

Detecting Renewal States in Chains of Variable Length via Intrinsic Bayes Factors

Victor Freguglia and Nancy L. Garcia
Department of Statistics – University of Campinas

Abstract

Markov chains with variable length are useful parsimonious stochastic models able to generate most stationary sequence of discrete symbols. The idea is to identify the suffixes of the past, called contexts, that are relevant to predict the future symbol. Sometimes a single state is a context, and looking at the past and finding this specific state makes the further past irrelevant. States with such property are called renewal states and they can be used to split the chain into independent and identically distributed blocks. In order to identify renewal states for chains with variable length, we propose the use of Intrinsic Bayes Factor to evaluate the hypothesis that some particular state is a renewal state. In this case, the difficulty lies in integrating the marginal posterior distribution for the random context trees for general prior distribution on the space of context trees, with Dirichlet prior for the transition probabilities, and Monte Carlo methods are applied. To show the strength of our method, we analyzed artificial datasets generated from different binary models and one example coming from the field of Linguistics.

Keywords: Variable length Markov Chains, Renewal States, Bayes Factor, intractable normalizing constant

1 Introduction

Markov Chains with variable length are useful stochastic models that provide a powerful framework for describing transition probabilities for finite-valued sequences due the possibility of capturing long-range interactions while keeping some parsimony in the number of free parameters. These models were introduced in the seminal paper of Rissanen (1983) for data compression and became known in the statistics literature as Variable Length Markov Chain (VLMC) by Bühlmann and Wyner (1999), and as Probabilistic Suffix Trees (PST) in the machine learning literature (Ron et al., 1996). The idea is that, for each past, only a finite suffix of the past is enough to predict the next symbol. Rissanen called *context*, the relevant ending string of the past. The set of all contexts can be represented by the set of leaves of a rooted tree if we require that no context is a proper suffix of another context. For a fixed set of contexts, estimation of

the transition probabilities can be easily achieved. The problems lies into estimating the set of contexts from the available data. In his seminal 1983 paper, Rissanen introduced the Context algorithm, which estimates the context tree by aggregating irrelevant states in the history of the process using a sequential procedure. A nice introductory guide to this type of models and particularly to the Context algorithm can be found in [Galves and Löcherbach \(2008\)](#).

Many of the tree model methods related to data compression tasks involve obtaining better predictions based on weighting over multiple models. A classical example is the Context-Tree Weighting (CTW) algorithm ([Willems et al., 1995](#)), which computes the marginal probability of a sequence by weighting over all context trees and all probability vectors using computationally convenient weights. Using CTW, [Csiszár and Talata \(2006\)](#) showed that context trees can be consistently estimated in linear time using the Bayesian information criterion (BIC). These weighting strategies can be translated to a Bayesian framework where unobserved parameters of a probabilistic system are treated as additional random components with given prior distribution and inference is based on integrating over the nuisance parameters, which is a form of weighting over these quantities based on the prior distribution. Nonetheless, inference performed following the Bayesian paradigm for VLMC models is a relatively recent topic of research. Some works that explicitly use Bayesian statistics in combination with VLMC models are [Dimitrakakis \(2010\)](#) which introduced an online prediction scheme by adding a prior, conditioned on context, on the Markov order of the chain, and [Kontoyiannis et al. \(2020\)](#) which provided more general tools such as posterior sampling through Metropolis-Hastings algorithm and Maximum a Posteriori context tree estimation focusing on model selection, estimation, and sequential prediction. A Bayesian approach for model selection in high-order Markov chains, allowing conditional probabilities to be partitioned into more general structures than the tree-based structures of VLMC models, is also proposed in [Xiong et al. \(2016\)](#).

As aforementioned, the effort was mostly concentrated in estimating the context tree structure. On the other hand, hypothesis testing for VLMC is a difficult topic, first tackled by [Balding et al. \(2009\)](#) and pursued further by [Busch et al. \(2009\)](#) using a Kolmogorov-Smirnov-type goodness-of-fit test, to compare if two samples come from the same distribution. Under the Bayesian paradigm, hypothesis testing is done through Bayes Factor, but computing Bayes Factors may depend on integrations that require enormous computational effort depending on the random objects and hypotheses involved. Particularly for VLMC problems, Bayes Factors require summing over the set of all possible context trees, which cardinality grows doubly exponentially with the maximum depth considered, quickly becoming intractable. To avoid such intractable quantities, we use the Monte Carlo approximations of the Intrinsic Bayes Factor ([Berger and Pericchi, 1996](#)), which is based on averaging over posterior distributions, that tend to be highly concentrated within a small set of context trees for VLMC models, and have been used recently in many different fields with the same purpose, such as [Cabras et al. \(2015\)](#), [Charitidou et al. \(2018\)](#) and [Villa and Walker \(2021\)](#). Alternatives to Bayes Factors based on using posterior distributions

instead of the prior distribution in integrations have been evolving over the past decades, the classical method using this strategy is the Posterior Bayes Factor [Aitkin \(1991\)](#), with applications in a variety of models such as [Aitkin \(1993\)](#) and [Aitkin et al. \(1996\)](#).

In this work, we focus on one characteristic of interest in a context tree, the presence of a renewal state or a renewal context. Renewal states play an important role in some computational methods frequently used in statistical analysis such as designing Bootstrap schemes and defining proper cross-validation strategies based on blocks. Therefore, having some methodology not only to detect renewal states, but also to quantify how plausible these assumptions are, can improve the robustness of analysis at the cost of some pre-processing computations. For example, [Galves et al. \(2012\)](#) proposed a constant free algorithm (Smallest Maximizer Criterion) to find the tree that maximizes the BIC based on a Bootstrap scheme that uses the renewal property of one of the states. To the best of our knowledge, using Bayes Factors for evaluating hypotheses involving probabilistic context trees is a topic that has not been explored.

2 Variable-Length Markov Chains

2.1 Model Description

Let \mathcal{A} be an *alphabet* of m symbols, and without loss of generality, consider $\mathcal{A} = \{0, 1, \dots, (m-1)\}$ for simplicity. For $t_2 > t_1$, a string $(z_{t_1}, \dots, z_{t_2}) \in \mathcal{A}^{t_2 - t_1 + 1}$ will be denoted by $\mathbf{z}_{t_1}^{t_2}$ and its length by $\ell(\mathbf{z}_{t_1}^{t_2}) = t_2 - t_1 + 1$. A sequence \mathbf{s}_{n-l}^n is a *suffix* of a string \mathbf{z}_1^n if $s_j = z_j$ for all $j = n-l, \dots, n$. If $l < n$ we say that \mathbf{s}_{n-l}^n is a *proper suffix* of the string \mathbf{z}_1^n .

