

Bayesian context trees: Modelling and exact inference for discrete time series

Ioannis Kontoyiannis¹  | Lambros Mertzanis² |
Athina Panotopoulou³ | Ioannis Papageorgiou⁴ | Maria Skoularidou⁵

¹Statistical Laboratory, Centre for Mathematical Sciences, University of Cambridge, Cambridge, UK

²Department of Electrical and Computer Engineering, University of Maryland, College Park, Maryland, USA

³Department of Computer Science, Dartmouth College, Hanover, New Hampshire, USA

⁴Department of Engineering, University of Cambridge, Cambridge, UK

⁵MRC-BSU, University of Cambridge, Cambridge, UK

Correspondence

Ioannis Kontoyiannis, Statistical Laboratory, Centre for Mathematical Sciences, University of Cambridge, Wilberforce Road, Cambridge CB3 0WB, UK.

Email: yiannis@maths.cam.ac.uk

Abstract

We develop a new Bayesian modelling framework for the class of higher-order, variable-memory Markov chains, and introduce an associated collection of methodological tools for exact inference with discrete time series. We show that a version of the context tree weighting algorithm can compute the prior predictive likelihood exactly (averaged over both models and parameters), and two related algorithms are introduced, which identify the a posteriori most likely models and compute their exact posterior probabilities. All three algorithms are deterministic and have linear-time complexity. A family of variable-dimension Markov chain Monte Carlo samplers is also provided, facilitating further exploration of the posterior. The performance of the proposed methods in model selection, Markov order estimation and prediction is illustrated through simulation experiments and real-world applications with data from finance, genetics, neuroscience and animal communication. The associated algorithms are implemented in the R package BCT.

KEYWORDS

Bayes factors, Bayesian context tree, context tree weighting, discrete time series, exact Bayesian inference, Markov chain Monte Carlo, Markov order estimation, model selection, prediction

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2022 The Authors. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* published by John Wiley & Sons Ltd on behalf of Royal Statistical Society.

1 | INTRODUCTION

Higher-order Markov chains are frequently the first and most natural modelling choice for discrete time series with significant—apparent or suspected—temporal structure, especially when no specific underlying data generating mechanism can be assumed. But the description of a full Markov chain of order d with values in a set of size m , say, requires the specification of $m^d(m-1)$ parameters, which makes the use of full Markov chains problematic in practice: As has been often noted (Bühlmann & Wyner, 1999; Raftery, 1985; Sarkar & Dunson, 2016), the dimension of the parameter space grows exponentially with the memory length, and the resulting model class lacks modelling wealth and flexibility. This lack of flexibility severely hinders, among other things, the important goal of balancing the bias-variance trade-off between more complex models that fit the data closely, and simpler models that generalise well.

To address these issues and to offer better solutions to a wealth of related scientific and engineering problems that arise in connection with discrete time series, numerous approaches have been developed since the mid-1980s.

The mixture transition distribution (MTD) models introduced by Raftery (1985) and later generalised in Raftery and Tavaré (1994) and Berchtold and Raftery (2002), allow for more parsimonious parametrisations of the transition distribution of some d th-order Markov chains, as mixtures of first-order transition matrices corresponding to different lags. MTD models make it possible to consider longer memory lengths d and to quantify the relative importance of different lags, but the resulting model class is still structurally poor.

A much more flexible class of Markov chain models with *variable memory* are the *tree sources* introduced in the celebrated work of Rissanen (1983a, 1983b, 1986a). The gist of this approach is that the length of the memory that determines the transition probability of the chain can depend on the exact pattern of the most recently occurring symbols. Initially tree sources received a lot of attention in the information-theoretic literature in connection with data compression (Weinberger et al., 1994; Willems et al., 1995). One of their first applications outside information theory was by Ron et al. (1996), who introduced the notion of a probabilistic suffix tree (PST) as an effective structure for representing variable-memory chains. The PST point of view, along with the associated model selection technique *Learn-PSA*, have been used for bioinformatics problems and other machine learning tasks (Bejerano & Yona, 2001; Gabadinho & Ritschard, 2016).

In the statistics literature, tree-structured models were examined by Bühlmann and Wyner (1999) and Bühlmann (2000). Their variable-length Markov chains (VLMCs) and the associated model selection tools are based, in part, on Rissanen's tree sources and his *CONTEXT* algorithm. The VLMC approach has been successful in applications (Busch et al., 2009; Mächler & Bühlmann, 2004) that include DNA modelling (Ben-Gal et al., 2005; Browning, 2006) and linguistics (Abakuks, 2012; Galves et al., 2012).

Tree sources and VLMC models group together certain patterns of past symbols that lead to the same conditional distribution for the chain, under the constraint that each such group consists of all patterns of length d that share a common suffix. A more general class of parsimonious Markov models, known as sparse Markov chains (SMC), arises when this constraint is removed. Originally introduced as 'minimal Markov models' by García and González-López (2011), they were later examined in more detail in Jääskinen et al. (2014), Xiong et al. (2016) and García and González-López (2017). But the lack of structure of this vast model class makes it difficult to identify appropriate models in practice.

A different interesting class of higher-order Markov models was more recently introduced by Sarkar and Dunson (2016), who used conditional tensor factorisation (CTF) to give parsimonious

representations of the full transition probability distribution (viewed as a high-dimensional tensor) of a Markov chain. These representations are based on extensions of earlier ideas on tensor factorisation for categorical regression (Yang & Dunson, 2016). CTF effectively shrinks the high-dimensional transition probability tensor to a lower-dimensional structure that can still capture high-order dependence. Unlike MTD, CTF accommodates complex interactions between the lags, and is accompanied by computational tools that allow for rich Bayesian inference.

In this work, we revisit the class of variable-memory Markov models. We introduce a new Bayesian framework for a version of these models, and we develop algorithmic tools that lead to very effective and efficient *exact* inference. Although our methods open the door to a wide range of statistical and machine learning applications—including anomaly detection, change-point estimation and pattern analysis—here we focus primarily on the more fundamental tasks of model selection, estimation and sequential prediction.

1.1 | Outline of contributions

In Section 2, we define a class of models for variable-memory Markov chains that admit natural representations as context trees. Given a finite set A and a maximal memory length $D \geq 0$, the class $\mathcal{T}(D)$ contains all variable-memory models of Markov chains with values in A and memory no longer than D . A new family of discrete prior distributions $\pi_D(T; \beta)$ on models $T \in \mathcal{T}(D)$ is introduced, indexed by a hyperparameter $\beta \in (0, 1)$. Roughly speaking, π_D penalises larger and more complex models by an exponential amount. Given a model T , we place independent Dirichlet priors $\pi(\theta|T)$ on the associated parameters θ . We refer to the models in $\mathcal{T}(D)$ equipped with this prior structure as *Bayesian context trees* (BCT).

Sections 3.1–3.3 contain our core methodological results, in the form of three exact inference algorithms for BCTs. First, we show that a version of the context tree weighting (CTW) algorithm (Tjalkens et al., 1994; Willems et al., 1995) can be used to not only evaluate the marginal likelihoods $P(x|T) = \int P(x|\theta, T)\pi(\theta|T)d\theta$ of observations x with respect to models T , which are easy to obtain, but also the *prior predictive likelihood* $P_D^*(x)$, averaged over all models,

$$P_D^*(x) := \sum_{T \in \mathcal{T}(D)} \pi_D(T; \beta) P(x|T) = \sum_{T \in \mathcal{T}(D)} \pi_D(T; \beta) \int P(x|\theta, T)\pi(\theta|T)d\theta. \quad (1)$$

The CTW algorithm computes $P_D^*(x)$ exactly, and its complexity is only linear in the length of the observed time series x . Since one of the basic obstacles to performing effective Bayesian inference is the difficulty to either sample from or obtain expectations with respect to the posterior distribution, in part stemming from the impossibility of computing its normalising factor $P_D^*(x)$, it is clear that the exact nature of the results produced by the CTW algorithm should facilitate the development of efficient methods for numerous core statistical tasks and related applications.

In Section 3.2, we describe the BCT algorithm and prove that it identifies the maximum a posteriori probability (MAP) model. This is a generalisation of the ‘context tree maximizing’ (CTM) algorithm of Willems and Volf (1994). And in Section 3.3, we show that a new algorithm, the k -BCT algorithm, can be used to identify the k a posteriori most likely tree models, for any $k \geq 1$. Despite the fact that the class $\mathcal{T}(D)$ is vast, consisting of doubly exponentially many models in the memory length D , the complexity of both the BCT and k -BCT algorithms is only linear in D and in the length of the observations x . But as a function of k , the complexity of k -BCT grows faster than linearly in k ; in fact, in its naive implementation, it grows like k^m , where m is the number

of possible values of the time series x . So its practical applicability is limited to relatively small values of k .

In order to enable broader exploration of the posterior distributions $\pi(T|x)$ and $\pi(\theta, T|x)$, in Section 3.5, we develop a new family of variable-dimension Markov chain Monte Carlo (MCMC) algorithms that obtain samples from $\pi(\theta, T|x)$. Their performance is illustrated in Section 5 on model selection, parameter estimation and Markov order estimation problems, on simulated and real data examples.

In Section 4, we present extensive model selection results, comparing the performance of the BCT framework with that of the corresponding VLMC and MTD methods, on both real and simulated data. We find that the BCT algorithm consistently performs at least as well as VLMC and MTD and usually gives a better fit on simulated data. Moreover, the k -BCT algorithm in combination with the MCMC samplers of Section 3.5 identify a number of candidate models for the observed data, also providing a quantitative measure of uncertainty for the selected models in the form of posterior probabilities. Using the CTW algorithm, these posterior probabilities can be computed exactly, as can the relevant Bayes factors and posterior odds for a variety of hypotheses of interest. In terms of complexity, the BCT algorithm is found to be computationally much more effective than MTD and VLMC. In fact, the linear complexity of CTW and BCT facilitates their use in big-data applications, as illustrated, for example in the analysis of a neural spike train data set of $\approx 4 \times 10^6$ samples, with memory lengths up to $D = 1500$.

In Section 6, we compare the natural predictor induced by the BCT framework with the predictors provided by the MTD, VLMC, SMC and CTF methodologies. The BCT predictor is seen to have two significant advantages, which lead to superior performance. The first is that the *posterior predictive distribution* can be computed exactly, as $P_D^*(x_{n+1}|x_1, \dots, x_n) = P_D^*(x_1, \dots, x_{n+1})/P_D^*(x_1, \dots, x_n)$, via the CTW algorithm. This way, the induced predictor is obtained by implicitly averaging over all models with respect to their exact posterior probabilities, thus avoiding the need to perform approximate model averaging via simulation or other numerical integration methods. The second advantage is that, because the CTW algorithm can be updated sequentially, so can the BCT predictor, so that it continues to ‘learn’ from the data even past the training phase. Results on both simulated and real time series illustrate the performance of the BCT predictor, confirming these observations.

1.2 | Further connections and comments

Variable-memory models, like the Bayesian context trees considered in this work, describe a flexible and rich class of higher-order Markov chains that admit parsimonious parametrisations and allow for natural graphical representations of important structural dependencies. The shape of the context tree can be easily interpreted and provides useful information about the regularities present in the data (Bejerano & Yona, 2001; Mächler & Bühlmann, 2004). Because BCTs are a vast model class, global model selection techniques based, for example on criteria like AIC and BIC, cannot be applied directly via, say, exhaustive search. But efficient tools like the BCT algorithm presented here make it possible to describe complex sequential data in a way that offers an effective balance for the simplicity-expressivity trade-off (Garivier & Leonardi, 2011).

Another point of view which naturally relates to the present development is Rissanen’s celebrated minimum description length (MDL) principle (Grünwald, 2007; Rissanen, 1987, 1989). The MDL principle provides a broad operational foundation for statistical inference, as well as constructive tools and appealing metaphors for selecting prior distributions (Chipman et al.,

2001). In particular, MDL considerations underpin much of the original work on the CTW algorithm (Tjalkens et al., 1994; Willems et al., 2002) and our own choice of priors in Section 2.

A method commonly used for model comparison is the Bayesian information criterion (BIC), derived by Schwarz (1978) as an asymptotic approximation to twice the logarithm of the Bayes factor $P(x|T)/P(x|T')$ between two models T, T' (Kass & Raftery, 1995). The form of the BIC and its familiar $(1/2) \log n$ -per-degree-of-freedom log-likelihood penalty also shares deep connections with the MDL principle, see for example the discussions by Barron et al. (1998) and Csiszár and Shields (2000). Further comments on this are given in Section 6 in connection with Theorem 5.

Finally, we note that there exist a number of alternative approaches to modelling discrete time series. An important collection of tools is provided by hidden Markov models. HMMs are a general and very broadly used model class, with a wide range of applications and a variety of associated methodological procedures for learning and inference (Bishop, 2006; Cappé et al., 2006). A more classical approach to discrete time series modelling is via discrete analogues of linear models, often using multinomial logit or probit regression (Yee, 2010). Such a linear-predictor approach is described in Zeger and Liang (1986), and a different parsimonious class of models with an emphasis on binary time series is given in Fahrmeir and Kaufmann (1987). A treatment of partial likelihood inference on generalised discrete linear models is presented in Fokianos and Kedem (2003), and an extension of the traditional ARMA methodology to integer autoregressive models for count time series is developed in Fokianos (2012).

2 | BAYESIAN CONTEXT TREES

The distribution of a full d th-order Markov chain $\{X_n\}$ with values in the finite state space, or *alphabet*, A , is identified by its conditional distributions,

$$\theta_s(a) := \Pr(X_n = a | X_{n-1} = x_{n-1}, X_{n-2} = x_{n-2}, \dots, X_{n-d} = x_{n-d}), \quad a \in A,$$

for every *context* $s = (x_{n-1}, x_{n-2}, \dots, x_{n-d}) \in A^d$ of length d . If the alphabet A has size m , the description of these conditional distributions requires the specification of $(m-1)m^d$ parameters. But suppose, for example, that $A = \{0, 1, 2\}$, and that when the most recent symbol x_{n-1} is a '1' the distribution of the next symbol is independent of the remaining values x_{n-2}, \dots, x_{n-d} , that is, the conditional probabilities,

$$\Pr(X_n = a | X_{n-1} = 1, X_{n-2} = x_{n-2}, \dots, X_{n-d} = x_{n-d}),$$

only depend on a . An example of how the distribution of such a fifth order, *variable-memory* Markov chain may be represented by a labelled tree is shown in Figure 1.

Each leaf of the tree corresponds to a string s containing between one and five symbols from A , determined by the labels along the path from the root λ to that leaf. Whenever the string s consists of the same sequence of symbols as the most recent x_i 's, the probability that the next value of the chain will be ' a ' is $\theta_s(a)$, given by the distribution θ_s marked on that leaf, for example

$$\Pr(X_n = 2 | X_{n-1} = 0, X_{n-2} = 1, X_{n-3} = 1, \dots) = \Pr(X_n = 2 | X_{n-1} = 0, X_{n-2} = 1) = \theta_{01}(2).$$

Note that, since the tree model for this chain has 13 leaves, instead of the full $2 \times 3^5 = 486$ parameters, it suffices to only specify $2 \times 13 = 26$.

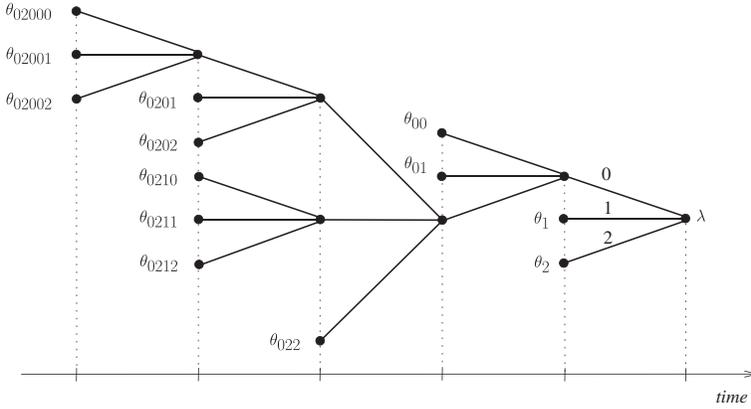


FIGURE 1 The tree model and parameters of a fifth-order variable-memory chain

2.1 | Variable-memory Markov chains

Let $\{X_n\}$ be a d th-order Markov chain, for some $d \geq 0$, with values in an alphabet A of size $m \geq 2$. Without loss of generality, we take $A = \{0, 1, \dots, m-1\}$ throughout. The *model* describing $\{X_n\}$ as a variable-memory chain will always be represented by a tree as in the example above.

