by clustering in the time domain the jump estimator can more easily be extended to a streaming setting. The demonstrated robustness of the jump-penalization approach shows promise for the possibility of extending it to an online estimator.

Acknowledgments

This work was supported by Vergstiftelsen and the Centre for IT-Intelligent Energy Systems (CITIES) project funded in part by Innovation Fund Denmark under Grant No. 1305-00027B. We are thankful to Stephen Boyd for introducing us to the framework for fitting jump models, to Kai Zheng for clarifying details about the simulation study in Zheng et al. (2019), and to two reviewers for valuable feedback.

References

- Amini, A. A., A. Chen, P. J. Bickel, and E. Levina. “Pseudo-likelihood methods for community detection in large sparse networks.” *Annals of Statistics*, vol. 41, no. 4 (2013), pp. 2097–2122.
- Andersson, S., T. Rydén, and R. Johansson. “Linear optimal prediction and innovations representations of hidden Markov models.” *Stochastic Processes and their Applications*, vol. 108, no. 1 (2003), pp. 131–149.
- Ang, A. and A. Timmermann. “Regime changes and financial markets.” *Annual Review of Financial Economics*, vol. 4, no. 1 (2012), pp. 313–337.
- Arthur, D. and S. Vassilvitskii. “k-means++: The advantages of careful seeding.” In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms* (2007), pp. 1027–1035.
- Bach, F. R. and M. I. Jordan. “Learning spectral clustering.” In *Advances in Neural Information Processing Systems* (2004), pp. 305–312.
- Bae, G. I., W. C. Kim, and J. M. Mulvey. “Dynamic asset allocation for varied financial markets under regime switching framework.” *European Journal of Operational Research*, vol. 234, no. 2 (2014), pp. 450–458.
- Bakis, R. “Continuous speech recognition via centisecond acoustic states.” *Journal of the Acoustical Society of America*, vol. 59, no. S1 (1976), p. S97.
- Baum, L. E., T. Petrie, G. Soules, and N. Weiss. “A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains.” *Annals of Mathematical Statistics*, vol. 41, no. 1 (1970), pp. 164–171.

- Beal, M. J., Z. Ghahramani, and C. E. Rasmussen. “The infinite hidden Markov model.” In *Advances in Neural Information Processing Systems* (2002), pp. 577–584.
- Bellman, R. E. *Dynamic programming*. Princeton University Press: Princeton (1957).
- Bemporad, A., V. Breschi, D. Piga, and S. Boyd. “Fitting jump models.” *Automatica*, vol. 96 (2018), pp. 11–21.
- Bhardwaj, S., V. Sharma, S. Srivastava, O. Sastry, B. Bandyopadhyay, S. Chandel, and J. Gupta. “Estimation of solar radiation using a combination of hidden Markov model and generalized fuzzy model.” *Solar Energy*, vol. 93 (2013), pp. 43–54.
- Blanco, I. and J. M. Morales. “An efficient robust solution to the two-stage stochastic unit commitment problem.” *IEEE Transactions on Power Systems*, vol. 32, no. 6 (2017), pp. 4477–4488.
- Box, G. E. P. “Robustness in the strategy of scientific model building.” In *Robustness in Statistics*, edited by R. L. Launer and G. N. Wilkinson. Academic Press: New York (1979), pp. 201–236.
- Boyd, S., E. Busseti, S. Diamond, R. N. Kahn, K. Koh, P. Nysttrup, and J. Speth. “Multi-period trading via convex optimization.” *Foundations and Trends in Optimization*, vol. 3, no. 1 (2017), pp. 1–76.
- Brodersen, K. H., C. S. Ong, K. E. Stephan, and J. M. Buhmann. “The balanced accuracy and its posterior distribution.” In *20th International Conference on Pattern Recognition* (2010), pp. 3121–3124.
- Bulla, J. “Hidden Markov models with t components. Increased persistence and other aspects.” *Quantitative Finance*, vol. 11, no. 3 (2011), pp. 459–475.
- Bulla, J. and I. Bulla. “Stylized facts of financial time series and hidden semi-Markov models.” *Computational Statistics & Data Analysis*, vol. 51, no. 4 (2006), pp. 2192–2209.
- Bulla, J., S. Mergner, I. Bulla, A. Sesboüé, and C. Chesneau. “Markov-switching asset allocation: Do profitable strategies exist?” *Journal of Asset Management*, vol. 12, no. 5 (2011), pp. 310–321.
- Chaudhuri, K., F. Chung, and A. Tsiatas. “Spectral clustering of graphs with general degrees in the extended planted partition model.” *Journal of Machine Learning Research*, vol. 35 (2012), pp. 1–23.
- Choo, K. H., J. C. Tong, and L. Zhang. “Recent applications of hidden Markov models in computational biology.” *Genomics, Proteomics & Bioinformatics*, vol. 2, no. 2 (2004), pp. 84–96.
- Cont, R. “Empirical properties of asset returns: stylized facts and statistical issues.” *Quantitative Finance*, vol. 1, no. 2 (2001), pp. 223–236.
- Davis, R. I. A. and B. C. Lovell. “Comparing and evaluating HMM ensemble training algorithms using train and test and condition number criteria.” *Pattern Analysis & Applications*, vol. 6, no. 4 (2004).
- Dempster, A. P., N. M. Laird, and D. B. Rubin. “Maximum likelihood from incomplete data via the EM algorithm.” *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 39, no. 1 (1977), pp. 1–38.