Definition 1 Let $L > 0$ and $\tau \subset \cup_{j=1}^L \mathcal{A}^j$ be a set of strings formed by symbols in \mathcal{A} . We say that τ satisfies the *suffix property* if, for every string $\mathbf{s}_{-j+j'}^{-1} = (s_{-j+j'}, \dots, s_{-1}) \in \mathcal{A}^{j-j'}$, $\mathbf{s}_{-j+j'}^{-1} \in \tau$ implies that $\mathbf{s}_{-j}^{-1} \notin \tau$ for $j > 1, j' = 1, \dots, j$.

Definition 2 Let $L > 0$ and $\tau \subset \cup_{j=1}^L \mathcal{A}^j$ be a set of strings formed by symbols in \mathcal{A} . We say that τ is an *irreducible tree* if, no string belonging to τ can be replaced by a proper suffix without violating the suffix property.

Definition 3 Let τ be an irreducible tree. We say that τ is *full* if, for each string $\mathbf{s} \in \tau$, any concatenation of a symbol $k \in \mathcal{A}$ and a suffix of \mathbf{s} is the suffix of a string $\mathbf{s}' \in \tau$.

Examples Suppose that we have a binary alphabet $\mathcal{A} = \{0, 1\}$, then:

- $\tau = \{0, 1, 11\}$ does **not satisfy the suffix property** because it contains both the strings 1 and 11.

- $\tau = \{0, 01\}$ is **not an irreducible tree**, because the string 01 can be replaced by its suffix, 1, without violating the suffix property, as the set $\{0, 1\}$ satisfies the suffix property.
- $\tau = \{0, 011, 111\}$ is **irreducible**, but it is **not full** because 1 is a suffix of a string in τ (either 011 or 111), but 01 (the concatenation of $0 \in \mathcal{A}$ and 1) is not.
- $\tau_1 = \{0, 100, 101, 110, 111\}$, $\tau_2 = \{01, 00, 10, 11\}$ and $\tau_3 = \{0, 10, 110, 1110, 1111\}$ are **full irreducible trees**.

A full irreducible tree τ can be represented by the set of leaves of a rooted tree with a finite set of labeled branches such that

- (1) The root node has no label,
- (2) each node has either 0 or m children (fullness) and
- (3) when a node has m children, each child has symbol of the alphabet \mathcal{A} as a label.

The elements of τ will be called *contexts* and we will refer to full irreducible trees as *context trees* henceforth. **Figure 1** presents 3 examples of contexts trees. The *depth* ℓ of a tree τ is given by the maximal length of a context belonging to τ , defined as

$$\ell(\tau) = \max\{\ell(\mathbf{z}); \mathbf{z} \in \tau\}.$$

In this work we will assume that the depth of the tree is bounded by an integer L . In this case, it is straightforward to conclude that, for any string $\mathbf{z}_{t_1}^{t_2}$ with at least $L + 1$ symbols, there exist a suffix $\mathbf{z}_{t_2-l}^{t_2}$ and a leaf of τ such that the symbols between the leaf (including) and the root node are exactly $\mathbf{z}_{t_2-l}^{t_2}$. Galves et al. (2012) referred to this property as the *properness* of a context tree.

For each context tree τ , we can associate a family of probability measures indexed by elements of τ ,

$$\mathbf{p} = \{p(\cdot|\mathbf{s}) : \mathcal{A} \rightarrow [0, 1]; \mathbf{s} \in \tau\}.$$

The pair (τ, \mathbf{p}) is called a *probabilistic context tree*.

Given a tree τ with the described properties and depth bounded by L , define a *suffix mapping function* $\eta_\tau : \cup_{j=L+1}^\infty \mathcal{A}^j \rightarrow \tau$ such that $\eta_\tau(\mathbf{z}_{t_1}^{t_2}) = \mathbf{z}_{t_2-l}^{t_2}$ is the unique suffix $\mathbf{z}_{t_2-l}^{t_2} \in \tau$.

Definition 4 A sequence of random variables $\mathbf{Z} = (Z_t)_{t=1}^T$ with state space \mathcal{A} is a *Variable Length Markov Chain (VLMC)* compatible with the probabilistic tree (τ, \mathbf{p}) if it satisfies

$$\mathbb{P}(Z_t = k | \mathbf{Z}_1^{t-1} = \mathbf{z}_1^{t-1}) = \mathbf{p}(k | \eta_\tau(\mathbf{z}_1^{t-1})), \quad (1)$$

for all $L < t \leq T$, $\mathbf{z}_1^{t-1} \in \mathcal{A}^{t-1}$, where $\eta_\tau(\mathbf{z}_1^{t-1}) \in \tau$ is the suffix of \mathbf{z}_1^{t-1} .

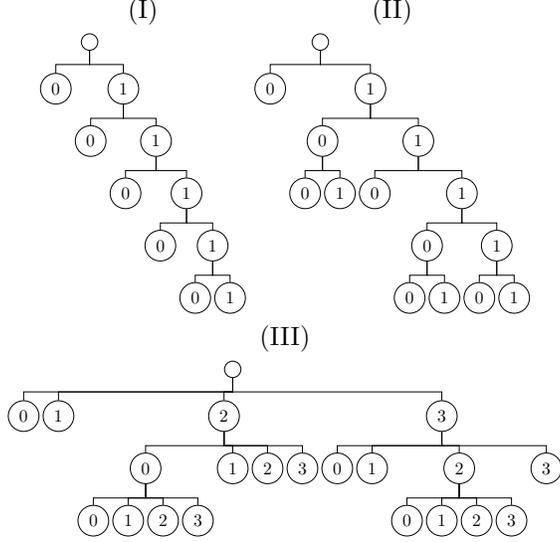


Figure 1: Examples of context trees.