Let T be an m -ary tree of depth no greater than d , which is *proper*, in that, if a node in T is not a leaf, then it has exactly m children. This means that every length- d context $s := (x_{-1}, x_{-2}, \dots, x_{-d}) \in A^d$ has a unique suffix (x_{-1}, \dots, x_{-j}) which is a leaf of T , for some $j \leq d$.

Let C be the function that maps each context $s \in A^d$ to a leaf of T . Viewing T as the collection of its leaves, it is the range of C ,

$$C : A^d \rightarrow T \subset \bigcup_{i=0}^d A^i,$$

with the convention that $A^0 = \{\lambda\}$ contains only the *empty string* λ .

For indices $i \leq j$, we write X_i^j for the vector of random variables $(X_i, X_{i+1}, \dots, X_j)$ and similarly $x_i^j \in A^{j-i+1}$ for a string $(x_i, x_{i+1}, \dots, x_j)$ representing a particular realisation of these random variables. Then the Markov property for a variable-memory chain $\{X_n\}$ with model T takes the form:

$$P(x_1^n | x_{-d+1}^0) := \Pr(X_1^n = x_1^n | X_{-d+1}^0 = x_{-d+1}^0) = \prod_{i=1}^n P(x_i | x_{i-d}^{i-1}) = \prod_{i=1}^n P(x_i | C(x_{i-d}^{i-1})). \quad (2)$$

The complete description of the distribution of $\{X_n\}$, in addition to the model T , requires the specification of a set of *parameters* $\theta = \{\theta_s; s \in T\}$: To every $s \in T$, we associate a probability vector,

$$\theta_s := (\theta_s(0), \theta_s(1), \dots, \theta_s(m-1)),$$

where the $\theta_s(j)$ are non-negative and sum to one. Then (2) can be written,

$$P(x_1^n | x_{-d+1}^0) = \prod_{i=1}^n \theta_{C(x_{i-d}^{i-1})}(x_i),$$

or, alternatively,

$$P(x_1^n | x_{-d+1}^0) = \prod_{s \in T} \prod_{j \in A} \theta_s(j)^{a_s(j)}, \quad (3)$$

where each element $a_s(j)$ of the *count vector* $a_s = (a_s(0), a_s(1), \dots, a_s(m-1))$ is,

$$a_s(j) := \# \text{ times symbol } j \in A \text{ follows context } s \text{ in } x_1^n. \quad (4)$$

2.2 | Bayesian context trees

Throughout the paper, we consider every d th-order, variable-memory chain $\{X_n\}$ as described by a (proper) tree model T with associated parameters $\theta = \{\theta_s; s \in T\}$. Here we define a family of prior distributions on (T, θ) .

Model prior. For a maximal depth $D \geq 0$ and an alphabet $A = \{0, 1, \dots, m-1\}$ of size $m \geq 2$, let $\mathcal{T}(D)$ denote the collection of all (proper) tree models T on A with depth no greater than D . Given an arbitrary $\beta \in (0, 1)$, we define the prior distribution,

$$\pi(T) := \pi_D(T; \beta) := \alpha^{|T|-1} \beta^{|T|-L_D(T)}, \quad T \in \mathcal{T}(D), \quad (5)$$

where $\alpha := (1 - \beta)^{1/(m-1)}$, $|T|$ denotes the number of leaves of T , and $L_D(T)$ denotes the number of leaves T has at depth D . The following lemma, proved in Section A of the supplementary material, states that (5) does indeed define a probability distribution.

Lemma 1 For any $D \geq 0$ and any $\beta \in (0, 1)$: $\sum_{T \in \mathcal{T}(D)} \pi_D(T; \beta) = 1$.

Prior on θ . Given a model $T \in \mathcal{T}(D)$, we place an independent Dirichlet prior with parameters $(1/2, 1/2, \dots, 1/2)$ on each θ_s so that, $\pi(\theta|T) = \prod_{s \in T} \pi(\theta_s)$, where,

$$\pi(\theta_s) = \pi(\theta_s(0), \theta_s(1), \dots, \theta_s(m-1)) = \frac{\Gamma(m/2)}{\pi^{m/2}} \prod_{j=0}^{m-1} \theta_s(j)^{-\frac{1}{2}} \alpha \prod_{j=0}^{m-1} \theta_s(j)^{-\frac{1}{2}}. \quad (6)$$

Although we take the parameter vector of the Dirichlet prior for all θ_s to be $(1/2, \dots, 1/2)$, corresponding to the Jeffreys prior, the extension of all our results to arbitrary Dirichlet parameters is straightforward as outlined in Section D of the supplementary material.

Finally, given the model T and associated parameters $\theta = \{\theta_s; s \in T\}$, from Equation (3) we have,

$$P(x_1^n | x_{-d+1}^0, \theta, T) = \prod_{s \in T} \prod_{j=0}^{m-1} \theta_s(j)^{a_s(j)}, \quad (7)$$

where a_s are the count vectors defined in Equation (4). By convention, when we write $\sum_{s \in T}$ or $\prod_{s \in T}$, we take the corresponding sum or product over all the *leaves* s of the tree, not all its nodes. Also, in order to avoid cumbersome notation, in what follows we often write x for the entire time series x_1^n and suppress the dependence on its initial context x_{-d+1}^0 , so that, for example, we denote,

$$P(x, \theta|T) = P(x_1^n, \theta | x_{-d+1}^0, T) = P(x_1^n | x_{-d+1}^0, \theta, T) \pi(\theta|T).$$

Choice of β . The first factor, $\alpha^{|T|-1}$, in the definition of the model prior $\pi_D(T; \beta)$ is the more important and easier one to interpret, showing that larger models are penalised by an exponential amount, while the second factor, $\beta^{|T|-L_D(T)}$, adds a less intuitive penalty to the leaves at depth strictly smaller than D .

Consider two models $S, T \in \mathcal{T}(D)$ such that S is a subtree of T . If T is produced by adding a single branch of m nodes to one of the leaves, s , say, of S , and s is at depth $D - 2$ or smaller, then $|T| = |S| + m - 1$ and $L_D(T) = L_D(S)$, so that,

$$\frac{\pi_D(T; \beta)}{\pi_D(S; \beta)} = (1 - \beta)\beta^{m-1},$$

which, as desired, is always less than one. But if T is produced from S by adding a branch of m nodes to a node s at depth $D - 1$, then $L_D(T) = L_D(S) + m$, and,

$$\frac{\pi_D(T; \beta)}{\pi_D(S; \beta)} = \frac{1 - \beta}{\beta}.$$

This is strictly decreasing in β , and for $\beta < 1/2$ it is greater than one, while for $\beta > 1/2$ it is smaller than 1.

Therefore, $\pi_D(T; \beta)$ penalises larger trees by an exponential amount as long as $\beta \geq 1/2$, and larger values of β make the penalisation more severe. Also, for larger alphabet sizes, $\alpha = (1 - \beta)^{1/(m-1)}$ becomes very close to 1 and the second factor dominates, an effect which is unintuitive and less desirable. Therefore, in practice, we will always take $\beta \approx 1 - 2^{-m+1}$, so that $\alpha \approx 1/2$, unless there are specific reasons for a different choice.

2.3 | Marginal likelihood and prior predictive likelihood

A useful property of the BCT framework is that the parameters θ can easily be integrated out, so that the *marginal likelihoods* $P(x|T)$ can be expressed in closed form. This result, stated without proof in Lemma 2, is based on a standard computation.

Lemma 2 *The marginal likelihood $P(x|T)$ of the observations x given a model T is,*

$$P(x|T) = \int P(x, \theta|T) d\theta = \int P(x|\theta, T) \pi(\theta|T) d\theta = \prod_{s \in T} P_e(a_s),$$

where the count vectors $a_s = (a_s(0), a_s(1), \dots, a_s(m-1))$ are defined in Equation (4) and the estimated probabilities $P_e(a_s)$ are defined by,

$$P_e(a_s) := \frac{\prod_{j=0}^{m-1} [(1/2)(3/2) \cdots (a_s(j) - 1/2)]}{(m/2)(m/2 + 1) \cdots (m/2 + M_s - 1)}, \quad (8)$$

where $M_s := a_s(0) + a_s(1) + \cdots + a_s(m-1)$, with the convention that any empty product is taken to be equal to 1.

In terms of inference, one of the main objects of interest is the model posterior distribution,

$$\pi(T|x) = \frac{P(x|T)\pi(T)}{P(x)},$$

where the denominator is the *prior predictive likelihood* $P(x) = P_D^*(x)$ in Equation (1). The difficulty in computing $P_D^*(x)$ comes, in part, from the fact that the class $\mathcal{T}(D)$ of variable-memory models is enormously rich, even for moderate (or even small) alphabet sizes m and maximal depths D : The size of $\mathcal{T}(D)$ grows doubly exponentially in D ; a simple computation shows that,

$$|\mathcal{T}(D)| \geq \sum_{d=0}^{D-1} 2^{m^d} - D. \quad (9)$$

In Sections 3.1–3.3, we describe three exact inference algorithms for the class of Bayesian context trees. We show that, perhaps somewhat surprisingly, $P_D^*(x)$ can be computed *exactly* and in a very efficient manner, and that the mode of the posterior $\pi(T|x)$ as well as the next few most likely models can be explicitly identified.

3 | METHODOLOGY

This section contains our core methodological contributions, which form the basis for most of the statistical and learning tasks that can be performed within the BCT framework.

The CTW algorithm is described in Section 3.1, and it is proved that it indeed computes the prior predictive likelihood $P_D^*(x)$ of a time series x . Section 3.2 introduces the BCT algorithm, which identifies the maximum a posteriori probability (MAP) tree model T_1^* given x . A more general algorithm, k -BCT, is described in Section 3.3; the k -BCT algorithm identifies not just the mode T_1^* of the posterior $\pi(T|x)$ on model space, but the k a posteriori most likely tree models, $T_1^*, T_2^*, \dots, T_k^*$, for any $k \geq 1$. All three algorithms are implemented in the publicly available R package BCT (Papageorgiou et al., 2020).

In Section 3.4, we discuss how the posterior probability $\pi(T|x)$ of any model T can be easily computed, and we identify the full conditional density of the parameters, $\pi(\theta|x, T)$. In order to explore the posterior distribution further, a family of variable-dimension MCMC samplers are developed in Section 3.5. Finally, in Section 3.6, we describe how the CTW, BCT and BCT algorithms can all be updated sequentially, and show that their complexity is only linear in the length of the time series x . In particular, the fact that the CTW algorithm can be implemented in a sequential fashion means that it can be effectively used for online prediction; cf. Section 6.

3.1 | CTW: The context tree weighting algorithm

The CTW algorithm takes as input: The size of the alphabet $m \geq 2$, the maximum context depth $D \geq 0$, a time series of observations x_1^n together with the initial context x_{-D+1}^0 , all taking values in the alphabet $A = \{0, 1, \dots, m-1\}$ and the value of the prior parameter $\beta \in (0, 1)$. It executes the following steps:

- (i) Build an m -ary tree, T_{MAX} , whose leaves are all the contexts x_{i-D}^{i-1} , $i = 1, 2, \dots, n$, that appear in the observations x_{-D+1}^n . If some node s of T_{MAX} is at depth $d < D$ and some but not all of its children are in T_{MAX} , then add all its remaining children as well, so that T_{MAX} is a proper tree.

- (ii) Compute the count vector a_s as in Equation (4), at every node s of the tree T_{MAX} (not only at the leaves), so that a_s will be the all-zero vector for the additional leaves included in the last step of (i).
- (iii) Compute the *estimated probability* $P_{e,s} := P_e(a_s)$ given by Equation (8), at each node s of the tree T_{MAX} , with the convention that $P_e(a_s) = 1$ when a_s is the all-zero count vector.
- (iv) Write sj for the concatenation of context s and symbol j , corresponding to the j th child of node s . Starting at the leaves and proceeding recursively towards the root, at each node s of the tree T_{MAX} compute the *mixture (or weighted) probabilities*:

$$P_{w,s} := \begin{cases} P_{e,s}, & \text{if } s \text{ is a leaf,} \\ \beta P_{e,s} + (1 - \beta) \prod_{j=0}^{m-1} P_{w,sj}, & \text{otherwise.} \end{cases}$$

- (v) Output the mixture probability $P_{w,\lambda}$ at the root λ .

Theorem 1 is proved in Section A of the supplementary material.

Theorem 1 *The mixture probability $P_{w,\lambda}$ at the root λ computed by CTW is exactly the prior predictive likelihood of the observations,*

$$P_{w,\lambda} = P_D^*(x) = \sum_{T \in \mathcal{T}(D)} \pi_D(T; \beta) \int_{\theta} P(x_1^D | x_{-D+1}^0, \theta, T) \pi(\theta | T) d\theta,$$

where the sum is over all proper context tree models of depth no greater than D .

3.2 | BCT: The Bayesian context tree algorithm

The BCT algorithm takes the same input as CTW and executes the following steps:

- (i) Build the m -ary tree T_{MAX} and compute the count vectors a_s and the estimated probabilities $P_{e,s} = P_e(a_s)$ at all nodes s of T_{MAX} , as in steps (i)–(iii) of CTW.
- (ii) Starting at the leaves and proceeding recursively towards the root, at each node s of the tree T_{MAX} compute the *maximal probabilities*,

$$P_{m,s} := \begin{cases} P_{e,s}, & \text{if } s \text{ is a leaf at depth } D, \\ \beta, & \text{if } s \text{ is a leaf at depth } < D, \\ \max \left\{ \beta P_{e,s}, (1 - \beta) \prod_{j=0}^{m-1} P_{m,sj} \right\}, & \text{otherwise.} \end{cases} \quad (10)$$

- (iii) Starting at the root and proceeding recursively with its descendants, for each node s : If the maximum in Equation (10) can be achieved by the first term, then prune all its descendants from the tree T_{MAX} ; otherwise, repeat the same process at each of the m children of node s .
- (iv) After all nodes have been exhausted in (iii), output the resulting tree T_1^* and the maximal probability at the root $P_{m,\lambda}$.

The following theorem is proved in Section A of the supplementary material.

Theorem 2 For all $\beta \geq 1/2$, the tree T_1^* produced by the BCT algorithm is the MAP tree model (or one of the MAP tree models, in case the maximum below is not uniquely achieved),

$$\pi(T_1^*|x) = \max_{T \in \mathcal{T}(D)} \pi(T|x), \quad (11)$$

and the maximal probability at the root satisfies,

$$P_{m,\lambda} = P(x|T_1^*)\pi_D(T_1^*; \beta) = P(x, T_1^*). \quad (12)$$

3.3 | k -BCT: The top- k Bayesian context trees algorithm

The k -BCT is one of the main novel contributions of this work. Although it is conceptually a natural generalisation of BCT, the precise description of its most efficient implementation is quite lengthy. So, for the sake of both clarity and brevity, we first describe here a less practical, idealised version of k -BCT, which is conceptually identical with its the practical version and only differs from it in the initialisation step of the leaves at depth $d < D$. The actual practical algorithm is given in Section B of the supplementary material, along with a discussion of its implementation complexity.

The idealised k -BCT algorithm takes the same input as CTW together with the number $k \geq 1$ of the a posteriori most likely models to be determined, and it executes the following steps:

- (i) Let T_{MAX} be the *complete* m -ary tree at depth D (this is the ‘idealised’ part), and compute the count vectors a_s and the estimated probabilities $P_{e,s} = P_e(a_s)$ at all nodes s of T_{MAX} as in steps (ii) and (iii) of CTW.
- (ii) Starting at the leaves and proceeding towards the root, at each node s compute a list of k maximal probabilities $P_{m,s}^{(i)}$ and k position vectors $c_s^{(i)} = (c_s^{(i)}(0), c_s^{(i)}(1), \dots, c_s^{(i)}(m-1))$, for $i = 1, 2, \dots, k$, where each $c_s^{(i)}(j)$ is an integer between 0 and k , recursively, as follows.
 - (iia) At each leaf s , we let $P_{m,s}^{(1)} = P_{e,s}$ and $c_s^{(1)} = (0, 0, \dots, 0)$, where the all-zero vector $c_s^{(1)}$ indicates that $P_{m,s}^{(1)}$ corresponds to the value of $P_{e,s}$ and does not depend on the children of s (since there are none). For $i = 2, 3, \dots, k$, we leave $P^{(i)}$ and $c_s^{(i)}$ undefined.
 - (iib) At each node s having only m descendants (which are necessarily leaves), we compute the two probability-position vector pairs $\beta P_{e,s}, (0, 0, \dots, 0)$ and $(1 - \beta) \prod_{j=0}^{m-1} P_{m,sj}^{(1)}, (1, 1, \dots, 1)$ (where the all-1 vector indicates that the latter probability only depends on the first maximal probability of each of the children), and sort them as $P_{m,s}^{(1)}, c_s^{(1)}$ and $P_{m,s}^{(2)}, c_s^{(2)}$ in order of decreasing probability. For $i = 3, 4, \dots, k$, we leave $P^{(i)}$ and $c_s^{(i)}$ undefined.
 - (iic) A general internal node s has m children, where each child sj has a list of k_j (for some $1 \leq k_j \leq k$) probability-vector pairs $P_{m,sj}^{(i)}, c_{sj}^{(i)}, 1 \leq i \leq k_j$. We compute the probability $\beta P_{e,s}$ with associated position vector $(0, 0, \dots, 0)$, and all possible probability-position vector pairs,

$$(1 - \beta) \prod_{j=0}^{m-1} P_{m,sj}^{(i_j)}, \quad (i_0, i_1, \dots, i_{m-1}), \quad (13)$$

for all possible combinations of indices $1 \leq i_j \leq k_j$ for $0 \leq j \leq m - 1$. We then sort these $k' = 1 + k_0 \times k_1 \times \dots \times k_{m-1}$ probabilities in order of decreasing probability, and rename the top k of them as $P_{m,s}^{(i)}$, for $i = 1, 2, \dots, k$, together with their associated position vectors $c_s^{(i)}$. [Of course, if $k' < k$, after sorting we leave the remaining $k - k'$ probability-position vector pairs undefined.]