- Dias, J. G. and S. B. Ramos. “Dynamic clustering of energy markets: An extended hidden Markov approach.” *Expert Systems with Applications*, vol. 41, no. 17 (2014), pp. 7722–7729.
- Dias, J. G., J. K. Vermunt, and S. B. Ramos. “Clustering financial time series: New insights from an extended hidden Markov model.” *European Journal of Operational Research*, vol. 243, no. 3 (2015), pp. 852–864.
- Fiecas, M., J. Franke, R. von Sachs, and J. T. Kamgaing. “Shrinkage estimation for multivariate hidden Markov models.” *Journal of the American Statistical Association*, vol. 112, no. 517 (2017), pp. 424–435.
- Fischer, I. and J. Poland. “Amplifying the block matrix structure for spectral clustering.” In *Proceedings of the 14th Annual Machine Learning Conference of Belgium and the Netherlands* (2005), pp. 21–28.
- Fox, E. B., E. B. Sudderth, M. I. Jordan, and A. S. Willsky. “A sticky HDP-HMM with application to speaker diarization.” *Annals of Applied Statistics*, vol. 5, no. 2A (2011), pp. 1020–1056.
- Gales, M. and S. Young. “The application of hidden Markov models in speech recognition.” *Foundations and Trends in Signal Processing*, vol. 1, no. 3 (2008), pp. 195–304.
- García-Galicia, M., A. A. Carsteanu, and J. B. Clempner. “Continuous-time reinforcement learning approach for portfolio management with time penalization.” *Expert Systems with Applications*, vol. 129 (2019), pp. 27–36.
- Ghahramani, Z. and M. I. Jordan. “Factorial hidden Markov models.” *Machine Learning*, vol. 29, no. 2–3 (1997), pp. 245–273.
- Guidolin, M. and A. Timmermann. “Asset allocation under multivariate regime switching.” *Journal of Economic Dynamics and Control*, vol. 31, no. 11 (2007), pp. 3503–3544.
- Gustafsson, F. *Adaptive Filtering and Change Detection*. Wiley: West Sussex (2000).
- Hallac, D., P. Nystrup, and S. Boyd. “Greedy Gaussian segmentation of multivariate time series.” *Advances in Data Analysis and Classification*, vol. 13, no. 3 (2019), pp. 727–751.
- Hardy, M. R. “A regime-switching model of long-term stock returns.” *North American Actuarial Journal*, vol. 5, no. 2 (2001), pp. 41–53.
- Hsu, D., S. M. Kakade, and T. Zhang. “A spectral algorithm for learning hidden Markov models.” *Journal of Computer and System Sciences*, vol. 78, no. 5 (2012), pp. 1460–1480.
- Huang, X., A. Acero, and H. W. Hon. *Spoken Language Processing: A Guide to Theory, Algorithm and System Development*. Prentice-Hall: New Jersey (2001).
- Jia, H., S. Ding, X. Xu, and R. Nie. “The latest research progress on spectral clustering.” *Neural Computing and Applications*, vol. 24, no. 7-8 (2013), pp. 1477–1486.
- Jiang, Y. and J. Ren. “Eigenvector sensitive feature selection for spectral clustering.” In *Machine Learning and Knowledge Discovery in Databases*, edited by D. Gunopulos, T. Hofmann, D. Malerba, and M. Vazirgiannis. Springer: Berlin (2011), pp. 114–129.