2.2 Likelihood Function

In order to extend the scope of VLMC models introduced previously to data involving multiple sequences, we define a VLMC dataset of size I , denoted as $\tilde{\mathbf{Z}}$, as a set of independent VLMC sequences $\tilde{\mathbf{Z}} = (\mathbf{Z}^{(i)})_{i=1, \dots, I}$ and $\tilde{\mathbf{z}} = (\mathbf{z}^{(i)})_{i=1, \dots, I}$ will denote its observed realizations.

For each sequence $\mathbf{Z}^{(i)}$, with length T_i , we will consider its first L elements as constant values, allowing us to write the joint probabilities as a product of transition probabilities in (1), without requiring additional parameters for consistently defining the probabilities of the first symbols in each sequence. Hence, the likelihood function is given by

$$f(\tilde{\mathbf{z}}|\tau, \mathbf{p}) = \prod_{\mathbf{s} \in \tau} \prod_{k=0}^{m-1} (p(k|\mathbf{s}))^{n_{\mathbf{s}k}(\tilde{\mathbf{z}})}, \quad (2)$$

where $n_{\mathbf{s}k}(\tilde{\mathbf{z}}) = \sum_{i=1}^I \sum_{t=L+1}^{T_i} \mathbb{1}(z_t^{(i)} = k, \eta_\tau(\mathbf{z}_1^{t-1(i)}) = \mathbf{s})$ counts the number of occurrences of the symbol k after strings with suffix \mathbf{s} across all sequences.

2.3 Renewal States

A symbol $a \in \mathcal{A}$ is called a *renewal state* if

$$\mathbb{P}(\mathbf{Z}_{t+1}^{t'} = \mathbf{z}_{t+1}^{t'} | \mathbf{Z}_1^{t-1} = \mathbf{z}_1^{t-1}, Z_t = a) = \mathbb{P}(\mathbf{Z}_{t+1}^{t'} = \mathbf{z}_{t+1}^{t'} | Z_t = a),$$

for all $t > L$, $t' > t + 1$, $\mathbf{z}_1^{t-1} \in \mathcal{A}^{t-1}$, $\mathbf{z}_{t+1}^{t'} \in \mathcal{A}^{t'-t}$. That is, conditioning on $Z_t = a$, the distribution of the chain after t , $(Z_u)_{u>t}$, is independent from the

past $(Z_u)_{u < t}$.

This property of conditional independences in a Markov Chain can be directly associated to the structure of the context tree of a VLMC model. For a VLMC model with associated context tree τ , a state $a \in \mathcal{A}$ is a renewal state if a does not appear in any inner node of the context tree. That is, for any context $\mathbf{s} \in \tau$, expressing \mathbf{s} as the concatenation of l symbols $\mathbf{s} = s_l \dots s_2 s_1$, $s_i \neq a$ for $i = 1, \dots, l - 1$. In this case, we say that the tree τ is *a-renewing*.

Two out of the three trees displayed in [Figure 1](#) present renewal states. Tree (I) has 0 as a renewal state, Tree (II) has no renewal states due to the presence of the contexts 001 and 00111, Tree (III) has only 1 as a renewal state. If, in Tree (III), the branch formed by contexts 002, 102, 202 and 302 was pruned and substituted by the context 02 only, then 0 would also be a renewal state. Note that a VLMC may contain multiple renewal states.

A remarkable consequence of a being a renewal state is that the random blocks between two occurrences of a are independent and identically distributed. This feature allows the use of block Bootstrap methods, enables a straightforward construction of cross-validation schemes and any other technique that relies on exchangeability properties.

3 Bayesian Renewal Hypothesis Evaluation

A VLMC model is fully specified by the probabilistic context tree (τ, \mathbf{p}) . The dimension of \mathbf{p} depends on the branches of τ . Both these unobserved components (τ, \mathbf{p}) can be treated as random elements with given prior distribution to carry out inference under the Bayesian paradigm.

From now on, we will use the following notation, for each $\mathbf{s} \in \tau$,

$$\mathbf{p}_{\mathbf{s}} = (p(0|\mathbf{s}), \dots, p(m-1|\mathbf{s})) \in \Delta_m$$

where Δ_m denotes the m -simplex,

$$\Delta_m = \left\{ \mathbf{x} \in \mathbb{R}^{\{0,1,\dots,m-1\}} : \sum_{k=0}^{m-1} x_k = 1 \text{ and } \forall j, x_j \geq 0 \right\}.$$

In this section we discuss the prior specification for the probabilistic context tree (τ, \mathbf{p}) , as well as the resultant posterior distribution and how to perform hypothesis testing using partial and intrinsic Bayes factor.

3.1 A Bayesian Framework for VLMC models

We consider a general prior distribution for τ proportional to any arbitrary non-zero function $h : \mathcal{T}_L \rightarrow [0, \infty)$ and, given τ , for each $\mathbf{s} \in \tau$, $\mathbf{p}_{\mathbf{s}}$ will have independent Dirichlet priors.

The complete Bayesian system can be described by the hierarchical structure

$$\begin{aligned}\tau &\sim \frac{h(\tau)}{\zeta(h, L)}, & \tau &\in \mathcal{T}_L, \\ \mathbf{p}|\tau &\sim \prod_{\mathbf{s} \in \tau} \frac{\Gamma(\sum_{k=0}^{m-1} \alpha_{\mathbf{s}k})}{\prod_{k=0}^{m-1} \Gamma(\alpha_{\mathbf{s}k})} \prod_{k=0}^{m-1} (p(k|\mathbf{s}))^{\alpha_{\mathbf{s}k}-1}, & \mathbf{p} &\in \Delta_m^{|\tau|}, \\ \tilde{\mathbf{Z}}|\tau, \mathbf{p} &\sim f(\tilde{\mathbf{z}}|\tau, \mathbf{p}), & \mathbf{z}^{(i)} &\in \mathcal{A}^{T_i},\end{aligned}\quad (3)$$

where

$$\zeta(h, L) = \sum_{\tau \in \mathcal{T}_L} h(\tau) \quad (4)$$

is the normalizing constant of the tree prior distribution and $f(\tilde{\mathbf{z}}|\tau, \mathbf{p})$ is given by (2). We are assuming that the prior distribution for the transition probabilities $\mathbf{p}_{\mathbf{s}}$ are independent Dirichlet distribution with hyperparameters $\boldsymbol{\alpha}_{\mathbf{s}} = (\alpha_{\mathbf{s}1}, \dots, \alpha_{\mathbf{s}(m-1)})$. Therefore, the joint distribution of $(\tau, \mathbf{p}, \tilde{\mathbf{Z}})$ is given by

$$\pi(\tau, \mathbf{p}, \tilde{\mathbf{z}}) = \frac{h(\tau)}{\zeta(h, L)} \prod_{\mathbf{s} \in \tau} \frac{\Gamma(\sum_{k=0}^{m-1} \alpha_{\mathbf{s}k})}{\prod_{k=0}^{m-1} \Gamma(\alpha_{\mathbf{s}k})} \prod_{k=0}^{m-1} (p(k|\mathbf{s}))^{n_{\mathbf{s}k}(\tilde{\mathbf{z}}) + \alpha_{\mathbf{s}k} - 1}.$$