- (iii) Having determined all maximal probabilities $P_{m,s}^{(i)}$ for all nodes s and $1 \leq i \leq k$, we now determine the ‘top k ’ trees $T_1^*, T_2^*, \dots, T_k^*$ from the corresponding position vectors $c_s^{(i)}$. For each i , we repeat the following process, starting at the root and proceeding until all available nodes of the tree T_{MAX} have been exhausted.
- (iiia) *Depth $d = 0$.* At the root node λ , we examine $c_\lambda^{(i)}$. If it is the all-zero vector, then T_i^* is the tree consisting of the root node only. Otherwise, we add to T_i^* the branch of m children starting at the root, and proceed to examine each of the nodes corresponding to the m children recursively.
 - (iiib) *Depth $d = 1$.* Reaching node $s = j$ corresponding to the j th child of the root, means that $t = c_\lambda^{(i)}(j)$ is nonzero. We examine $c_s^{(i)}$: If it is the all-zero vector, then we prune from T_i^* all the descendants of s and move to the next unexamined node; otherwise, we add to T_i^* the branch of m children starting at s , and proceed to examine each of the nodes corresponding to the m children recursively.
 - (iiic) *General depth $1 \leq d \leq D - 1$.* Reaching a node sj at depth d from its parent node s means that we decided to visit sj because $t = c_s^{(i)}(j)$ is nonzero for the appropriate index u (corresponding to the position vector $c_s^{(u)}$ that was examined at node s). We examine $c_{sj}^{(i)}$: If it is the all-zero vector, then we prune from T_i^* all the descendants of sj and move to the next unexamined node; otherwise, we add to T_i^* the branch of m children starting at sj , and proceed to examine each of the nodes corresponding to the m children recursively.
 - (iiid) *Depth $d = D$.* Reaching a node s at depth D means we have reached a leaf of T_{MAX} , so we simply add s to T_i^* and proceed to the next unexamined node.
- (iv) Output the resulting trees T_i^* and maximal probabilities at the root $P_{m,\lambda}^{(i)}$, $1 \leq i \leq k$.

Theorem 3 is proved in Section C of the supplementary material.

Theorem 3 *For any $\beta \geq 1/2$, the trees $T_1^*, T_2^*, \dots, T_k^*$ produced by the k -BCT algorithm are the k a posteriori most likely tree models: For each $i = 1, 2, \dots, k$,*

$$\pi(T_i^* | x) = \max_{T \in \mathcal{T}(D)}^{(i)} \pi(T | x), \quad (14)$$

where $\max_{t \in \mathcal{T}}^{(i)} f(t)$ of a function f defined on a discrete set of arguments $t \in \mathcal{T}$ denotes the i th largest value of f . Moreover, the i th maximal probability at the root satisfies,

$$P_{m,\lambda}^{(i)} = P(x | T_i^*) \pi_D(T_1^*, \beta) = P(x, T_i^*), \quad i = 1, 2, \dots, k.$$

Once again we note that, if some of the maxima in Equation (11) are not uniquely achieved, then T_1^*, \dots, T_k^* are one of the equivalent collections of k a posteriori most likely models.

3.4 | Computation of posterior probabilities

In addition to the prior predictive likelihood $P_D^*(x) = P_{w,\lambda}$ computed by CTW, and the a posteriori most likely models identified by the BCT and k -BCT algorithms, there are several other useful quantities that can easily be obtained through the results of these algorithms.

For a fixed maximal depth $D \geq 0$ and a fixed $\beta \in [1/2, 1)$, let x denote a time series x_1^n with initial context x_{-D+1}^0 .

Model posterior probabilities. For any model $T \in \mathcal{T}(D)$,

$$\pi(T|x) = \frac{P(x|T)\pi(T)}{P_D^*(x)} = \frac{\pi(T) \prod_{s \in T} P_e(a_s)}{P_{w,\lambda}}, \quad (15)$$

where $P_D^*(x) = P_{w,\lambda}$ is the prior predictive likelihood computed by CTW, and the numerator is given as in Lemma 2. If a model T has a leaf s that is not included in the tree T_{MAX} generated by CTW, so that no corresponding count-vector a_s is available, we follow the convention of setting $a_s = (0, 0, \dots, 0)$ and $P_e(a_s) = 1$.

Full conditional density of θ . Conditional on a model T and observations x , the density of the parameters $\theta = \{\theta_s; s \in T\}$ is,

$$\pi(\theta|x, T) \propto P(x|\theta, T)\pi(\theta|T) \propto \left(\prod_{s \in T} \prod_{j=0}^{m-1} \theta_s(j)^{a_s(j)} \right) \left(\prod_{s \in T} \prod_{j=0}^{m-1} \theta_s(j)^{-\frac{1}{2}} \right),$$

where we have used the definition of the prior on θ in Equation (6) and the expression for the likelihood (7). Therefore, the full conditional density $\pi(\theta|x, T)$ is the product of Dirichlet densities,

$$\pi(\theta|x, T) \sim \prod_{s \in T} \text{Dir}(a_s(0) + 1/2, a_s(1) + 1/2, \dots, a_s(m-1) + 1/2). \quad (16)$$

3.5 | MCMC samplers

The k -BCT algorithm can identify the k a posteriori most likely models, including in cases where the posterior $\pi(T|x)$ is multimodal, but, as discussed in Section 3.6, its complexity grows faster than linearly in k , which makes it impractical for large values of k . Next, we describe a family of effective variable-dimension MCMC samplers that make it possible to explore $\pi(T|x)$ further, and to sample from the posterior $\pi(\theta, T|x)$ jointly on models and parameters.

The random walk (RW) sampler for $\pi(T|x)$ is a Metropolis–Hastings algorithm with a proposal distribution that, at each step, either adds or removes an m -tuple of leaves from the current tree.

RW sampler. It takes as input the same parameters as BCT, and also: An initial model $T^{(0)} \in \mathcal{T}(D)$, the required number $N \geq 1$ of MCMC iterations, and the tree T_{MAX} together with the estimated probabilities $P_{e,s}$ at all nodes of T_{MAX} , as computed by CTW. It executes the following steps at each iteration $t = 1, 2, \dots, N - 1$:

- (i) Given $T^{(t)}$, propose a new tree T' as follows:
 - (a) If $T^{(t)}$ is the empty tree $\Lambda := \{\lambda\}$ consisting only of the root node, let T' be the complete tree at depth 1, $T_c(1)$.

- (b) If $T^{(t)}$ is the complete m -ary tree at depth D , $T_c(D)$, choose uniformly at random one of the internal nodes s at depth $D - 1$ and let T' be the same as $T^{(t)}$ but with the m children of s removed.
 - (c) Otherwise, with probability $1/2$ decide to propose a larger tree T' , formed by choosing uniformly at random one of the leaves of $T^{(t)}$ at depths $\leq D-1$ and adding its m children to form T' ;
 - (d) Or, with probability $1/2$ decide to propose a smaller tree T' , formed by choosing uniformly at random one of the internal nodes s of $T^{(t)}$ that only have m descendants and removing the m leaves that stem from s , to form T' .
- (ii) Either accept T' and set $T^{(t+1)} = T'$, or reject it and set $T^{(t+1)} = T^{(t)}$, with corresponding probabilities $\alpha(T^{(t)}, T') = \min\{1, r(T^{(t)}, T')\}$ and $1 - \alpha(T^{(t)}, T')$, respectively; explicit expressions for the ratios $r(T, T')$ are given in Section E of the supplementary material.

The jump sampler for $\pi(T|x)$ is a modification of the RW sampler, which, in addition to nearest neighbour moves, also allows for jumps to any one of the k most likely models. This way we overcome the common difficulty of RW samplers to move between separated modes of multimodal posterior distributions.

Jump sampler. It takes as input the same parameters as the RW sampler, and also the value of a jump parameter $p \in (0, 1)$ and the collection $\mathcal{T}^* = \{T_1^*, \dots, T_k^*\}$ of the top k trees computed by k -BCT. It executes the following steps at each iteration $t = 1, 2, \dots, N - 1$:

- (i) Given $T^{(t)}$, propose a new tree T' as follows:
 - (a) With probability $(1 - p)$, propose a new tree T' as in steps (ia)–(id) of the RW sampler;
 - (b) Or, with probability p , propose a jump move: Let T' be one of the top k trees T_i^* , uniformly chosen from \mathcal{T}^* .
- (ii) Either accept T' and set $T^{(t+1)} = T'$, or reject it and set $T^{(t+1)} = T^{(t)}$, with corresponding probabilities $\alpha(T^{(t)}, T') = \min\{1, r(T^{(t)}, T')\}$ and $1 - \alpha(T^{(t)}, T')$, where the ratios $r(T, T')$ are given explicitly in Section E of the supplementary material.

Note that a jump move to one of the top k models T_i^* may be proposed from any state $T^{(t)}$ of the sampler, but it only has a nonzero probability of being accepted if $T^{(t)}$ itself is either one of the T_i^* or a neighbour of one of them. This suggests that the jump parameter should be chosen so that $(1 - p)$ is not too small; in all of our experiments below we take $p = 1/2$.

MCMC convergence. The target distribution of both the RW and jump samplers is the posterior $\pi(T|x)$, $T \in \mathcal{T}(D)$, on model space. Unlike with most MCMC samplers used in Bayesian inference, here we can in fact compute the value of the target distribution $\pi(T|x)$ exactly, for any specific $T \in \mathcal{T}(D)$, as noted in Section 3.4. Nevertheless, because of the enormity of the space $\mathcal{T}(D)$ and the fact that it does not possess any easily exploitable structure, it is still practically impossible to sample from $\pi(T|x)$ directly, hence we resort to MCMC. On the other hand, knowing the posterior probabilities $\pi(T|x)$ precisely (not only up to a multiplicative constant) means that it is easy to obtain a good first indication of whether the sampler has converged, or at least whether it has spent the ‘right’ amount of time in the important areas of the support of $\pi(T|x)$ near its mode(s): Simply compute the frequency of each of the top k models T_i^* in the MCMC sample, and compare it with its actual posterior probability $\pi(T_i^*|x)$.

Although we have not observed convergence issues in our experiments on simulated or real data, we note that more sophisticated MCMC methods can also be used, for example tempering the likelihood (Robert & Casella, 2004) or using a Wang–Landau-style algorithm to force the

sampler to spend a specified proportion of time at models of each depth (Atchade & Liu, 2004).

Joint sampler. Being able to obtain MCMC samples $\{T^{(t)}\}$ for $\pi(T|x)$, and knowing the full conditional density $\pi(\theta|x, T)$ of the parameters explicitly as in Equation (16), it is simple to obtain a corresponding sequence of samples $\{(\theta^{(t)}, T^{(t)})\}$ for the posterior $\pi(\theta, T|x)$ jointly on models and parameters. This can be done by drawing a conditionally independent sample $\theta^{(t)} \sim \pi(\theta|x, T^{(t)})$ at each MCMC iteration.

3.6 | Sequential updates, prediction and complexity

As more observations become available, the results of all three exact inference algorithms in Sections 3.1–3.3 can be updated sequentially. This facilitates their online use in applications where it is important that data be processed sequentially rather than in large blocks.

For CTW, having computed $P_D^*(x_1^n | x_{-D+1}^0)$ and given an additional sample x_{n+1} , the new prior predictive likelihood $P_D^*(x_1^{n+1} | x_{-D+1}^0)$ can be obtained in $O(D)$ operations by updating steps (ii)–(iv) of CTW as follows. Let s_D, s_{D-1}, \dots, s_0 be the contexts of lengths $D, D-1, \dots, 0$, respectively, immediately preceding x_{n+1} , so that in particular $s_D = x_{n-D+1}^n$ and $s_0 = \lambda$. Suppose $x_{n+1} = j$; at each of the nodes s_D, s_{D-1}, \dots, s_0 in the tree T_{MAX} already constructed (in that order, and *only* there):

- (ii') Update $a_s(j)$ and M_s by increasing each of their values by 1.
- (iii') Update $P_{e,s} = P_e(a_s)$ by multiplying its earlier value by $(a_s(j) - 1/2)/(m/2 + M_s - 1)$, for the updated values of $a_s(j)$ and M_s .
- (iv') Re-compute the probability $P_{w,s}$.

The required result $P_D^*(x_1^{n+1} | x_{-D+1}^0)$ is the (updated) mixture probability $P_{w,\lambda}$ at the root.

The corresponding update rules for BCT and k -BCT are analogous.

The ability to compute the prior predictive likelihood sequentially makes it easy to perform online prediction. The canonical Bayesian rule for predicting the next observation x_{n+1} given the past x_1^n , is given by the *posterior predictive distribution*,

$$P_D^*(x_{n+1} | x_1^n) = \sum_T \int_{\theta} P(x_{n+1} | x_1^n, \theta, T) \pi(\theta, T | x_1^n) d\theta, \quad (17)$$

where we suppress the dependence on the initial context x_{-D+1}^0 for brevity. Although it is common that $P_D^*(x_{n+1} | x_1^n)$ can only be estimated (e.g. using MCMC sampling or approximate model averaging), here its value can be computed exactly and sequentially by CTW, via,

$$P_D^*(x_{n+1} | x_1^n) = \frac{P_D^*(x_1^{n+1})}{P_D^*(x_1^n)}. \quad (18)$$

Applications of this methodology on both simulated and real data are given in Section 6.

Next we briefly discuss the implementation complexity of the three algorithms in Sections 3.1–3.3, as a function of the parameters n, m, D , and k . The complexity of CTW is linear in each of n, m and D , and in fact it is of order $O(nmD)$. To see this, observe that in step (i), for each $x_i, 1 \leq i \leq n$, a new node is created for each of the $(D+1)$ contexts of x_i , of lengths $0, 1, \dots, D$. This requires $O(nD)$ operations and produces the tree T_{MAX} which has no more than $nD + 1$ nodes.

The second and third steps, where the count-vectors a_s and the probabilities $P_e(a_s)$ are computed, can be integrated into the first one. For each i , when we visit each of the $(D + 1)$ contexts s of x_i , we increase the corresponding counts $a_s(x_i)$ by one and update the values of $P_e(a_s)$, using a constant number of operations. Therefore, the additional complexity of steps (ii) and (iii) is again $O(nD)$. Lastly, to compute the mixture probabilities in step (iv), at each of the (no more than $nD + 1$) nodes of T_{MAX} we perform $O(m)$ operations, so that the overall complexity is $O(nD) + O((nD + 1)m) = O(nmD)$ operations.

A similar argument shows that, as a function of n and D , the complexity of both the BCT and k -BCT algorithms is also $O(nD)$. But as a function of k the complexity of k -BCT grows faster than linearly and increases very substantially as we require more information about the area near the mode of the posterior $\pi(T|x)$ in model space; more details are given in the relevant discussion in Section B of the supplementary material. This, in part, was the motivation for the MCMC algorithms described in Section 3.5.

From the above discussion and the algorithms' descriptions, it is easy to see that the memory requirements of CTW and BCT are of order $O(nmD)$, and for k -BCT of order $O(nmDk)$.