- Jin, R., F. Kang, and C. H. Ding. “A probabilistic approach for optimizing spectral clustering.” In *Advances in Neural Information Processing Systems* (2006), pp. 571–578.
- Johnson, M. J. and A. S. Willsky. “Bayesian nonparametric hidden semi-Markov models.” *Journal of Machine Learning Research*, vol. 14 (2013), pp. 673–701.
- Joseph, A. and B. Yu. “Impact of regularization on spectral clustering.” *Annals of Statistics*, vol. 44, no. 4 (2016), pp. 1765–1791.
- Kang, M., J. Ahn, and K. Lee. “Opinion mining using ensemble text hidden Markov models for text classification.” *Expert Systems with Applications*, vol. 94 (2018), pp. 218–227.
- Katz, I. and K. Crammer. “Outlier-robust convex segmentation.” In *Twenty-Ninth AAAI Conference on Artificial Intelligence* (2015).
- Kim, S. J., K. Koh, S. Boyd, and D. Gorinevsky. “ ℓ_1 trend filtering.” *SIAM Review*, vol. 51, no. 2 (2009), pp. 339–360.
- Langrock, R. and W. Zucchini. “Hidden Markov models with arbitrary state dwell-time distributions.” *Computational Statistics & Data Analysis*, vol. 55, no. 1 (2011), pp. 715–724.
- Lloyd, S. P. “Least squares quantization in PCM.” *IEEE Transactions on Information Theory*, vol. 28, no. 2 (1982), pp. 129–137.
- Meila, M. and J. Shi. “Learning segmentation by random walks.” In *Advances in Neural Information Processing Systems* (2001), pp. 873–879.
- Nefian, A. V. and M. H. Hayes. “Hidden Markov models for face recognition.” In *Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing*, vol. 5. IEEE (1998), pp. 2721–2724.
- Ng, A. Y., M. I. Jordan, and Y. Weiss. “On spectral clustering: Analysis and an algorithm.” In *Advances in Neural Information Processing Systems* (2002), pp. 849–856.
- Nielsen, M. G., J. M. Morales, M. Zugno, T. E. Pedersen, and H. Madsen. “Economic valuation of heat pumps and electric boilers in the Danish energy system.” *Applied Energy*, vol. 167 (2016), pp. 189–200.
- Nystrup, P., S. Boyd, E. Lindström, and H. Madsen. “Multi-period portfolio selection with draw-down control.” *Annals of Operations Research*, vol. 282, no. 1–2 (2019), pp. 245–271.
- Nystrup, P., B. W. Hansen, H. O. Larsen, H. Madsen, and E. Lindström. “Dynamic allocation or diversification: A regime-based approach to multiple assets.” *Journal of Portfolio Management*, vol. 44, no. 2 (2017a), pp. 62–73.
- Nystrup, P., B. W. Hansen, H. Madsen, and E. Lindström. “Regime-based versus static asset allocation: Letting the data speak.” *Journal of Portfolio Management*, vol. 42, no. 1 (2015a), pp. 103–109.
- . “Detecting change points in VIX and S&P 500: A new approach to dynamic asset allocation.” *Journal of Asset Management*, vol. 17, no. 5 (2016), pp. 361–374.