Since our interest lies in making inferences about the dependence structure represented by τ rather than the transition probabilities, we can simplify our analysis by marginalizing the joint probability function over \mathbf{p} , $\pi(\tau, \tilde{\mathbf{z}}) = \int \pi(\tau, \mathbf{p}, \tilde{\mathbf{z}}) d\mathbf{p}$, obtaining a function that depends only on the context tree and the data. The product of Dirichlet densities, assigned as the prior distribution of \mathbf{p} , conjugates to the likelihood function, allowing us to express the integrated distribution in closed-form as

$$\pi(\tau, \tilde{\mathbf{z}}) = \frac{h(\tau)}{\zeta(h, L)} \prod_{\mathbf{s} \in \tau} \frac{\Gamma(\sum_{k=0}^{m-1} \alpha_{\mathbf{s}k})}{\prod_{k=0}^{m-1} \Gamma(\alpha_{\mathbf{s}k})} \frac{\prod_{k=0}^{m-1} \Gamma(n_{\mathbf{s}k}(\tilde{\mathbf{z}}) + \alpha_{\mathbf{s}k})}{\Gamma(\sum_{k=0}^{m-1} n_{\mathbf{s}k}(\tilde{\mathbf{z}}) + \alpha_{\mathbf{s}k})}, \quad (5)$$

obtained by multiplying the appropriate normalizing constant to achieve Dirichlet densities with parameters $(\alpha_{\mathbf{s}0} + n_{\mathbf{s},0}(\tilde{\mathbf{z}}), \dots, \alpha_{\mathbf{s}(m-1)} + n_{\mathbf{s},m-1}(\tilde{\mathbf{z}}))$ for each $\mathbf{s} \in \tau$, so that the integration is done on a proper density. For a less convoluted notation, we shall denote

$$q(\tau, \tilde{\mathbf{z}}) = \prod_{\mathbf{s} \in \tau} \frac{\Gamma(\sum_{k=0}^{m-1} \alpha_{\mathbf{s}k})}{\prod_{k=0}^{m-1} \Gamma(\alpha_{\mathbf{s}k})} \frac{\prod_{k=0}^{m-1} \Gamma(n_{\mathbf{s}k}(\tilde{\mathbf{z}}) + \alpha_{\mathbf{s}k})}{\Gamma(\sum_{k=0}^{m-1} n_{\mathbf{s}k}(\tilde{\mathbf{z}}) + \alpha_{\mathbf{s}k})}, \quad (6)$$

and use, from now on, the shorter expression $\pi(\tau, \tilde{\mathbf{z}}) = \frac{h(\tau)}{\zeta(h, L)} q(\tau, \tilde{\mathbf{z}})$.

Finally, the model evidence (marginal likelihood) can now be obtained by summing (5) over all trees in \mathcal{T}_L ,

$$\mathcal{E}(\tilde{\mathbf{z}}; h) = \sum_{\tau \in \mathcal{T}_L} \pi(\tau, \tilde{\mathbf{z}}) = \sum_{\tau \in \mathcal{T}_L} \frac{h(\tau)}{\zeta(h, L)} q(\tau, \tilde{\mathbf{z}}). \quad (7)$$

Note that we explicitly describe the model evidence in terms of the prior distribution h as we will be interested in evaluating hypotheses based on different prior distributions.

3.2 Bayes Factors for Renewal State Hypothesis

Let $\tilde{\mathbf{z}} = (\mathbf{z}^{(i)})_{i=1,\dots,I}$ be a VLMC sample compatible with a probabilistic context tree (τ, \mathbf{p}) where τ has maximum depth L . We will call *maximal tree* the complete tree with depth L and let $a \in \mathcal{A}$ be a fixed state of the alphabet. Our goal is to use Bayes Factors (Kass and Raftery, 1995) to evaluate the evidence in favor of the null hypothesis H_a that τ is *a-renewing* against an alternative hypothesis $H_{\bar{a}}$ that τ is not *a-renewing*. We denote $\mathcal{T}_L^a \subset \mathcal{T}_L$ the set of *a-renewing* trees with depth no more than L and $\bar{\mathcal{T}}_L^a$ the set of trees with a as an inner node and, consequently, a is not a renewal state for those trees.

We are interested in defining a metric for evaluating the hypothesis $H_a : \tau \in \mathcal{T}_L^a$ against its complement $H_{\bar{a}} : \tau \in \bar{\mathcal{T}}_L^a$ in a Bayesian framework. These hypotheses can be expressed in terms of special prior distributions proportional to functions h_a and $h_{\bar{a}}$, respectively, such that $h_a(\tau) = 0$ if, and only if, $\tau \in \bar{\mathcal{T}}_L^a$. Similarly, $h_{\bar{a}}(\tau) = 0$ if, and only if, $\tau \in \mathcal{T}_L^a$.

The Bayes Factor for H_a against $H_{\bar{a}}$ is defined as

$$\text{BF}_{a,\bar{a}}(\tilde{\mathbf{z}}) = \frac{\mathcal{E}(\tilde{\mathbf{z}}; h_a)}{\mathcal{E}(\tilde{\mathbf{z}}; h_{\bar{a}})} = \frac{\zeta(h_{\bar{a}}, L) \sum_{\tau \in \mathcal{T}_L} h_a(\tau) q(\tau, \tilde{\mathbf{z}})}{\zeta(h_a, L) \sum_{\tau \in \mathcal{T}_L} h_{\bar{a}}(\tau) q(\tau, \tilde{\mathbf{z}})}, \quad (8)$$

where $\zeta(\cdot, L)$ is given by (4) and q is given by (6).