3.7 | Bibliographical remarks

The CTW algorithm was first introduced for data compression in 1993. It was originally described for binary observations ($m = 2$) and only in the special case of $\beta = 1/2$ by Willems et al. (1993c). The general version of CTW for chains on non-binary alphabets and arbitrary β was introduced by Tjalkens et al. (1994), without reference to Bayesian inference. The connection between CTW and the prior predictive likelihood $P_D^*(x)$ established in Theorem 1 was given in the restricted setting $m = 2, \beta = 1/2$ in the unpublished work Willems et al. (1993a), and the outline of a corresponding argument in the case of general β (still only for binary time series) was later described in Willems et al. (2002). The result of Theorem 1 at the level of generality stated here is new, as is the class of prior distributions $\pi_D(T; \beta)$.

A special case of the BCT algorithm, termed the CTM algorithm, was first introduced, again in the context of data compression, by Willems et al. (1993a) and Willems and Volf (1994, 1995), for binary observations ($m = 2$) and only in the special case of $\beta = 1/2$. An extension for arbitrary β (still only for $m = 2$) was later given in Willems et al. (2002); the general version of the BCT algorithm presented here is new. The fact that the BCT algorithm identifies the MAP tree model was established in the restricted setting $m = 2, \beta = 1/2$ in Willems et al. (1993a, 2000), and some generalisations (still restricted to $m = 2$) are discussed in Willems et al. (2002). Theorem 2 at the level of generality stated here is new.

The k -BCT algorithm, Theorem 3, and the MCMC samplers in Section 3.5 are new.

4 | MODEL SELECTION

Here, we compare the model selection performance of the BCT algorithms of Section 3 with the VLMC and MTD approaches described in the Introduction. In the rest of this section, we give brief descriptions of how these three different techniques will be used. In Section 4.1, we compare them on simulated data and discuss some preliminary conclusions. Section 4.2 contains corresponding results on real data.

Bayesian context trees. The BCT framework provides a consistent foundation for learning and evaluating appropriate models for a given data set x . The BCT and k -BCT algorithms can be used to identify the k a posteriori most likely models T_i^* , $i = 1, 2, \dots, k$, for some reasonable k , and we can further explore $\pi(T|x)$ using the MCMC samplers of Section 3.5. With the interpretation of $\pi(T|x)$ as a measure of the ‘truth’ of a model T (Chipman et al., 2001), the posterior probability provides a quantitative confidence measure for the resulting models.

Let $\{X_n\}$ be a variable-memory chain with model $T \in \mathcal{T}(D)$. The specific model T that describes the chain is in general not unique. For example, every independent and identically distributed (i.i.d.) sequence $\{X_n\}$ can also trivially be described as a first-order Markov chain, and adding m children to any leaf of T which is not at maximal depth and giving each of them the same parameters as their parent, leaves the distribution of the chain unchanged. Naturally, the goal in model selection is to identify the smallest model that can fully describe the distribution of the chain. We call a model $T \in \mathcal{T}(D)$ *minimal* with respect to the parameter vector $\theta = \{\theta_s; s \in T\}$, if T is either equal to $\Lambda = \{\lambda\}$ or, if $T \neq \Lambda$, then every m -tuple of leaves $\{sj; j = 0, 1, \dots, m-1\}$ in T contains at least two with non-identical parameters, that is, there are $j \neq j'$ such that $\theta_{sj} \neq \theta_{sj'}$. It is easy to see that every D th-order Markov chain $\{X_n\}$ has a unique minimal model $T^* \in \mathcal{T}(D)$.

The following result, established by Willems et al. (1993a, 1993b) in the context of the MDL principle, offers a partial frequentist justification for the BCT framework. It says that, in the special case of binary chains with $\beta = 1/2$, the MAP model is always eventually correct.

Theorem 4 *Let $\{X_n\}$ be an ergodic, variable-memory chain, with alphabet size $m = 2$ and minimal model $T^* \in \mathcal{T}(D)$. For $\beta = 1/2$, the MAP model $T^*(n)$ based on the random sample X_{-D+1}^n is eventually almost surely unique, and in fact¹:*

$$T^*(n) = T^*, \text{ eventually, a.s.}$$

Variable-length Markov chains. The VLMC class (Bühlmann, 2000; Bühlmann & Wyner, 1999) consists of tree models very similar to those in the BCT class $\mathcal{T}(D)$, except for the fact that VLMC trees are not necessarily proper. The associated VLMC model selection algorithm also has similarities with the pruning procedure of the BCT algorithm: First a version of the tree T_{MAX} is constructed and the count vectors a_s are computed, as in the BCT, and then T_{MAX} is pruned, based on a *cut-off parameter* K that plays a role analogous to β , producing the final model. Theoretical justifications for the resulting VLMC model are provided in Bühlmann and Wyner (1999), where general conditions for asymptotic consistency are established.

The VLMC results in our experiments below are obtained using the implementation in the R package VLMC. Since the algorithm that uses the default value of the cut-off parameter K (‘default-VLMC’) generally gives significantly inferior results, we also examine the results obtained by optimising the choice of K in order to minimise the BIC or the AIC score (‘best-BIC-VLMC’ and ‘best-AIC-VLMC’). But these parameter optimisations are computationally very costly, as we point out in more detail at the end of Section 4.1.

MTD models. The original ‘single-matrix’ version of the MTD model (Raftery, 1985) is based on a different way of parsimoniously representing the D th-order transition probabilities of a Markov chain, as a mixture of the form,

$$P(x_0|x_{-D}^{-1}) = \sum_{d=1}^D \lambda_d Q(x_0|x_{-d}),$$

¹As usual, ‘eventually a.s.’ means that there is a possibly random $1 \leq N < \infty$ such that, almost surely, $T^*(n) = T^*$ for all $n \geq N$.

where $(\lambda_1, \lambda_2, \dots, \lambda_D)$ are lag parameters and $Q = (Q(j|i))$ is a stochastic matrix. Numerical methods for fitting this model via approximate maximum likelihood were developed by Raftery and Tavaré (1994), and a generalisation, the *multi-matrix MTD model*, or MTDg, was introduced by Berchtold (1996, 1998). The MTDg model is based on the more general representation,

$$P(x_0|x_{-D}^{-1}) = \sum_{d=1}^D \lambda_d Q^{(d)}(x_0|x_{-d}),$$

where a different stochastic matrix $Q^{(d)} = (Q^{(d)}(j|i))$ is used for each lag $d = 1, 2, \dots, D$.

In our experiments below, we use the MTD implementation in the R package MARCH. For each data set, we run the MTD algorithm for a range of possible depths D , and choose the value of D that maximises the corresponding BIC or AIC score. We refer to the resulting models as the best-BIC-MTD and best-AIC-MTD models. Similarly we obtain the best-BIC-MTDg and best-AIC-MTDg models from the multimatrix version.

Model comparison. A natural and logically consistent way to compare different models from $\mathcal{T}(D)$ is to compare their posterior probabilities $\pi(T|x)$, but this is not possible for models produced by different methods. In such cases, we follow the common practice of comparing their AIC (Akaike, 1973) and BIC (Schwarz, 1978) scores. BIC is asymptotically consistent when dealing with a fixed number of finite-dimensional parametric models (Hannan & Quinn, 1979; Wei, 1992) whereas AIC is typically asymptotically efficient and minimax optimal in infinite-dimensional, non-parametric settings (Barron et al., 1999; Shibata, 1980); see also Dziak et al. (2019). This dichotomy suggests that, for our purposes, BIC is the more relevant criterion, as confirmed by our findings in Section 4.1 and Section F.1 of the supplementary material.

4.1 | Simulated data

We compare the performance of the three different model selection approaches described above on a relatively short synthetic time series. Analogous comparisons on simulated data from two more chains—one from the VLMC paper (Bühlmann, 2000) and one from the MTD paper (Berchtold & Raftery, 2002)—are carried out in Section F.1 of the supplementary material.

Consider the fifth-order variable-memory chain $\{X_n\}$ on the alphabet $A = \{0, 1, 2\}$ of $m = 3$ letters, with model given by the tree T shown in figure 1 of Section 2.2; the associated parameter vector $\theta = \{\theta_s; s \in T\}$ is given in Section E of the supplementary material.

With x_1^n consisting of $n = 10,000$ simulated observations (x_1, x_2, \dots, x_n) and with a given initial context $x_{-D+1}^0 = (x_{-D+1}, \dots, x_0)$, the MAP model T_1^* obtained by the BCT algorithm with depth $D = 10$ and $\beta = 1 - 2^{-m+1} = 3/4$ is the true underlying model with respect to which the data were generated. Its posterior probability is $\pi(T_1^*|x) \approx 0.3946$ while its prior probability is $\pi(T_1^*) \approx 5.8 \times 10^{-6}$. Although it may seem quite unremarkable that the ‘correct’ model is identified based on a series of 10,000 samples, it is worth noting that, with $D = 10$, there are more than 10^{5900} different models in $\mathcal{T}(D)$.

With $n = 1000$ samples, $D = 10$, and $\beta = 1 - 2^{-m+1} = 3/4$, the five a posteriori most likely models produced by the BCT algorithm are shown in Figure 2. The MAP model is a depth-4 subtree of the true underlying model, with prior probability $\pi(T_1^*) \approx 2.9 \times 10^{-4}$ and posterior $\pi(T_1^*|x) \approx 0.2702$. The true model appears as the fourth most likely tree, T_4^* , with posterior probability $\pi(T_4^*|x) \approx 0.0213$. The sum of the posterior probabilities of the top 5 models is approximately 0.4737.

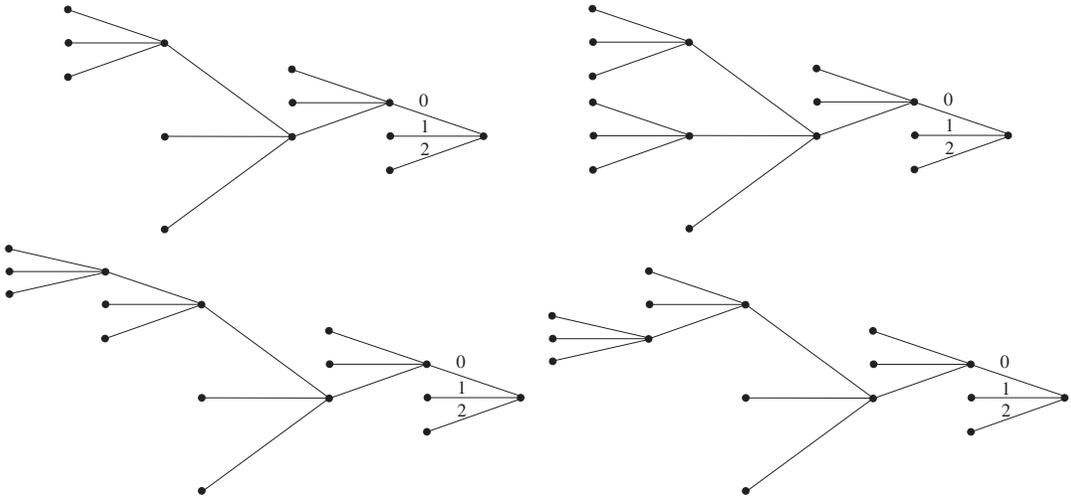


FIGURE 2 The first, second, third and fifth a posteriori most likely models T_1^* , T_2^* , T_3^* , T_5^* produced by k -BCT with $n = 1000$ samples. The posterior odds $\pi(T_i^*|x)/\pi(T_1^*|x)$ with respect to the MAP model T_1^* are approximately 2.369, 5.358, 12.69 and 15.24, for $i = 2, 4, 5$, respectively

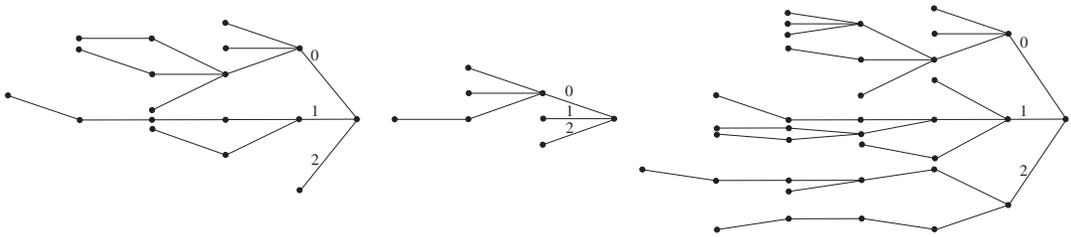


FIGURE 3 The models produced by the default-VLMC, the best-BIC-VLMC and the best-AIC-VLMC methods, with $n = 1,000$ samples

The default-VLMC model is the first tree shown in Figure 3; only about half of its nodes appear in the true underlying model. It has a worse BIC score and a better AIC score than the MAP model. The best-BIC-VLMC produces the small tree of depth 3 shown second in Figure 3, which is a subtree of the true model; it has a good BIC score and a poor AIC score. In sharp contrast, the best-AIC-VLMC produces a clearly overfitted model of depth 6, shown third in Figure 3. Although it has a poor BIC score, it has a very good AIC score, as expected.

Finally, the best-BIC-MTD and the best-BIC-MTDg both give $D = 0$, whereas the best-AIC-MTD gives $D = 3$ and the best-AIC-MTDg gives $D = 2$. Their scores (both BIC and AIC) are generally quite a bit worse than those of the MAP model or the models produced by VLMC.

Overall, the BCT and the best-BIC-VLMC algorithms achieve the best performance in terms of scores, and they learn part of the true underlying tree. VLMC has a marginally better BIC score, whereas the BCT algorithm has a slightly better AIC score (by approximately 1% in both cases). More importantly, the BCT MAP model T_1^* has an additional full branch at depth 4 that reveals more of the true underlying structure, and the k -BCT identifies the true model as T_4^* .

Discussion. In the results on simulated data above and in the two examples in Section F.1 of the supplementary material, the BCT and k -BCT algorithms consistently give the most accurate

model fit, with the best-BIC version of VLMC often giving similar results. The BCT framework has the advantage of identifying not just a single model, but the top k a posteriori most likely models, together with their exact prior and posterior probabilities. And, importantly, in terms of complexity, the BCT algorithm is more efficient than either VLMC or MTD, typically by at least 2 orders of magnitude, since it does not require any tuning.

VLMC. The default and best-AIC versions of VLMC generally gave similar or identical models, usually much larger than the true underlying model, in rather typical examples of overfitting. The best-BIC-VLMC was found to be much more accurate in revealing significant parts of the true model, as expected in view of the earlier AIC-vs-BIC discussion. The resulting models were smaller, which is consistent with the observation that optimising the BIC score led to larger values for the cut-off parameter K and more frequent pruning. For the best-AIC and best-BIC versions, we executed VLMC approximately 500 times with different values of $K \in [1, 10]$. Although a smaller range ($2.8 \leq K \leq 6$) was used by Mächler and Bühlmann (2004), we found it often necessary to look further; for example for the financial data in Section F.2 of the supplementary material the best value was found to be close to 10, and for the genetic data in Section 4.2 it was larger than 20. In most cases, using fewer runs resulted in different models giving poor fits.

BCT. The MAP models produced by the BCT algorithm were usually similar to those produced by the best-BIC-VLMC, they had good AIC and BIC scores, and they generally learned the most accurate approximations of the underlying model among all methods considered. The additional models obtained by k -BCT offered further indications of the type of structure present in the data, and they were accompanied by posterior probabilities, indicating the level of ‘posterior confidence’ we may have in these models. Also, BCT and k -BCT only require a single run with the default value of the hyperparameter β .

MTD. None of the MTD methods performed as well as the BCT or VLMC algorithms, partly because they gave (by design) less flexible models. As their only output in terms of the model is the Markov order D and the number of nonzero lag parameters λ_d , which can only take a few discrete values, the best-BIC and best-AIC results were often same. MTDg in most cases had worse scores than MTD, but even the MTD’s overall performance was not competitive with that of BCT and VLMC. Finally, MTD had very high complexity, since the implicit maximum likelihood computation is over a non-convex space defined by a large number of non-linear constraints (Raftery & Tavaré, 1994). For this reason, it appears infeasible in practice to use values for D significantly larger than $D = 10$.

4.2 | Real data

In view of the above discussion, for the comparisons in the real-world data examples we only consider the best-BIC versions of VLMC, MTD and MTDg. One more example with a financial time series is given in Section F.2 of the supplementary material.

SARS-CoV-2 genome. The severe acute respiratory syndrome coronavirus 2, *SARS-CoV-2*, is the novel coronavirus responsible for the COVID-19 global pandemic in 2019-20. Here we examine the SARS-CoV-2 genome, available in the GenBank database (Clark et al., 2016) as the sequence MN908947.3 identified in Wu et al. (2020). It consists of $n = 29,903$ base pairs, and we translate the four-letter DNA alphabet to $\{0, 1, 2, 3\}$ via the map $(A, C, G, T) \mapsto (0, 1, 2, 3)$.