- Nystrup, P., H. Madsen, and E. Lindström. “Stylised facts of financial time series and hidden Markov models in continuous time.” *Quantitative Finance*, vol. 15, no. 9 (2015b), pp. 1531–1541.
- . “Long memory of financial time series and hidden Markov models with time-varying parameters.” *Journal of Forecasting*, vol. 36, no. 8 (2017b), pp. 989–1002.
- . “Dynamic portfolio optimization across hidden market regimes.” *Quantitative Finance*, vol. 18, no. 1 (2018), pp. 83–95.
- Oh, K. J. and I. Han. “Using change-point detection to support artificial neural networks for interest rates forecasting.” *Expert Systems with Applications*, vol. 19, no. 2 (2000), pp. 105–115.
- Petropoulos, A., S. P. Chatzis, and S. Xanthopoulos. “A novel corporate credit rating system based on student’s- t hidden Markov models.” *Expert Systems with Applications*, vol. 53 (2016), pp. 87–105.
- . “A hidden Markov model with dependence jumps for predictive modeling of multidimensional time-series.” *Information Sciences*, vol. 412-413 (2017), pp. 50–66.
- Pinson, P., L. Christensen, H. Madsen, P. Sørensen, M. Donovan, and L. Jensen. “Regime-switching modelling of the fluctuations of offshore wind generation.” *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 96, no. 12 (2008), pp. 2327–2347.
- Qin, T. and K. Rohe. “Regularized spectral clustering under the degree-corrected stochastic block-model.” In *Advances in Neural Information Processing Systems* (2013), pp. 3120–3128.
- Robinson, W. N. and A. Aria. “Sequential fraud detection for prepaid cards using hidden Markov model divergence.” *Expert Systems with Applications*, vol. 91 (2018), pp. 235–251.
- Rousseeuw, K., Émilie Poisson Caillault, A. Lefebvre, and D. Hamad. “Hybrid hidden Markov model for marine environment monitoring.” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 8, no. 1 (2015), pp. 204–213.
- Rydén, T. “EM versus Markov chain Monte Carlo for estimation of hidden Markov models: A computational perspective.” *Bayesian Analysis*, vol. 3, no. 4 (2008), pp. 659–688.
- Rydén, T., T. Teräsvirta, and S. Åsbrink. “Stylized facts of daily return series and the hidden Markov model.” *Journal of Applied Econometrics*, vol. 13, no. 3 (1998), pp. 217–244.
- Samé, A., F. Chamroukhi, G. Govaert, and P. Aknin. “Model-based clustering and segmentation of time series with changes in regime.” *Advances in Data Analysis and Classification*, vol. 5, no. 4 (2011), pp. 301–321.
- Sedoc, J., J. Rodu, D. Foster, and L. Ungar. “Multiscale hidden Markov models for covariance prediction.” (2018).
- Shamshad, A., M. A. Bawadi, W. M. A. W. Hussin, T. A. Majid, and S. A. M. Sanusi. “First and second order Markov chain models for synthetic generation of wind speed time series.” *Energy*, vol. 30, no. 5 (2005), pp. 693–708.
- Shi, J. and J. Malik. “Normalized cuts and image segmentation.” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 8 (2000), pp. 888–905.

- Song, L., B. Boots, S. M. Siddiqi, G. Gordon, and A. Smola. “Hilbert space embeddings of hidden Markov models.” In *Proceedings of the 27th International Conference on Machine Learning* (2010), pp. 991–998.
- Verbeek, J. J., N. Vlassis, and B. Kröse. “Efficient greedy learning of Gaussian mixture models.” *Neural Computation*, vol. 15, no. 2 (2003), pp. 469–485.
- Viterbi, A. J. “Error bounds for convolutional codes and an asymptotically optimum decoding algorithm.” *IEEE Transactions on Information Theory*, vol. 13, no. 2 (1967), pp. 260–269.
- von Luxburg, U. “A tutorial on spectral clustering.” *Statistics and Computing*, vol. 17, no. 4 (2007), pp. 395–416.
- Weiss, Y. “Segmentation using eigenvectors: a unifying view.” In *Proceedings of the Seventh IEEE International Conference on Computer Vision*, vol. 2. IEEE (1999), pp. 975–982.
- Zelnik-Manor, L. and P. Perona. “Self-tuning spectral clustering.” In *Advances in Neural Information Processing Systems* (2005), pp. 1601–1608.
- Zhang, X., J. Li, and H. Yu. “Local density adaptive similarity measurement for spectral clustering.” *Pattern Recognition Letters*, vol. 32, no. 2 (2011), pp. 352–358.
- Zheng, K., Y. Li, and W. Xu. “Regime switching model estimation: spectral clustering hidden Markov model.” *Annals of Operations Research* (2019). URL <http://dx.doi.org/10.1007/s10479-019-03140-2>.