Kass and Raftery (1995) proposed the following interpretation for the quantity $\log_{10}(\text{BF}_{a,\bar{a}}(\tilde{\mathbf{z}}))$ as a measure of evidence provided by the data $\tilde{\mathbf{z}}$ in favor of the hypothesis that corresponds to *a-renewing* trees as opposed to the alternative one. A value between 0 and 1/2 is considered to provide evidence that is “Not worth more than a bare mention”, “Substantial” for values between 1/2 and 1, “Strong” if they are between 1 and 2 and “Decisive” for values greater than 2. By symmetry, these intervals with negative sign provide the same amount of evidence but reversing the hypotheses considered. Therefore, the sign of $\log_{10}(\text{BF}_{a,\bar{a}}(\tilde{\mathbf{z}}))$ provides a straightforward measure whether the data provides more evidence that the chain is compatible with a context tree τ that is *a-renewing* or that τ belongs to $\bar{\mathcal{T}}_L^a$.

3.3 Metropolis-Hastings algorithm for context tree posterior sampling

Before further development of methods to compute the Bayes Factors from (8), we need to introduce a Metropolis-Hastings algorithm for sampling from the marginal posterior distribution of context trees, $\pi(\tau|\tilde{\mathbf{z}})$. From (5) and the Bayes rule we obtain

$$\pi(\tau|\tilde{\mathbf{z}}) = \frac{\frac{h(\tau)}{\zeta(h,L)} q(\tau, \tilde{\mathbf{z}})}{\mathcal{E}(\tilde{\mathbf{z}}; h)} \propto h(\tau) q(\tau, \tilde{\mathbf{z}}), \quad (9)$$

which has a simple expression up to the intractable proportionality terms, suggesting that the Metropolis-Hastings algorithm (Hastings, 1970; Chib and

(Greenberg, 1995) is an appropriate strategy to obtain an empirical sample from the posterior distribution given by (9).

The main step for constructing the algorithm is defining a suitable proposal kernel $\kappa(\tau'|\tau), \tau', \tau \in \mathcal{T}_L$ to move to new context trees from a current tree τ . We propose the use of a graph-based kernel that can be viewed as a modification of the Monte Carlo Markov Chain Model Composition (MC³) method from Madigan et al. (1995) by defining a neighborhood system \mathcal{N} over \mathcal{T}_L and constructing a proposal kernel that allows transitions only between neighboring trees only.

We first specify a set **directed** edges \mathcal{N}_d such that, an edge (τ, τ') , from τ to τ' , is included if, and only if, $|\tau \nabla \tau'| = m + 1$ and $|\tau'| > |\tau|$, where ∇ denotes the symmetric difference operator $A \nabla B = (A \cap B^c) \cup (A^c \cap B)$. An equivalent definition is that an edge from τ to τ' is obtained substituting one of the contexts $\mathbf{s} \in \tau$, by the contexts associated with its m children nodes, $\{ks, k \in \mathcal{A}\}$, in τ' . We refer to this substitution as *growing a branch* from \mathbf{s} .

Additionally, we define the *grow* (\oplus) and *prune* (\ominus) operators, as

$$\begin{aligned}\oplus(\tau, h) &= \{\tau' \in \mathcal{T}_L : (\tau, \tau') \in \mathcal{N}_d \text{ and } h(\tau') > 0\}, \\ \ominus(\tau, h) &= \{\tau' \in \mathcal{T}_L : (\tau', \tau) \in \mathcal{N}_d \text{ and } h(\tau') > 0\}.\end{aligned}$$

The operator \oplus maps a tree τ to the set of trees with positive prior distribution that can be obtained by growing new branches from τ , whereas \ominus maps τ to the set of trees in \mathcal{T}_L from which τ can be obtained after growing a branch.

Some important properties that can be easily checked are

1. For every $\tau \in \mathcal{T}_L$, if $\tau' \in \oplus(\tau, h)$ and $h(\tau) > 0$, then $\tau \in \ominus(\tau', h)$.
2. For every $\tau \in \mathcal{T}_L$, if $\tau' \in \ominus(\tau, h)$ and $h(\tau) > 0$, then $\tau \in \oplus(\tau', h)$.
3. For any finite sequence $\tau^{(1)}, \tau^{(2)}, \dots, \tau^{(N)}$ such that $h(\tau^{(1)}) > 0$ and $\tau^{(n+1)} \in \oplus(\tau^{(n)}, h) \cup \ominus(\tau^{(n)}, h)$, we have $\tau^{(n)} \in \oplus(\tau^{(n+1)}, h) \cup \ominus(\tau^{(n+1)}, h)$.

It follows from Properties 1 and 2 that any context tree τ can be recovered by applying sequentially grow and prune operations. Property 3 is a direct consequence of Properties 1 and 2 and means that any sequence of context trees obtained by a sequence of grow or prune operations can also be visited in reverse order with a sequence of grow and prune operations. These properties also suggest that combining \oplus and \ominus for constructing a set of transitions with positive probabilities in a proposal kernel is a good strategy in order to achieve the irreducibility condition $\kappa(\tau|\tau') > 0$ if, and only if, $\kappa(\tau'|\tau) > 0$.

We define a transition kernel κ as

$$\kappa(\tau'|\tau) = \begin{cases} \frac{1}{|\oplus(\tau, h)|} \mathbf{1}(\tau' \in \oplus(\tau, h)), & \text{if } \ominus(\tau, h) = \emptyset, \\ \frac{1}{|\ominus(\tau, h)|} \mathbf{1}(\tau' \in \ominus(\tau, h)), & \text{if } \oplus(\tau, h) = \emptyset, \\ \frac{1}{2} \frac{1}{|\oplus(\tau, h)|} \mathbf{1}(\tau' \in \oplus(\tau, h)) + \frac{1}{2} \frac{1}{|\ominus(\tau, h)|} \mathbf{1}(\tau' \in \ominus(\tau, h)), & \text{o/w,} \end{cases}$$