The top three models obtained by the k -BCT algorithm with $D = 10$, $\beta = 1 - 2^{-m+1} = 7/8$ and $k = 3$ are shown as the first three trees in Figure 4. The MAP model has posterior probability

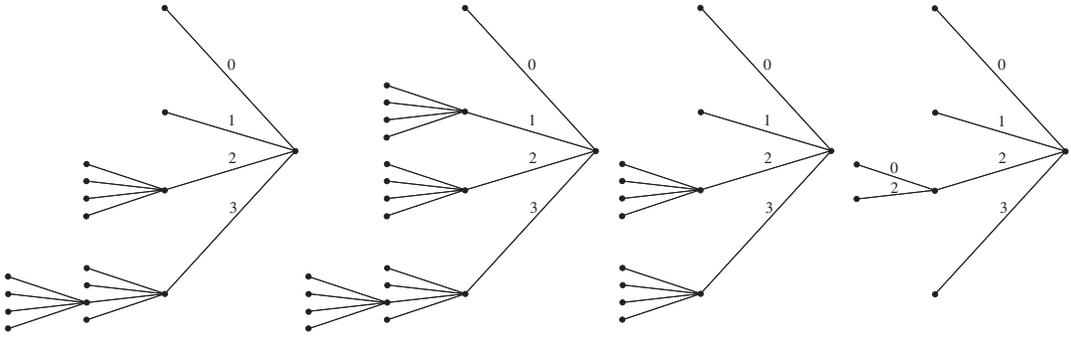


FIGURE 4 First three trees: The a posteriori most likely models T_1^* , T_2^* , T_3^* obtained by the k -BCT algorithm on the SARS-CoV-2 genome. The posterior odds $\pi(T_1^*|x)/\pi(T_2^*|x)$ and $\pi(T_1^*|x)/\pi(T_3^*|x)$ are approximately 35.75 and 101.4, respectively. Last tree: The best-BIC-VLMC model for the SARS-CoV-2 genome; its AIC and BIC scores are both within 0.1% of the corresponding scores of the BCT model T_1^*

$\pi(T_1^*|x) \approx 0.963$ and prior $\pi(T_1^*) \approx 4.3 \times 10^{-5}$. The sum of the posterior probabilities of these three models is ≈ 0.9994 .

The VLMC model is the depth-2 subtree of T_1^* shown last in Figure 4. The optimisation of the cut-off parameter K required to find the best-BIC model, involved a rather laborious search, resulting in the choice $K = 22.14$. Both MTD and MTDg give $D = 1$ as the optimal depth, corresponding to a simple first-order Markov chain.

The AIC and BIC scores of all models generated by all three approaches are within 0.3% of each other. It is interesting that the MAP model has such high posterior probability, and that k -BCT gives models of depth 3 with very high confidence. Although it is not possible to otherwise verify the significance of the bigger depth of the BCT models T_i^* compared to those produced by the other methods, it may be that the BCT finds evidence of the fact that DNA naturally gets encoded into triplets of bases to form codons that specify particular amino acids.

Pewee birdsong. The twilight song of the wood pewee bird, studied extensively in Craig (1943), can be described as a sequence consisting of an arrangement of musical phrases taken from an alphabet of three specific, distinct phrases. The data in this section consist of a single contiguous song by a wood pewee, of length $n = 1327$ phrases, first recorded by Craig (1943). It was analysed by Raftery and Tavaré (1994) and Berchtold (2001) using MTD models, by Kharin and Piatlitski (2011) using a different MTD-type model representation, and by Sarkar and Dunson (2016) using CTF models; it is also contained in the R package MARCH.

It is well known that the pewee birdsong contains significant variability but it is also fairly structured, with specific repeating patterns. The most common of these, described by Saunders (1944) as ‘the commonest and most pleasing sentence’, is the string $s^* = 0201$, which dominates the data set, appearing 266 times and occupying 1064 positions or a little over 80% of the data.

The MAP tree T_1^* obtained by k -BCT with $m = 3$, $\beta = 1 - 2^{-m+1} = 3/4$, $D = 10$ and $k = 5$ is shown in Figure 5. Its prior probability is $\pi(T_1^*) \approx 4.1 \times 10^{-5}$, and its posterior $\pi(T_1^*|x) \approx 0.1244$.

Clearly the importance of $s^* = 0201$ is captured in the MAP tree T_1^* , which includes the entire prefix 020 necessary to predict the appearance of $s^* = 0201$. The VLMC produces an interesting model, also shown in Figure 5, which is quite similar to T_1^* . Although it does not include the 020 context, it does include the context 02, which contains most of the predictive power regarding $s^* = 0201$: While s^* appears 266 times in the data, the context $s = 201$ appears 267 times, therefore, we can ‘statistically’ almost identify s^* with s . Compared to the VLMC model, the MAP tree T_1^* has

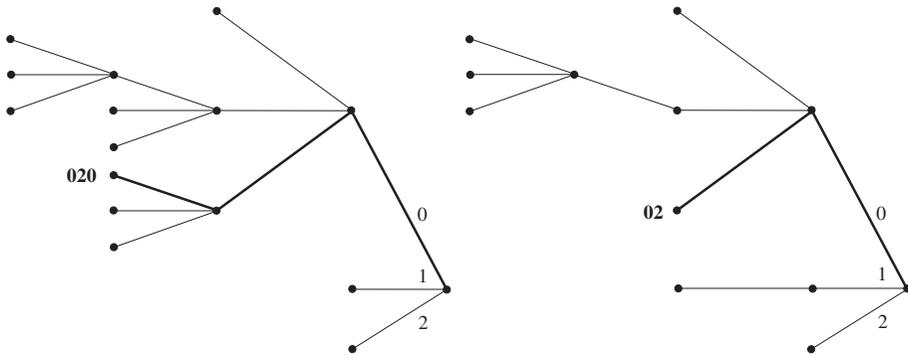


FIGURE 5 Left: The MAP model T_1^* for the pewee birdsong data. Right: The corresponding VLMC model

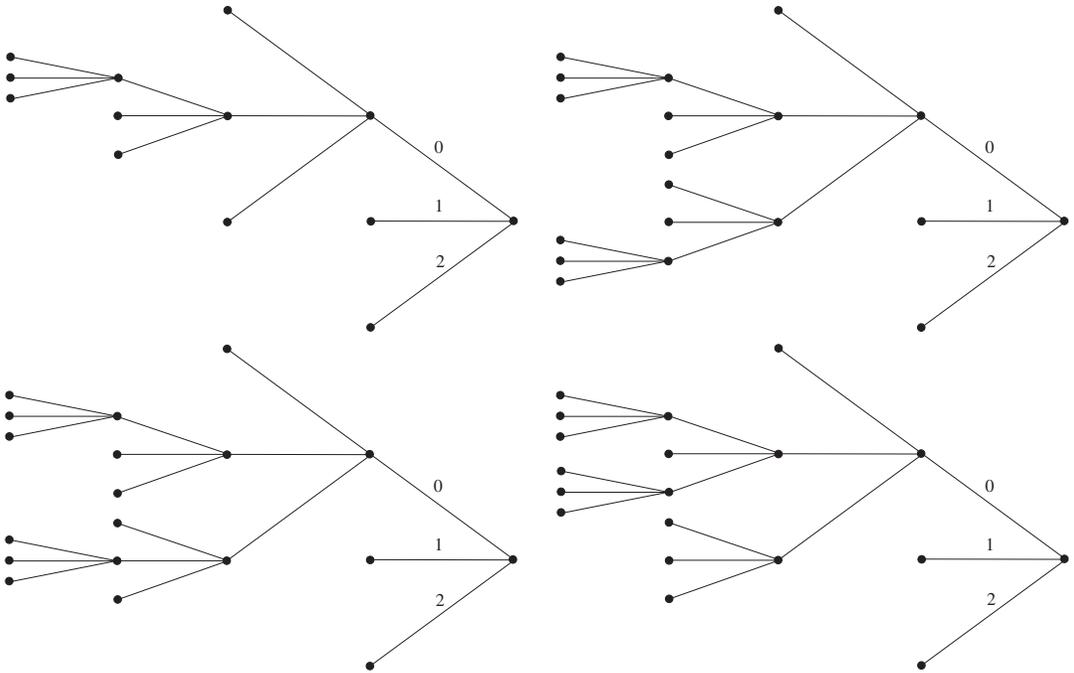


FIGURE 6 The four a posteriori most likely models T_i^* , $i = 2, 3, 4, 5$, after T_1^* , obtained by the k -BCT algorithm from the pewee birdsong data. The corresponding posterior odds are $\pi(T_1^*|x)/\pi(T_2^*|x) \approx 5.727$ and $\pi(T_1^*|x)/\pi(T_i^*|x) \approx 7.111$, for $i = 3, 4, 5$. The sum of the posteriors of the top five models is ≈ 0.1985

a marginally better AIC score and a slightly worse BIC score, but it is clear that the two methods learn much of the same structure from the data.

The four a posteriori most likely models after T_1^* are shown in Figure 6. They are all quite similar to the MAP tree, each one differing from T_1^* in only a single branch.

MTD gives $D = 6$ as the optimal depth, and MTDg gives $D = 3$. The resulting MTDg model has better AIC and BIC scores than the one from MTD, but they are still much worse (by 39% for AIC and 27% for BIC) than those of the MAP model. This confirms the findings of both Raftery

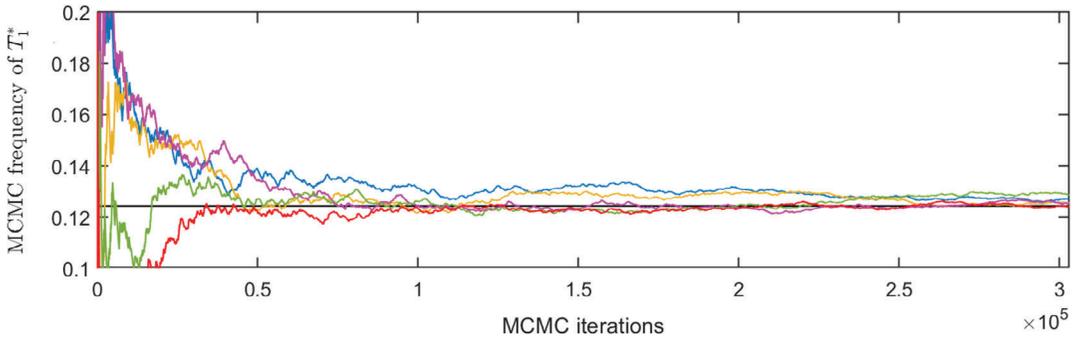


FIGURE 7 Markov chain Monte Carlo frequency of T_1^* . The five graphs plotted correspond to five independent repetitions of the experiment with $N = 3 \times 10^5$ iterations; the horizontal line is the limiting frequency, $\pi(T_1^* | x)$

and Tavaré (1994) and Berchtold (2001), where it was noted that the MTD family of models is not appropriate for this data set, in large part due to the high significance of individual patterns like s^* .

The top five models here carry less than 20% of the total posterior mass. Therefore, in order to get a better sense of the posterior distribution on model space, we employed the RW MCMC sampler described in Section 3.5 to produce $N = 10^6$ samples from $\pi(T|x)$. The acceptance rate was 57.8%, a total of 274,721 unique trees were visited, and the sum of their posterior probabilities was 61.2%. The MCMC frequency of T_1^* in Figure 7 indicates that the sampler converged quite quickly.

The 100 most visited models have a total posterior probability of 38.3%. They all have depths between 4 and 7, with 48 of them having depth 5. Together with the results of k -BCT and VLMC, this suggests that there is significant fourth- and fifth-order structure in the data.

Neural spike trains. Next we consider a long binary time series x that consists of $n = 3,919,361$ bits, describing the spike train recorded from a single neuron in the V4 region of a monkey's brain. The recording was made during the experiment described in Gregoriou et al. (2009, 2012), while the monkey was performing an attention task. The recorded signal was discretised into one-millisecond bins (with $x_i = 1$ if there was a spike in bin i and $x_i = 0$ otherwise), corresponding to a trial lasting a little over 65 min. As we have not been able to find implementations of VLMC or MTD that can operate on data sets of this length, in this section, we only present the results obtained by the BCT and k -BCT algorithms.

With $D = 100$, $\beta = 1 - 2^{-m+1} = 1/2$ and $k = 5$, the MAP model T_1^* is shown in Figure 8. It has depth 98 and, with the exception of two additional branches at depths 9 and 90, it is very similar to a renewal model, qualitatively similar to the first chain in Section F.1 of the supplementary material. Its prior probability is $\pi(T_1^*) \approx 3.1 \times 10^{-61}$ and its posterior $\pi(T_1^* | x) \approx 2.1 \times 10^{-8}$. Although this probability is small, we note that there are more than 79^{1029} possible models of depth no more than 100, cf. (9), and that this posterior is still larger than the corresponding prior probability by more than 50 orders of magnitude. Therefore, the observations x offer significant support for this model.

The next four a posteriori most likely models shown in Figure 8 are also very similar to simple renewal models, offering a partial justification to the biological intuition behind the elementary Poisson/renewal models commonly used in neuroscience (Dayan & Abbott, 2001; Rieke et al., 1999), and confirming relevant earlier findings (Gao et al., 2006, 2008). In other words, we have learned from the data that the statistically most significant factor in determining

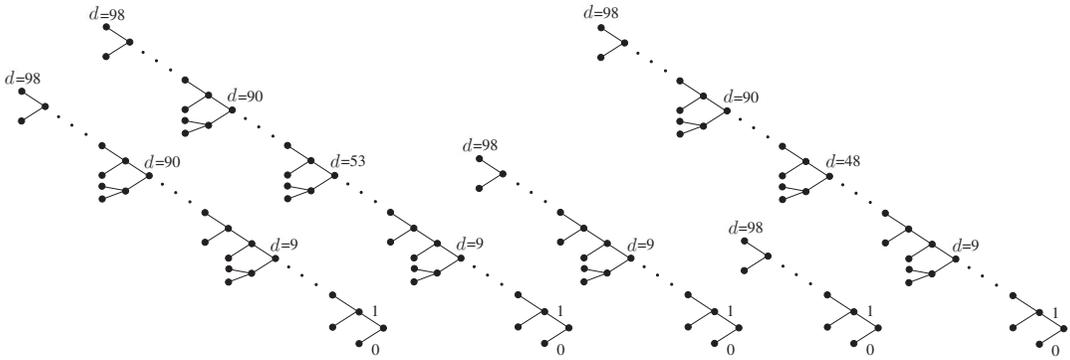


FIGURE 8 The top $k = 5$ models obtained by k -BCT on the spike train data. The posterior odds $\pi(T_1^* | x) / \pi(T_i^* | x)$ for $i = 2, 3, 4, 5$, are $\approx 1.2, 1.43, 1.52$ and 1.53 , respectively

whether a neuron will fire, given its past history, is the time when it most recently fired before.

Although the sum of the posterior probabilities of the top five models is less than 10^{-7} , in this case it would not make sense to employ an MCMC sampler to explore the posterior further. The reason is that, given the very small values of the posterior probabilities of the top five models, even with 10^6 MCMC samples visiting 10^6 distinct models, we would still only visit around 1% of the support of the posterior, at best. On the other hand, increasing the value of k can give a better idea of the shape of the posterior near its mode T_1^* . With $k = 50$, k -BCT produced 50 trees, all of depth 98, and all of them being small variations of the renewal model T_4^* : All the resulting models T_i^* had between one and five additional branches at various depths. The sum of their posterior probabilities is $\approx 4.1 \times 10^{-7}$.

As a final test of the scalability of the BCT algorithm on a large data set, we obtained the MAP tree with maximum depth $D = 1500$. Interestingly, it is the same as the MAP tree for $D = 100$, and with only a slightly smaller posterior probability of $\approx 1.6 \times 10^{-8}$.

5 | POSTERIOR EXPLORATION AND ESTIMATION

We present two examples that illustrate the utility of the MCMC samplers in Section 3.5 for posterior exploration, and the application of the BCT methodology to parameter estimation and Markov order estimation.

Example 1 (Daily changes in S&P 500). We consider the daily changes in Standard & Poor's index, from 2 January 1928 until 7 October 2016 (available at <http://www.finance.yahoo.com/quote/GSPC>), quantised to $m = 7$ values: If the change between two successive trading days, day $(i - 1)$ and day i , is smaller than -3% , x_i is set equal to 0; for changes in the intervals $(-3\%, -2\%]$, $(-2\%, -1\%]$, $(-1\%, +1\%]$, $(1\%, 2\%]$, and $(2\%, 3\%]$, x_i is set equal to 1, 2, 3, 4 and 5, respectively; and for changes greater than 3% , $x_i = 6$.

Based on the resulting $n = 22900$ points x_i , the top $k = 5$ a posteriori most likely models obtained by the k -BCT algorithm with maximum tree depth $D = 260$ (corresponding to approximately one calendar year's trading days), are described in Figure 9.

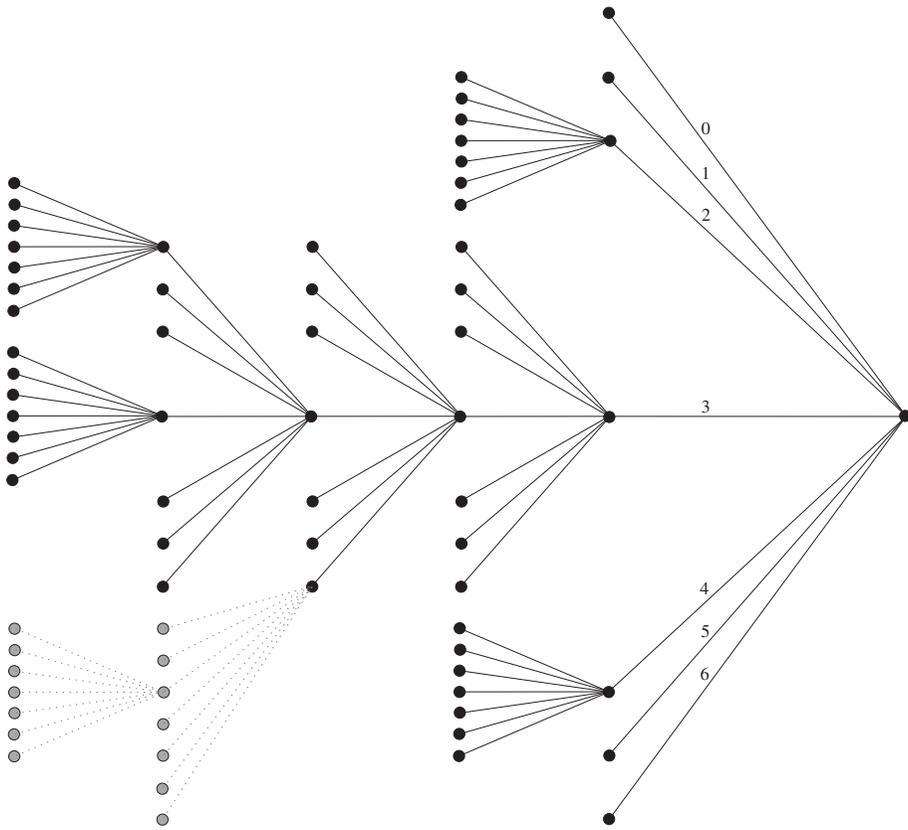


FIGURE 9 The tree shown *without* the two dotted branches is the MAP model T_1^* obtained from $n = 22900$ observations, with $D = 260$. Its posterior probability $\pi(T_1^*|x) \approx 0.0174$ and its prior $\pi(T_1^*) \approx 5.7 \times 10^{-11}$. The whole tree shown is the fifth a posteriori most likely model T_5^* , and T_2^* , T_3^* and T_4^* were found to be small variations around T_1^* and T_5^* , all with depth 5. The corresponding posterior odds $\pi(T_i^*|x)/\pi(T_1^*|x)$, for $i = 2, 3, 4, 5$, are 1.094, 1.367, 1.496 and 2.467, respectively

The shape of the MAP model T_1^* contains significant information. Since its maximal depth is 5, in order to determine the distribution of the next sample we never have to look more than five days back—corresponding to a week of trading days. The smaller the changes in the most recent S&P values, the further back we need to look in order to predict tomorrow's value. For example, if the difference between yesterday and today is larger than $\pm 2\%$, we need look no further than yesterday; if it is between 1% and 2%, we need to look at the day before yesterday as well; and if it is smaller than $\pm 1\%$, we need to look even further back, but no more than a week.

The sum of the posterior probabilities of the top $k = 5$ models is less than 6.5%. But with $n = 22900$ data points, $m = 7$ and $D = 260$, the complexity of k -BCT becomes prohibitive for large values of k . In order to explore $\pi(T|x)$ further, we ran the RW sampler with $T^{(0)} = T_1^*$ for $N = 10^6$ iterations. The acceptance rate was $\approx 44.9\%$, and a total of 356531 different models were visited. The MCMC frequencies of the 25 most visited models shown in Figure 10 indicate that the sampler has converged after $N = 10^6$ iterations.

The MCMC output can also be used for Markov order estimation, by providing an approximation to the posterior distribution on model depth. The empirical distributions of the model depths obtained in five repetitions of the same experiment are shown in Figure 10.

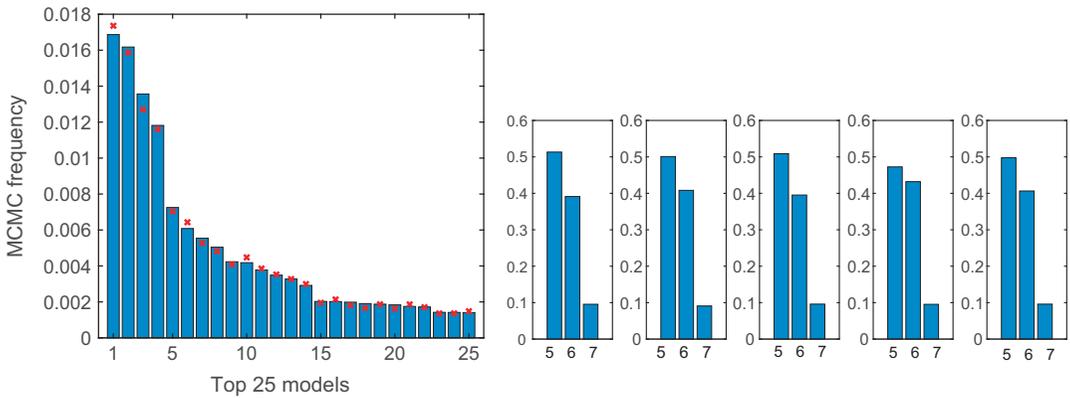


FIGURE 10 Left: Markov chain Monte Carlo (MCMC) empirical frequencies of the 25 most frequently visited models, after $N = 10^6$ iterations. The corresponding true posterior probabilities are marked with a red ‘x’. Right: Markov order estimation: The five histograms show the empirical frequencies of the depths of the 1000 most visited models after 10^6 MCMC iterations, in five independent repetitions of the same experiment

Example 2 (A bimodal posterior). Consider a third-order chain $\{X_n\}$ on the alphabet $A = \{0, 1, 2, 3, 4, 5\}$ with the property that each X_n depends on the past $(X_{n-1}, X_{n-2}, X_{n-3})$ only via X_{n-3} . Specifically, suppose that, for $i, j, a, b \in A$, $\Pr(X_n = j | X_{n-1} = a, X_{n-2} = b, X_{n-3} = i) = Q_{ij}$, where the transition matrix $Q = (Q_{ij})$ is given in Section E of the supplementary material.

The model of $\{X_n\}$ viewed as a variable-memory chain is clearly the complete tree of depth 3, but the dependence of each X_n on its past is only meaningful if it can extend at least three time steps back: The first and second most recent symbols are independent of X_n .

Therefore, it is not surprising that the MAP model identified by the k -BCT algorithm (with $k = 5$, based on $n = 1850$ samples) is simply the root, $T_1^* = \Lambda = \{\lambda\}$, with T_2^* a complex tree of depth 3 (with 54 leaves at depth 3) being a close second; their posterior probabilities are $\pi(T_1^* | x) \approx 0.3512$ and $\pi(T_2^* | x) \approx 0.2373$, respectively. The next three a posteriori most likely models T_3^*, T_4^*, T_5^* were found to be small variations of T_2^* , also of depth 3.

Running the RW sampler with $T^{(0)} = T_1^*$, we found that it never visited any of the other top five models after 10^6 iterations, and similarly starting at T_2^* it never visited T_1^* . Although T_2^* can theoretically be reached from T_1^* in just 12 MCMC steps, most models between them have extremely small posterior probabilities; for example the complete tree of depth one, $T_c(1)$, which must necessarily be visited in order to move between T_1^* and T_2^* , has $\pi(T_c(1) | x) \approx 3.2 \times 10^{-19}$.

In contrast, the jump sampler with jump parameter $p = 1/2$, made frequent jumps between T_1^* and $\{T_2^*, T_3^*, T_4^*, T_5^*\}$. Starting with $T^{(0)} = T_1^*$, after $N = 10^5$ MCMC iterations it appears to have explored the bulk of the posterior distribution, having visited all of the significant parts of its support. The empirical frequencies of the top five models were very close to their actual posterior probabilities (Figure 11), the total number of unique models visited was 39, and the sum of their posterior probabilities was $\approx 99.8\%$.

In terms of Markov order estimation, although the MAP model T_1^* has depth 0, the mode of the posterior distribution of the depth parameter is 3: The $N = 10^5$ MCMC samples obtained above consist entirely of models having depths either 0 or 3, with corresponding MCMC frequencies of approximately 35.13% and 64.87%, respectively.

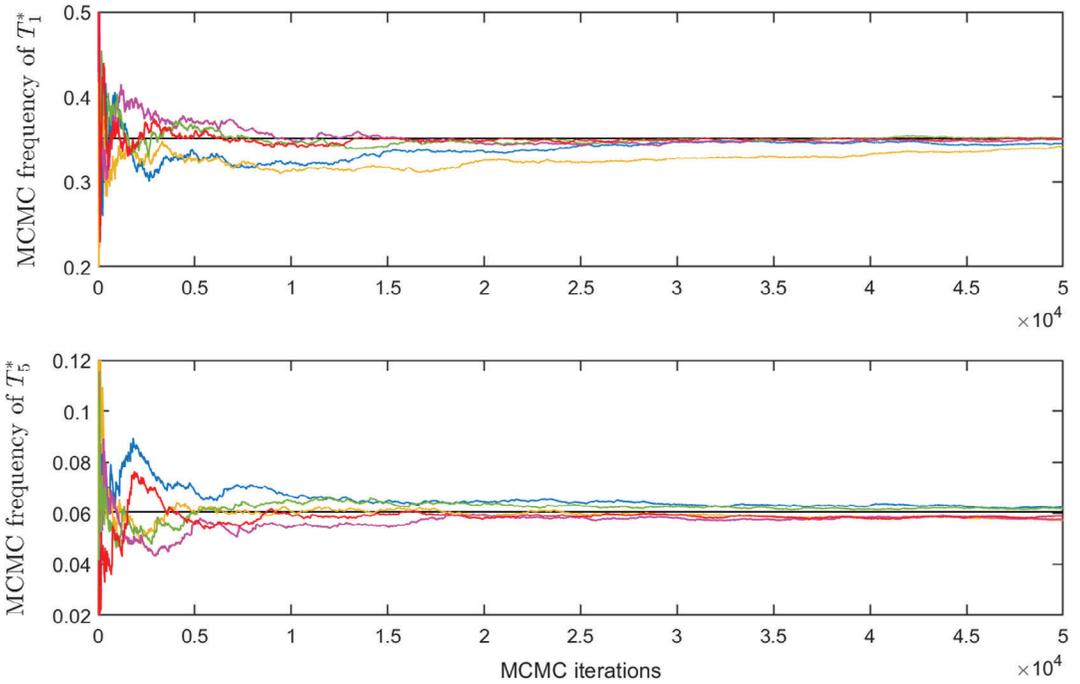


FIGURE 11 Results of the jump sampler. The top plot shows the empirical frequency with which T_1^* appears in the Markov chain Monte Carlo (MCMC) sample, as a function of the number of MCMC iterations. The five graphs plotted correspond to five independent repetitions of the experiment with 50,000 iterations; the horizontal line is the limiting frequency, $\pi(T_1^*|x)$. The bottom plot shows corresponding results for T_5^*

More generally, for any statistic $F(\theta, T)$, we can obtain MCMC samples $\{(\theta^{(t)}, T^{(t)})\}$ using the joint sampler of Section 3.5, and use the values $\{F(\theta^{(t)}, T^{(t)})\}$ to provide an approximation to the posterior $\pi(F|x)$. Standard Bayesian methods (Bernardo & Smith, 1994; Geisser, 1993; Gelman et al., 2014) can then be applied to provide point estimates, credible sets, and other relevant information.

For example, suppose we wish to estimate the parameter $\theta_s(j) = \Pr(X_0 = j | X_{-D}^{-1} = s)$, for the specific context $s = 020$ and $j = 5$. The maximum likelihood estimate (MLE) is $\hat{\theta}_s^{\text{MLE}}(j) = a_s(j)/M_s$, and a commonly used Bayesian counterpart to the MLE, $\hat{\theta}_s^{\text{MAP}}(j)$, is the mode of the full conditional density $\pi(\theta_s(j)|x, T)$ in Equation (16) of $\theta_s(j)$ given the data x and the MAP model $T = T_1^*$. Another Bayesian alternative to the MLE is the posterior mean, which can be approximated as,

$$\hat{\theta}_s^{\text{MCMC}}(j) := \frac{1}{N} \sum_{t=1}^N \theta_s^{(t)}(j),$$

and an estimator with smaller variance can also be obtained via Rao-Blackwellization (Gelfand & Smith, 1990), using (16):

$$\hat{\theta}_s^{\text{RB}}(j) := \frac{1}{N} \sum_{t=1}^N E(\theta_s(j)|x, T^{(t)}).$$

In this example, we obtain the values:

$\hat{\theta}_s^{\text{MLE}}(j)$	$\hat{\theta}_s^{\text{MAP}}(j)$	$\hat{\theta}_s^{\text{MCMC}}(j)$	$\hat{\theta}_s^{\text{RB}}(j)$	true $\theta_s(j)$
0.1290	0.1871	0.1512	0.1516	0.15

The estimates $\hat{\theta}_s^{\text{MCMC}}(j)$ and $\hat{\theta}_s^{\text{RB}}(j)$ are based on $N = 10^5$ MCMC samples obtained by the jump version of the joint sampler, with jump parameter $p = 1/2$. From the same samples, we get an estimate of the posterior distribution of $\theta_{020}(5)$, shown in Figure 12.

Note that the values of the four estimates above are quite different, which is common, especially for small data sizes n . On the other hand, the fact that the Bayesian estimates are more accurate than the MLE is by no means universal.

Finally, in Figure 12, we also show the results of five different MCMC experiments, indicating that the ergodic averages in the estimators $\hat{\theta}_s^{\text{MCMC}}(j)$ and $\hat{\theta}_s^{\text{RB}}(j)$ converge quickly, and that the variance of $\hat{\theta}_s^{\text{RB}}(j)$ appears to be smaller, as expected.

6 | PREDICTION

Given a *training sequence* $x_1^t = (x_1, x_2, \dots, x_t)$ of length t from a discrete time series x with values in the alphabet $A = \{0, 1, \dots, m-1\}$, we wish to sequentially predict the next n values of the *test sequence* $x_{t+1}^{t+n} = (x_{t+1}, x_{t+2}, \dots, x_{t+n})$. At each step $i = 1, 2, \dots, n$, given the past samples x_1^{t+i-1} , the prediction of the next sample x_{t+i} is expressed as a conditional distribution $\hat{Q}_i(z|x_1^{t+i-1})$, $z \in A$, and the performance of a predictor $Q = \{\hat{Q}_i; i = 1, 2, \dots, n\}$ is evaluated by the normalised, cumulative log-loss.

In the remainder of this section, we describe the four prediction methods that will be used below, together with the natural predictor induced by the BCT framework discussed in Section 3.6. In Sections 6.1 and 6.2, we report the results obtained by all five methods on simulated and real data sets, and we discuss them in Section 6.3.