Algorithm 1: Metropolis-Hastings algorithm for sampling context trees from $\pi(\tau|\tilde{\mathbf{z}})$ under a tree prior distribution proportional to h .

```

1 Set an initial tree  $\tau^{(0)} \in \mathcal{T}_L$ ;
2 for  $t = 1, \dots, n_{\text{iter}}$  do
3   Sample an operation  $\oplus$  or  $\ominus$  which leads to a non-empty set when
   applied to  $\tau^{(t-1)}$ , with equal probabilities;
4   Randomly pick a proposed tree  $\tau'$  from  $\oplus(\tau^{(t-1)}, h)$  or  $\ominus(\tau^{(t-1)}, h)$ ;
5   Compute the acceptance ratio
      
$$A(\tau'|\tau^{(t-1)}) = \min\left(\frac{h(\tau')q(\tau', \tilde{\mathbf{z}})}{h(\tau^{(t-1)})q(\tau^{(t-1)}, \tilde{\mathbf{z}})} \frac{\kappa(\tau^{(t-1)}|\tau')}{\kappa(\tau'|\tau^{(t-1)})}, 1\right);$$

      Generate a random variable  $U \sim \text{Unif}(0, 1)$ ;
6   if  $U < A(\tau'|\tau^{(t-1)})$  then
7     |  $\tau^{(t)} \leftarrow \tau'$ 
8   else
9     |  $\tau^{(t)} \leftarrow \tau^{(t-1)}$ 
10  end if
11 end for

```

which allows us to propose a tree τ' in a simple two-step process. First, pick the operator to be applied to τ , \oplus or \ominus with probabilities 1/2 if both lead to non-empty set of trees, otherwise pick the operation that produces a non-empty set. Then, pick τ' from $\oplus(\tau, h)$ or $\ominus(\tau, h)$ with uniform probabilities.

The idea of a proposal kernel for context trees based on growing and pruning nodes of trees was already used in [Kontoyiannis et al. \(2020\)](#) for a specific prior distribution h . The complete Metropolis-Hastings algorithm described in [Algorithm 1](#), not only formally defines the construction in terms of graphs, but also extends it to accommodate arbitrary prior distributions.

3.4 Partial Bayes Factors and Intrinsic Bayes Factors

While $\text{BF}_{a,\bar{a}}(\tilde{\mathbf{z}})$ given by (8) provides a measure of the plausibility of one hypothesis with respect to another, computing this quantity may present an enormous computational cost as the sum over \mathcal{T}_L involves a doubly exponential number of terms. In fact, even the normalizing constant of the tree prior distribution $\zeta(h, L)$ is intractable in the general case for moderate values L , hindering the evaluation of the model evidence $\mathcal{E}(\tilde{\mathbf{z}}; h)$.

Previous algorithms proposed in the literature that use similar ideas to compute the marginal likelihood, like the Context Tree Weighting (CTW) algorithm from [Willems et al. \(1995\)](#), are not suitable for our purposes. While we are aiming to compute (7) for an arbitrary prior, their weighting of context trees correspond to a very specific choice of prior distributions as $h(\tau)$ such that (7)

can be computed recursively based on the nodes of the maximal tree rather than every subtree. To overcome this difficulty, we consider the Partial Bayes Factor (PBF) described in O’Hagan (1995) as an alternative approach for model comparison.

The methodology consists of dividing the data $\tilde{\mathbf{z}}$ into two independent chunks, $\tilde{\mathbf{z}}_{\text{train}}$ and $\tilde{\mathbf{z}}_{\text{test}}$ and then computing the Bayes Factor based on part of the data, $\tilde{\mathbf{z}}_{\text{test}}$, conditioned on $\tilde{\mathbf{z}}_{\text{train}}$ as follows,

$$\text{PBF}_{a,\bar{a}}(\tilde{\mathbf{z}}_{\text{test}}|\tilde{\mathbf{z}}_{\text{train}}) = \frac{\sum_{\tau \in \mathcal{T}_L} \pi_a(\tau|\tilde{\mathbf{z}}_{\text{train}})q(\tau, \tilde{\mathbf{z}}_{\text{test}})}{\sum_{\tau \in \mathcal{T}_L} \pi_{\bar{a}}(\tau|\tilde{\mathbf{z}}_{\text{train}})q(\tau, \tilde{\mathbf{z}}_{\text{test}})}, \quad (10)$$

where $\pi_a(\tau|\tilde{\mathbf{z}}_{\text{train}})$ and $\pi_{\bar{a}}(\tau|\tilde{\mathbf{z}}_{\text{train}})$ are the posterior distributions of τ conditioned on the training data $\tilde{\mathbf{z}}_{\text{train}}$ under the hypotheses H_a and $H_{\bar{a}}$, respectively.

Although the original goal of using PBF is to avoid undefined behaviors when evaluating the ratio of terms involving improper priors, that are replaced by the posterior distributions conditioned on the training sample, we can see that the same strategy is very useful to avoid the intractable normalizing constant from the prior distribution.

Note that, even though (10) still involves sums over \mathcal{T}_L , the terms

$$\sum_{\tau \in \mathcal{T}_L} \pi_a(\tau|\tilde{\mathbf{z}}_{\text{train}})q(\tau, \tilde{\mathbf{z}}_{\text{test}}) \quad \text{and} \quad \sum_{\tau \in \mathcal{T}_L} \pi_{\bar{a}}(\tau|\tilde{\mathbf{z}}_{\text{train}})q(\tau, \tilde{\mathbf{z}}_{\text{test}})$$

can be written as expected values $\mathbb{E}_{H_a}(q(\tau, \tilde{\mathbf{z}}_{\text{test}})|\tilde{\mathbf{z}}_{\text{train}})$ and $\mathbb{E}_{H_{\bar{a}}}(q(\tau, \tilde{\mathbf{z}}_{\text{test}})|\tilde{\mathbf{z}}_{\text{train}})$, which can be obtained from ergodic Markov Chains $(\tau_a^{(t)})_{t \geq 1}$ and $(\tau_{\bar{a}}^{(t)})_{t \geq 1}$ with invariant measures $\pi_a(\tau|\tilde{\mathbf{z}}_{\text{train}})$ and $\pi_{\bar{a}}(\tau|\tilde{\mathbf{z}}_{\text{train}})$, respectively.

Therefore, we can use MCMC methods to approximate Partial Bayes Factors by sampling two Markov Chains $(\tau_a^{(t)})$ and $(\tau_{\bar{a}}^{(t)})$ and using the ratio of empirical averages instead of the expected values

$$\widehat{\text{PBF}}_{a,\bar{a}}(\tilde{\mathbf{z}}_{\text{test}}|\tilde{\mathbf{z}}_{\text{train}}) = \frac{\sum_{t=1}^{n_{\text{iter}}} q(\tau_a^{(t)}, \tilde{\mathbf{z}}_{\text{test}})}{\sum_{t=1}^{n_{\text{iter}}} q(\tau_{\bar{a}}^{(t)}, \tilde{\mathbf{z}}_{\text{test}})}. \quad (11)$$

To avoid the arbitrary segmentation of the dataset into train and test subsets, Berger and Pericchi (1996) proposed the Intrinsic Bayes Factor (IBF), which averages PBFs obtained using different segmentations, based on minimal training samples. Denote by \mathcal{I}_v the collection of subsets of $\{1, \dots, I\}$ of size v , i.e.,

$$\mathcal{I}_v = \{\{i_1, i_2, \dots, i_v\} \subset \{1, 2, \dots, I\}\}.$$

The dataset is divided into minimal training samples, which we will consider a v -tuple of sequences $i_1^v = \{i_1, i_2, \dots, i_v\} \in \mathcal{I}_v$, denoted $\tilde{\mathbf{z}}^{(i_1^v)}$ and the remaining $I-v$ sequences, denoted as $\tilde{\mathbf{z}}^{(-i_1^v)}$ to be used as the test sample. For each possible subset of v sequences i_1^v , we compute the Monte Carlo approximation of the PBF in (11) and take either the arithmetic average to obtain the Arithmetic Intrinsic Bayes Factor (AIBF) or the geometric average for the Geometric Intrinsic Bayes Factor (GIBF).

Denoting by $(\tau_{a,i_1^v}^{(t)})$ and $(\tau_{\bar{a},i_1^v}^{(t)})$, $t = 1, \dots, n_{\text{iter}}$ the Markov Chains of context trees obtained using [Algorithm 1](#) with target distribution $\pi_a(\tau|\mathbf{z}^{(i_1^v)})$ and $\pi_{\bar{a}}(\tau|\mathbf{z}^{(i_1^v)})$ (considering prior distribution proportional to h_a and $h_{\bar{a}}$), respectively, the AIBF and GIBF are defined as

$$\text{AIBF}_{a,\bar{a}}(\tilde{\mathbf{z}}) = \sum_{i_1^v \in \mathcal{I}_v} \frac{1}{\binom{I}{v}} \left(\frac{\sum_{t=1}^{n_{\text{iter}}} q\left(\tau_{a,i_1^v}^{(t)}, \tilde{\mathbf{z}}^{(-i_1^v)}\right)}{\sum_{t=1}^{n_{\text{iter}}} q\left(\tau_{\bar{a},i_1^v}^{(t)}, \tilde{\mathbf{z}}^{(-i_1^v)}\right)} \right),$$

and

$$\text{GIBF}_{a,\bar{a}}(\tilde{\mathbf{z}}) = \prod_{i_1^v \in \mathcal{I}_v} \left(\frac{\sum_{t=1}^{n_{\text{iter}}} q\left(\tau_{a,i_1^v}^{(t)}, \tilde{\mathbf{z}}^{(-i_1^v)}\right)}{\sum_{t=1}^{n_{\text{iter}}} q\left(\tau_{\bar{a},i_1^v}^{(t)}, \tilde{\mathbf{z}}^{(-i_1^v)}\right)} \right)^{\binom{I}{v}^{-1}}.$$

The complete procedure for obtaining these quantities for a given VLMC dataset is described in [Algorithm 2](#).

Note that, while a single sequence ($v = 1$) is theoretically sufficient to identify the context tree and can be considered a minimal training sample, computing posterior distributions using small datasets may result in posterior distributions that assign very low probabilities to context trees with long branches due to smaller total counts on those longer branches. Therefore, using more sequences (higher value for v), may lead to more consistent results as the posterior distribution used in each PBF is more likely to capture long-range contexts. On the other hand, the number of PBFs to be computed is $|\mathcal{I}_v| = \binom{I}{v}$ which quickly becomes prohibitive when v increases. The choice of v is a trade-off between computational cost and the deepness of contexts to be captured by partial posterior distributions.

4 Simulation Studies and Application

To show the strength of our method, we analyzed artificial VLMC datasets generated from two binary models and a real one coming from the field of Linguistics.

4.1 Simulation for binary models

In this section, the primary goal is to examine the performance of the AIBF and GIBF for evaluating the evidence in favor of a null hypothesis of τ being *a-renewing* considering aspects of the effect of the number of independent samples, the size of each chain, and discrimination ability when similar trees are considered. We consider simulations for binary VLMC models with three different sample sizes $I = 3, 10, 25$. For each scenario we sample I chains of equal length, but three different values $T_i = 1000, 2500, 5000$. A dataset was simulated for each combination of I and T_i , resulting in 9 datasets.

The two models considered are presented in [Figure 2](#). Model 1 has a depth equal to 6 with $a = 0$ being a renewal state, while Model 2 is a modified version

Algorithm 2: Complete algorithm for computing AIBF and GIBF with MCMC approximations of Partial Bayes Factors for a dataset $\tilde{\mathbf{z}}$ based on v -tuples.

- 1 **for** $i_1^v \in \mathcal{I}_v$ **do**
- 2 Generate $(\tau_{a,i_1^v}^{(t)})$, $t = 1, \dots, n_{\text{iter}}$ with the Context Tree Metropolis Hastings algorithm with target distribution $\pi_a(\tau|\mathbf{z}^{(i_1^v)})$;
- 3 Generate $(\tau_{\bar{a},i_1^v}^{(t)})$, $t = 1, \dots, n_{\text{iter}}$ with target distribution $\pi_{\bar{a}}(\tau|\mathbf{z}^{(i_1^v)})$;
- 4 Compute the Partial Bayes Factor for the v -tuple i_1^v

$$\widehat{\text{PBF}}_{a,\bar{a}}(\tilde{\mathbf{z}}^{(-i_1^v)}|\mathbf{z}^{(i_1^v)}) = \frac{\sum_{t=1}^{n_{\text{iter}}} q(\tau_{a,i_1^v}^{(t)}, \tilde{\mathbf{z}}^{(-i_1^v)})}{\sum_{t=1}^{n_{\text{iter}}} q(\tau_{\bar{a},i_1^v}^{(t)}, \tilde{\mathbf{z}}^{(-i_1^v)})}.$$
- 5 **end for**
- 6 Return the averages