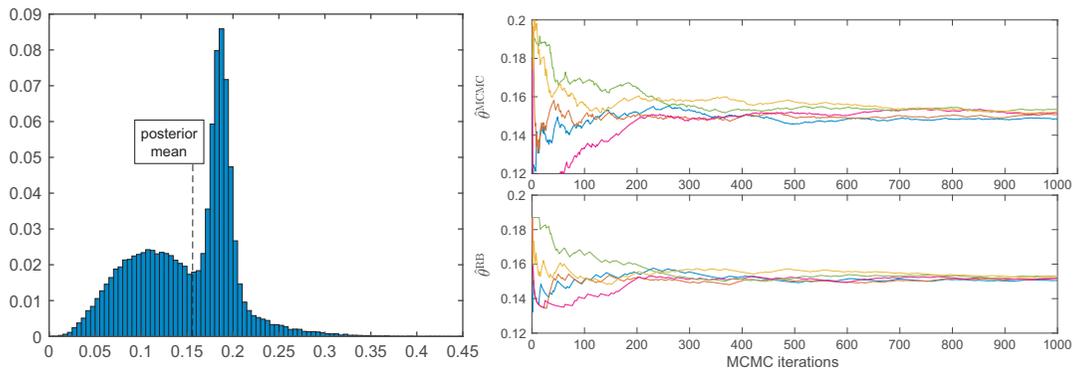


FIGURE 12 Left: Histogram of the $N = 10^5$ Markov chain Monte Carlo (MCMC) samples $\theta_s^{(j)}(j)$, with $D = 3$, $\beta = 0.95$, $s = 020$ and $j = 5$. Right: Corresponding MCMC estimates. The top plot shows the ergodic averages $\hat{\theta}_s^{\text{MCMC}}(j)$ as a function of the number of MCMC iterations. The five graphs plotted correspond to five independent repetitions with only $N = 1000$ MCMC iterations each. The bottom plot shows corresponding results for $\hat{\theta}_s^{\text{RB}}(j)$

A common approach to prediction is to first use the training data to learn a model and associated parameters, and then use the conditional distributions of this model as predictors. Indeed, this is exactly the form of the predictors proposed by all methods, other than BCT, described below.

Bayesian context trees. As outlined in Section 3.6, the canonical predictor within the BCT framework is the one based on the posterior predictive distribution, $\hat{Q}_i(x_{t+i}|x_1^{t+i-1}) = P_D^*(x_{t+i}|x_1^t)$, cf. (17), which can easily be computed sequentially via the CTW algorithm. By averaging over all models, $P_D^*(x_{t+i}|x_1^t)$ incorporates the model uncertainty and it avoids the problem of model selection by replacing it with model averaging.

When the data x are generated by a stationary, irreducible variable-memory chain, in the special case $\beta = 1/2$ it is shown by Jiao et al. (2013) that the BCT predictor is consistent with probability one, asymptotically in the size of the test data, for any finite training sequence. Moreover, as the following discussion indicates, the BCT predictor essentially achieves the optimal minimax regret with respect to log-loss. Theorem 5 gives a nonasymptotic lower bound, for the special case of binary chains and $\beta = 1/2$; it was established in Willems et al. (1993b, 1995).

Theorem 5 *For any binary, variable-memory chain model $T \in \mathcal{T}(D)$ and associated parameters $\theta = \{\theta_s; s \in T\}$, for any data sequence x_1^n of arbitrary length n , and any initial context x_{-D+1}^0 , the prior predictive likelihood for $\beta = 1/2$ satisfies,*

$$\log P_D^*(x_1^n|x_{-D+1}^0) \geq \log P(x_1^n|x_{-D+1}^0, \theta, T) - \frac{|T|}{2} \log n + C, \quad (19)$$

with an explicit constant $C = C(T)$, independent of n and θ .

The lower bound (19) is asymptotically tight up to and including the $\log n$ term, both in expectation and for individual sequences: Corresponding upper bounds follow from the fundamental ‘converse’ theorem by Rissanen (1984, 1986b), and from the general results by Weinberger et al. (1994). These results clearly indicate that the BCT predictor indeed essentially achieves the optimal minimax regret (Takeuchi & Barron, 2014).

Variable length Markov chains. For prediction using variable-length Markov chains we use the prediction function in the R package VLMLC. This selects a VLMLC model based on the training data, and then predicts future samples by the maximum likelihood parameters associated with this model. Here we only show results for the best-BIC-VLMLC model, as it was found to be the most effective choice in practice; see also the relevant comments in Section 4.1. When the data x are generated by a stationary and ergodic variable-memory chain, the results of Bühlmann and Wyner (1999) imply that the VLMLC predictor is asymptotically consistent, but for consistency it is the size of the training data that needs to grow to infinity. Further methodology in connection with VLMLC prediction is developed in Bühlmann (2000).

Mixture transition distribution. The MTD approach to prediction is again based on first selecting a model and then performing prediction using an approximation to the maximum likelihood parameters associated with that model. Among the four MTD versions discussed in Section 4, we only report results obtained by the best-BIC-MTDg, since the performance of the other three methods was either identical or inferior.

Sparse Markov chains. As described in the Introduction, SMCs are a generalisation of the BCT models in $\mathcal{T}(D)$: Each SMC model consists of a partition of the set of all contexts of depth D into different states, such that all contexts in the same state induce the same conditional distribution on the next symbol of the underlying chain. In Jääskinen et al. (2014), a class of priors for SMCs

are defined, and algorithms for selecting approximate MAP versions of models and parameters are developed in the subsequent work Xiong et al. (2016). In our experiments, prediction is performed based on the resulting model and parameters obtained using the code publicly available at <https://www.helsinki.fi/bsg/filer/SMCD.zip>.

Conditional tensor factorisation. The CTF models of Sarkar and Dunson (2016), described in the Introduction, use CTF to represent the full transition probability tensor of a higher-order chain. In our experiments, we use the CTF software publicly available at <https://github.com/david-dunson/bnphomc>. Prediction is again performed in two steps. First, a model is selected via the results of an MCMC sampler on model space, and appropriate parameters are chosen by an approximation to their posterior mean via Gibbs sampling. The induced predictive distributions are then obtained from the selected model and parameters.

Throughout our experiments, we take the maximal depth to be $D = 10$ for BCT, MTD, SMC and CTF. Following standard practice (Begleiter et al., 2004) in most cases, we split the data 50-50 into a training set and a test set. One difficulty that occasionally arises with the VLMC and MTD predictors is that they use maximum likelihood parameter estimates for their models, which in some cases means that they assign zero conditional probability to certain symbols, and which in turn results in poor performance and an infinite log-loss.

Finally, we note that different versions of the CTW algorithm have been used for prediction in earlier work, including Ron et al. (1996); Begleiter et al. (2004); Dimitrakakis (2010).

6.1 | Simulated data

The experiments here are based on 1,000 samples simulated from the fifth-order chain with alphabet size $m = 3$ in Section 4.1. Figure 13 shows the log-loss achieved by all five predictors as a function of the number of predicted samples in three different cases.

The results in the first plot are based on $t = 100$ training samples and $n = 900$ test samples. The BCT is seen to perform consistently better than the other four methods, and the difference between BCT and the second best method, CTF, increases as more data get predicted, reaching a significant $\approx 3.7\%$ by the end. This is in part due to the fact that the BCT predictor continues to get updated past the training stage, whereas the models employed by the other four methods based on only 100 training samples are too simple to be effective. SMC, VLMC and MTD all produce the empty model corresponding to i.i.d. data, with VLMC and MTD also having the same parameters, inducing the exact same predictors (hence shown as a single graph). CTF, on the other hand, uses a first-order Markov model, leading to slightly better performance.

The second plot shows the log-loss obtained with $t = 500$ training and $n = 500$ test samples. The results are similar, however, the difference between BCT and the other methods is smaller. Since the training set is larger, the other methods have a better chance to capture more of the structure present in the data. The VLMC and MTD models both correspond to simple, second-order chains. MTD again produces an i.i.d. model, and SMC, which is the closest to BCT, with a log-loss difference of $\approx 2.5\%$, identifies a second-order model with five states.

With $t = 900$ training and $n = 100$ test samples, BCT again appears to have the best performance, although there are larger fluctuations due to the smaller test data size. The results of SMC, VLMC and CTF are almost identical, and significantly better than MTD. SMC uses a second-order model with four states, the CTF model also has order 2, MTD selects the i.i.d. model, and VLMC, which is closest to BCT by the end (by a 2.2% difference), produces a tree model of depth 3 with 5 leaves.

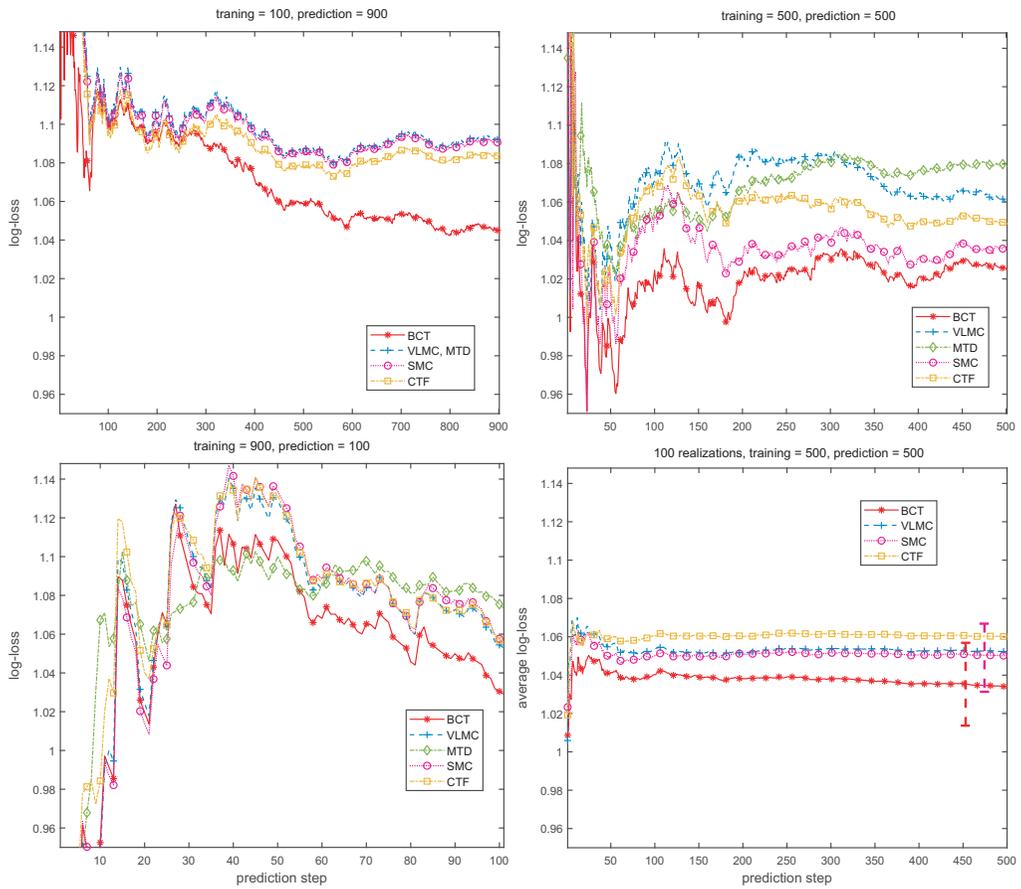


FIGURE 13 Prediction results on simulated data. First three plots: Log-loss achieved by each of the five methods, as a function of the size of the test data. Bottom right: Log-loss achieved by BCT, VLMLC, SMC and CTF on $n = 500$ test samples with $t = 500$ training samples, averaged over 100 independent repetitions of the same experiment. One-standard-deviation error bars are also plotted for BCT and SMC near the end of the test data. The vertical (log-loss) axis has the same range in all four plots

Finally, we examine the log-loss of these predictors averaged over 100 independent realisations simulated from this chain, each time with a 50-50 split between training and test data. The last plot in Figure 13 shows the results obtained by all methods except MTD, which was based on an i.i.d. model every time. The average performance of BCT is better than that of the other methods; after the first 10 samples or so, the log-loss of the BCT stays consistently approximately 1.6% lower than that of the second best method, SMC.

6.2 | Real data

SARS-CoV-2 gene. Here we examine the spike (S) gene, in positions 21,563–25,384 of the SARS-CoV-2 genome; cf. Section 4.2. The importance of this gene is that it codes for the surface glycoprotein whose function was identified in Yan et al. (2020); Lan et al. (2020) as critical, in that it binds onto the angiotensin-converting enzyme 2 (ACE2) receptor on human epithelial cells, giving the virus access to the cell and thus facilitating the COVID-19 disease.

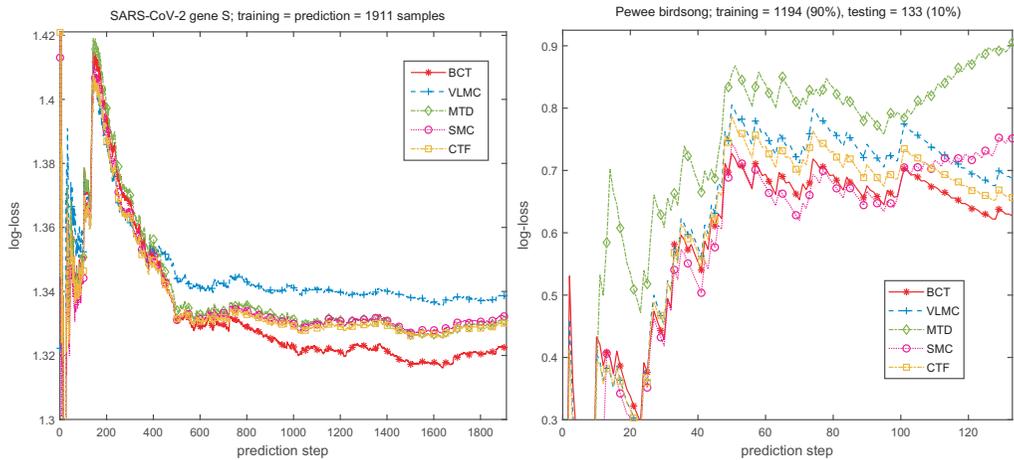


FIGURE 14 Prediction results on real data. Left: SARS-CoV-2 gene (S), with $t = 1911$ training samples and $n = 1911$ test samples. Right: Pewee birdsong data, with $t = 1194$ training samples and $n = 133$ test samples, corresponding to a 90%-training/10%-testing split

The data, consisting of a 3822 bp-long gene sequence, was again split 50-50 into a training set and test set. Figure 14 shows the prediction results obtained by all five methods. BCT consistently achieves the smallest log-loss. On the training data, the MAP tree produced by BCT corresponds to the full first-order Markov model with posterior probability of $\approx 98\%$, while on the entire data set the MAP tree has depth 2. The MAP model T_1^* now has $\pi(T_1^*|x) \approx 49.5\%$, while the first-order chain which appears as T_2^* has $\pi(T_2^*|x) \approx 48\%$. This ability of BCT to perform sequential updates and adaptive model averaging explains, in part, its superior performance.

The second best method, CTF, uses a first-order Markov model, as does MTD, which achieves near-identical performance. SMC uses a first-order model that groups $\{G, T\}$ into a single state, and VLMC similarly produces a tree of depth 1 that groups $\{A, G, T\}$ into a single state.

Pewee birdsong. The data, discussed in Section 4.2, consists of 1327 samples in a three-letter alphabet. When half of it was used for training and the other half for testing, VLMC and MTD both had an infinite log-loss after the 420th symbol in the test sequence; this was mentioned earlier as a possible drawback of methods that use maximum likelihood estimates for the parameters. Among the methods that give a finite log-loss, BCT was found to have the best performance. SMC produced a second-order model with five states, and by the end of the test data its log-loss was larger than that of the BCT by 13.2%. CTF used a fourth-order model for prediction, resulting in a log-loss larger than that of BCT by 1.2%.

To illustrate the relative performance of all five methods, in Figure 14, we report the prediction results on the last 10% of the samples, with the first 90% used for training. The BCT has smallest log-loss overall, approximately 4.8% lower than that of the second best method, CTF, at the end of the experiment. The CTF and MTD predictors both use a fourth-order model, SMC finds a second-order model with five states, and VLMC uses a simple tree model of depth 4.

6.3 | Discussion

The BCT predictor was found to have consistently better performance than the other four methods considered, achieving a log-loss that was between 1.2% and 4.8% better than that of the second best

method in each case. This is due, in part, to the fact that the BCT predictor is based on averaging over all models and parameters with respect to their posterior distribution, and that it allows for sequential updates beyond the training stage.

The method that performed the closest to BCT in most cases was CTF, which usually identified the same Markov order as the other methods. The VLMC and MTD predictors were found to be consistently and significantly less effective than BCT, and as noted earlier they sometimes assigned zero probability to the occurrence of certain symbols in the data, resulting in an infinite log-loss. Another difficulty with CTF, VLMC and MTD is that they required much more computational effort in the training stage.

The SMC predictor was found in most cases to have performance similar to VLMC. As the context tree models in $\mathcal{T}(D)$ are a subset of all SMCs, the fact that BCT performs better than SMC highlights the power of the full Bayesian BCT framework. In contrast, SMC uses a single (and often poor) estimate of its MAP model, leading to much less satisfactory results.

7 | CONCLUDING REMARKS

The Bayesian context trees framework introduced in this work was found to be effective for several core statistical tasks in the analysis of discrete time series. Variable-memory Markov chains are a rich class of models that capture important aspects of the higher-order temporal structure naturally present in many types of real-world data, and the associated exact inference algorithms in Sections 3.1–3.3 provide efficient tools for modelling, estimation and prediction. These algorithms, along with the variable-dimensional MCMC samplers of Section 3.5, also facilitate accurate posterior computations that offer a quantitative measure of uncertainty for the estimated models and parameters. The resulting methods were found to outperform several of the most commonly used approaches on both simulated and real-world data.

The exact nature of the algorithms in Sections 3.1–3.3, particularly the ability to compute the prior predictive likelihood, opens the door to numerous other applications that are the object of ongoing work, including anomaly detection, segmentation and change-point detection. Similarly, the low complexity of the algorithms opens the door to their effective use in problems involving ‘big data’. As has been noted in before connection with VLMCs, the main practical limitation in the application of these tools is when the alphabet size is large.

While our focus has been exclusively on discrete-valued time series, the ideas and techniques introduced in this work can be generalised to real- or vector-valued observations. Important directions of ongoing relevant research include the extension of the exact inference algorithms to various types of tree-structured autoregressive models, and the development of corresponding tools for inference with hidden Bayesian context trees.

ACKNOWLEDGEMENTS

We gratefully acknowledge interesting conversations with Petros Dellaportas, Marina Doufexis, Chris Gioran, Carl Rasmussen, Stephen Souras and Frans Willems. We are also grateful to Georgia Gregoriou for providing us with the spike train data examined in Section 4.2, and to the Associate Editor and anonymous reviewers for numerous helpful suggestions.

ORCID

Ioannis Kontoyiannis  <https://orcid.org/0000-0001-7242-6375>

REFERENCES

- Abakuks, A. (2012) The synoptic problem: on Matthew's and Luke's use of Mark. *Journal of the Royal Statistical Society: Series A*, 175(4), 959–975.
- Akaike, H. (1973) Information theory and an extension of the maximum likelihood principle. In: Petrov, B. & Csaki, F. (Eds.) *2nd international symposium on information theory*, Budapest, Hungary: Akademia Kiado, pp. 267–281.
- Atchade, Y. & Liu, J. (2004) The Wang-Landau algorithm for Monte Carlo computation in general state spaces. *Statistica Sinica*, 20, 209–233.
- Barron, A., Rissanen, J. & Yu, B. (1998) The minimum description length principle in coding and modeling. *IEEE Transactions on Information Theory*, 44(6), 2743–2760. Information theory: 1948–1998.
- Barron, A., Birgé, L. & Massart, P. (1999) Risk bounds for model selection via penalization. *Probability Theory and Related Fields*, 113(3), 301–413.
- Begleiter, R., El-yaniv, R. & Yona, G. (2004) On prediction using variable order Markov models. *Journal of Artificial Intelligence Research*, 22, 385–421.
- Bejerano, G. & Yona, G. (2001) Variations on probabilistic suffix trees: statistical modeling and prediction of protein families. *Bioinformatics*, 17(1), 23–43.
- Ben-Gal, I., Shani, A., Gohr, A., Grau, J., Arviv, S., Shmilovici, A. et al. (2005) Identification of transcription factor binding sites with variable-order Bayesian networks. *Bioinformatics*, 21(11), 2657–2666.
- Berchtold, A. (1996) Modélisation autorégressive des chaînes de Markov: utilisation d'une matrice différente pour chaque retard. *Revue de Statistique Appliquée*, 44(3), 5–25.
- Berchtold, A. (1998) *Chaînes de Markov et modèles de transition: applications aux sciences sociales*. Paris: Hermes.
- Berchtold, A. (2001) Estimation in the mixture transition distribution model. *Journal of Time Series Analysis*, 22(4), 379–397.
- Berchtold, A. & Raftery, A. (2002) The mixture transition distribution model for high-order Markov chains and non-Gaussian time series. *Statistical Science*, 17(3), 328–356.
- Bernardo, J. & Smith, A. (1994) *Bayesian theory*. New York: John Wiley & Sons.
- Bishop, C. (2006) *Pattern recognition and machine learning*. New York: Springer.
- Browning, S. (2006) Multilocus association mapping using variable-length Markov chains. *The American Journal of Human Genetics*, 78(6), 903–913.
- Bühlmann, P. (2000) Model selection for variable length Markov chains and tuning the context algorithm. *Annals of the Institute of Statistical Mathematics*, 52(2), 287–315.
- Bühlmann, P. & Wyner, A. (1999) Variable length Markov chains. *Annals of Statistics*, 27(2), 480–513.
- Busch, J., Ferrari, P., Flesia, A., Fraiman, R., Grynberg, S. & Leonardi, F. (2009) Testing statistical hypothesis on random trees and applications to the protein classification problem. *The Annals of Applied Statistics*, 3(2), 542–563.
- Cappé, O., Moulines, E. & Ryden, T. (2006) *Inference in hidden Markov models*. New York: Springer.
- Chipman, H., George, E., McCulloch, R., Clyde, M., Foster, D. & Stine, R. (2001) The practical implementation of Bayesian model selection. In: *Model selection*, Volume 38 of *IMS Lecture Notes Monogr. Ser.*, pp. 65–134. Inst. Math. Statist., Beachwood, OH. With discussion by M. Clyde, Dean P. Foster, and Robert A. Stine, and a rejoinder by the authors.
- Clark, K., Karsch-Mizrachi, I., Lipman, D., Ostell, J. & Sayers, E. (2016) GenBank. *Nucleic Acids Research*, 44(D1), D67–D72. Available from: <http://www.ncbi.nlm.nih.gov>
- Craig, W. (1943) *The song of the wood pewee (Myiochanes virens Linnaeus): a study of bird music*. New York State Museum Bulletin No. 334. Albany, NY: University of the State of New York.
- Csiszár, I. & Shields, P. (2000) The consistency of the BIC Markov order estimator. *Annals of Statistics*, 28(6), 1601–1619.
- Dayan, P. & Abbott, L. (2001) *Theoretical neuroscience: computational and mathematical modeling of neural systems*. Cambridge, MA: MIT Press.
- Dimitrakakis, C. (2010) Bayesian variable order Markov models. In *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics (AISTATS)*, Volume 9 of *JMLR: W&CP*, Chia Laguna Resort, Sardinia, Italy.
- Dziak, J., Coffman, D., Lanza, S., Li, R. & Jermiin, L. (2019) Sensitivity and specificity of information criteria. *Briefings in Bioinformatics*, 21(2), 553–565.

- Fahrmeir, L. & Kaufmann, H. (1987) Regression models for non-stationary categorical time series. *Journal of Time Series Analysis*, 8(2), 147–160.
- Fokianos, K. (2012) Count time series models. In: Rao, T.S., Rao, S.S. & Rao, C.R. (Eds.) *Handbook of statistics*, volume 30, Amsterdam: Elsevier, pp. 315–347.
- Fokianos, K. & Kedem, B. (2003) Regression theory for categorical time series. *Statistical Science*, 18(3), 357–376.
- Gabadinho, A. & Ritschard, G. (2016) Analyzing state sequences with probabilistic suffix trees: the PST R package. *Journal of Statistical Software*, 72(3), 1–39.
- Galves, A., Galves, C., Garcia, J., Garcia, N. & Leonardi, F. (2012) Context tree selection and linguistic rhythm retrieval from written texts. *The Annals of Applied Statistics*, 6(1), 186–209.
- Gao, Y., Kontoyiannis, I. & Bienenstock, E. (2006) From the entropy to the statistical structure of spike trains. In: *IEEE International Symposium on Information Theory (ISIT)*, Seattle, WA.
- Gao, Y., Kontoyiannis, I. & Bienenstock, E. (2008) Estimating the entropy of binary time series: methodology, some theory and a simulation study. *Entropy*, 10(2), 71–99.
- García, J. & González-López, V. (2011) Minimal Markov models. In: Rissanen, J., Myllymäki, P., Roos, T., Tabus, I. & Yamanishi, K. (Eds.) *Fourth Workshop on Information Theoretic Methods in Science and Engineering*. Helsinki, Finland: Department of Computer Science, University of Helsinki, p. 25.
- García, J. & González-López, V. (2017) Consistent estimation of partition Markov models. *Entropy*, 19(4), 160.
- Garivier, A. & Leonardi, F. (2011) Context tree selection: a unifying view. *Stochastic Processes and Their Applications*, 121(11), 2488–2506.
- Geisser, S. (1993) *Predictive inference: an introduction*. New York: CRC Press.
- Gelfand, A. & Smith, A. (1990) Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association*, 85(410), 398–409.
- Gelman, A., Carlin, J., Stern, H., Dunson, D., Vehtari, A. & Rubin, D. (2014) *Bayesian data analysis*. Boca Raton, FL: Chapman & Hall/CRC.
- Gregoriou, G., Gotts, S., Zhou, H. & Desimone, R. (2009) High-frequency, long-range coupling between prefrontal and visual cortex during attention. *Science*, 324(5931), 1207–1210.
- Gregoriou, G., Gotts, S. & Desimone, R. (2012) Cell-type-specific synchronization of neural activity in FEF with V4 during attention. *Neuron*, 73(3), 581–594.
- Grünwald, P. (2007) *The minimum description length principle*. Cambridge, MA: MIT Press.
- Hannan, E. & Quinn, B. (1979) The determination of the order of an autoregression. *Journal of the Royal Statistical Society: Series B*, 41(2), 190–195.
- Jääskinen, V., Xiong, J., Corander, J. & Koski, T. (2014) Sparse Markov chains for sequence data. *Scandinavian Journal of Statistics*, 41(3), 639–655.
- Jiao, J., Permuter, H., Zhao, L., Kim, Y.-H. & Weissman, T. (2013) Universal estimation of directed information. *IEEE Transactions on Information Theory*, 59(10), 6220–6242.
- Kass, R. & Raftery, A. (1995) Bayes factors. *Journal of the American Statistical Association*, 90(430), 773–795.
- Kharin, Y. & Piatlitski, A. (2011) Statistical analysis of discrete time series based on the MC(s,r)-model. *Austrian Journal of Statistics*, 40(1&2), 75–84.
- Lan, J., Ge, J., Yu, J., Shan, S., Zhou, H., Fan, S. et al. (2020) Structure of the SARS-CoV-2 spike receptor-binding domain bound to the ACE2 receptor. *Nature*, 581, 215–220.
- Mächler, M. & Bühlmann, P. (2004) Variable length Markov chains: methodology, computing, and software. *Journal of Computational and Graphical Statistics*, 13(2), 435–455.
- Papageorgiou, I., Lungu, V. & Kontoyiannis, I. (2020) BCT: Bayesian Context Trees for discrete time series. R package version 1.1. Available from: [CRAN.Rproject.org/package=BCT](https://CRAN.R-project.org/package=BCT).
- Raftery, A. (1985) A model for high-order Markov chains. *Journal of the Royal Statistical Society: Series B*, 47(3), 528–539.
- Raftery, A. & Tavaré, S. (1994) Estimation and modelling repeated patterns in high order Markov chains with the mixture transition distribution model. *Journal of the Royal Statistical Society: Series C*, 43(1), 179–199.
- Rieke, F., Warland, D., de Ruyter van Steveninck, R. & Bialek, W. (1999) *Spikes*. Cambridge, MA: MIT Press. Exploring the neural code, computational neuroscience.
- Rissanen, J. (1983a) A universal data compression system. *IEEE Transactions on Information Theory*, 29(5), 656–664.

- Rissanen, J. (1983b) A universal prior for integers and estimation by minimum description length. *Annals of Statistics*, 11(2), 416–431.
- Rissanen, J. (1984) Universal coding, information, prediction, and estimation. *IEEE Transactions on Information Theory*, 30(4), 629–636.
- Rissanen, J. (1986a) Complexity of strings in the class of Markov sources. *IEEE Transactions on Information Theory*, 32(4), 526–532.
- Rissanen, J. (1986b) Stochastic complexity and modeling. *Annals of Statistics*, 14(3), 1080–1100.
- Rissanen, J. (1987) Stochastic complexity. *Journal of the Royal Statistical Society: Series B*, 49(3), 223–239, 253–265. With discussion.
- Rissanen, J. (1989) *Stochastic complexity in statistical inquiry*. Singapore: World Scientific.
- Robert, C. & Casella, G. (2004) *Monte Carlo statistical methods*, 2nd edition, New York: Springer-Verlag.
- Ron, D., Singer, Y. & Tishby, N. (1996) The power of amnesia: learning probabilistic automata with variable memory length. *Machine Learning*, 25, 117–149.
- Sarkar, A. & Dunson, D. (2016) Bayesian nonparametric modeling of higher order Markov chains. *Journal of the American Statistical Association*, 111(516), 1791–1803.
- Saunders, A. (1944) The song of the wood pewee (book review). *The Auk*, 61(4), 658–660.
- Schwarz, G. (1978) Estimating the dimension of a model. *Annals of Statistics*, 6(2), 461–464.
- Shibata, R. (1980) Asymptotically efficient selection of the order of the model for estimating parameters of a linear process. *Annals of Statistics*, 8(1), 147–164.
- Takeuchi, J. & Barron, A. (2014) Stochastic complexity for tree models. In: *2014 IEEE Workshop on Information Theory (ITW)*, Hobart, Tasmania, Australia, pp. 222–226.
- Tjalkens, T., Willems, F. & Shtarkov, Y. (1994) Multi-alphabet universal coding using a binary decomposition context tree weighting algorithm. In: *15th Symposium on Information Theory in the Benelux*, Louvain-la-Neuve, Belgium.
- Wei, C.-Z. (1992) On predictive least squares principles. *Annals of Statistics*, 20(1), 1–42.
- Weinberger, M., Merhav, N. & Feder, M. (1994) Optimal sequential probability assignment for individual sequences. *IEEE Transactions on Information Theory*, 40(2), 384–396.
- Willems, F. & Volf, P. (1994) Context maximizing: Finding MDL decision trees. In: *15th Symposium on Information Theory in the Benelux*, Louvain-la-Neuve, Belgium.
- Willems, F. & Volf, P. (1995) A study of the context tree maximizing method. In: *16th Symposium on Information Theory in the Benelux*, Nieuwerkerk Ijsel, The Netherlands.
- Willems, F., Shtarkov, Y. & Tjalkens, T. (1993a) Context tree weighting: Basic properties. Unpublished manuscript. Available from: <http://www.sps.tue.nl/wp-content/uploads/2015/09/ctw1submission.pdf>
- Willems, F., Shtarkov, Y. & Tjalkens, T. (1993b) Context tree weighting: Redundancy bounds and optimality. In: *6th Joint Swedish-Russian International Workshop on Information Theory*, Mölle, Sweden.
- Willems, F., Shtarkov, Y. and Tjalkens, T. (1993c) Context weighting: General finite context sources. In: *14th Symposium on Information Theory in the Benelux*, Veldhoven, The Netherlands.
- Willems, F., Shtarkov, Y. & Tjalkens, T. (1995) The context tree weighting method: basic properties. *IEEE Transactions on Information Theory*, 41(3), 653–664.
- Willems, F., Shtarkov, Y. & Tjalkens, T. (2000) Context tree maximizing. In: *2000 Conference on Information Sciences and Systems*, Princeton, NJ.
- Willems, F., Nowbahkt-irani, A. & Volf, P. (2002) Maximum a-posteriori probability tree models. In: *4th International ITG Conference on Source and Channel Coding*, Berlin, Germany.
- Wu, F., Zhao, S., Yu, B., Chen, Y.M., Wang, W., Song, Z.G. et al. (2020) A new coronavirus associated with human respiratory disease in China. *Nature*, 579(7798), 265–269.
- Xiong, J., Jääskinen, V. & Corander, J. (2016) Recursive learning for sparse Markov models. *Bayesian Analysis*, 11(1), 247–263.
- Yan, R., Zhang, Y., Li, Y., Xia, L., Guo, Y. & Zhou, Q. (2020) Structural basis for the recognition of SARS-CoV-2 by full-length human ACE2. *Science*, 367(6485), 1444–1448.
- Yang, Y. & Dunson, D. (2016) Bayesian conditional tensor factorizations for high-dimensional classification. *Journal of the American Statistical Association*, 111(514), 656–669.
- Yee, T. (2010) The VGAM package for categorical data analysis. *Journal of Statistical Software*, 32(10), 1–34.

Zeger, S. & Liang, K. (1986) Longitudinal data analysis for discrete and continuous outcomes. *Biometrics*, 42(1), 121–130.

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of the article at the publisher's website.

How to cite this article: Kontoyiannis, I., Mertzanis, L., Panotopoulou, A., Papageorgiou, I. & Skoularidou, M. (2022) Bayesian context trees: Modelling and exact inference for discrete time series. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 1–37. Available from: <https://doi.org/10.1111/rssb.12511>