$$\text{AIBF}_{a,\bar{a}}(\tilde{\mathbf{z}}) = \frac{1}{\binom{I}{v}} \sum_{i_1^v \in \mathcal{I}_v} \widehat{\text{PBF}}_{a,\bar{a}}(\tilde{\mathbf{z}}^{(-i_1^v)}|\mathbf{z}^{(i_1^v)})$$

$$\text{GIBF}_{a,\bar{a}}(\tilde{\mathbf{z}}) = \prod_{i_1^v \in \mathcal{I}_v} \left(\widehat{\text{PBF}}_{a,\bar{a}}(\tilde{\mathbf{z}}^{(-i_1^v)}|\mathbf{z}^{(i_1^v)}) \right)^{\binom{I}{v}^{-1}}$$

with an additional branch grown from the node 01111, which is substituted by the two suffixes 001111 and 101111. Therefore, in Model 2, 0 is no longer a renewal state although both trees are very similar.

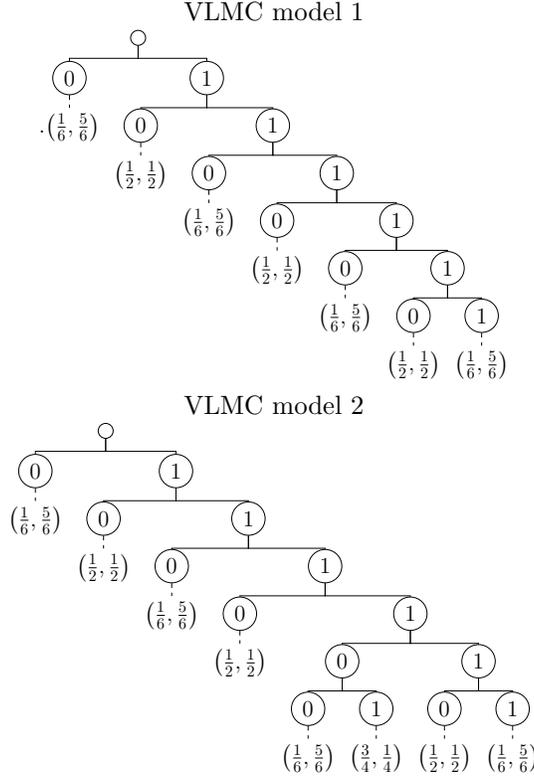


Figure 2: Probabilistic Context Trees for Model 1 and Model 2. The pair of values below each leaf corresponds to the transition probabilities for the suffix associated with that leaf.

For each possible renewal state $a = 0$ or $a = 1$, we compute both Intrinsic Bayes Factors (AIBF and GIBF) using [Algorithm 2](#) considering prior distributions proportional to

$$h_a(\tau) = \mathbf{1}(\tau \in \mathcal{T}_L^a) \quad \text{and} \quad h_{\bar{a}}(\tau) = \mathbf{1}(\tau \in \bar{\mathcal{T}}_L^a)$$

which correspond to the uniform distribution in the space of context trees that are allowed under H_a and $H_{\bar{a}}$, respectively.

For the hyper-parameter α , we choose $\alpha_{sk} = 0.001$ for all s and k , resulting in symmetrical prior distributions for the transition probabilities and a higher density for vectors that are more concentrated.

In each scenario, we considered $v = 1$ and $v = 2$ as the number of sequences

Table 1: AIBF and GIBF computed in \log_{10} scale for simulations for Model 1.

a	I	v	$T_i = 1000$		$T_i = 2500$		$T_i = 5000$	
			AIBF	GIBF	AIBF	GIBF	AIBF	GIBF
0	3	1	4.28	2.39	2.35	2.03	1.69	1.65
0	3	2	1.04	0.66	1.54	1.20	1.88	1.86
0	10	1	25.32	6.00	1.97	1.94	2.55	2.53
0	10	2	4.05	2.31	1.96	1.90	2.49	2.46
0	25	1	68.98	11.87	4.40	2.90	2.59	2.57
0	25	2	67.29	3.54	2.63	2.60	2.59	2.57
1	3	1	-46.87	-49.03	-137.09	-147.48	-285.28	-287.18
1	3	2	-17.82	-20.41	-61.80	-68.78	-135.72	-138.05
1	10	1	-248.24	-265.45	-683.99	-695.28	-1401.88	-1411.24
1	10	2	-234.44	-238.54	-596.78	-616.72	-1237.09	-1252.95
1	25	1	-688.51	-741.06	-1921.76	-1951.27	-3834.79	-3862.61
1	25	2	-705.66	-719.25	-1822.97	-1869.77	-3653.59	-3701.06

to be used for the minimal training sample and ran $n_{\text{iter}} = 10^5$ Metropolis-Hastings steps for each PBF Monte Carlo approximation.

From [Table 1](#), we can see that, for Model 1, both AIBF and GIBF lead to the correct decision for both renewal states tested. For $a = 0$, which was in fact a renewal state, all cases returned a value, in \log_{10} scale, greater than 1 (except one equals 0.66) The discrepancy between AIBF and GIBF diminishes as T_i and I increase converging to a value around 2.5 which was considered Decisive Evidence in the Kass-Raftery scale. On the other hand, for $a = 1$, which was not a renewal state, all cases reported AIBF and GIBF in \log_{10} scale smaller than -15 , converging to values around -3000 when $T_i = 500$ and $I = 25$.

The results for Model 2, which includes a new pair of contexts causing 0 to be no longer a renewal state, presented in [Table 2](#), showed a similar performance rejecting the 1-renewing hypothesis when compared to Model 1, which is expected due to the similarity of both models with respect to the short-length nodes that do not involve 1. The main difference occurred in the computed values for the 0-renewing hypothesis, where the computed GIBFs were positive in logarithmic scale for $T_i = 1000$, for all values of I and v . Moreover, we can see that for $T_i = 5000$, GIBF was negative for all scenarios whereas AIBF was strongly positive for larger values of I . For intermediate value of the sequences sizes ($T_i = 2500$), AIBF pointed to the wrong direction in all scenarios while GIBF identified the right hypothesis in three cases ($I = 10$ and $v = 1$ and 2 and $I = 25$ and $v = 2$). This suggests that for smaller sample sizes, the posterior distributions obtained from the training samples were insufficient to capture the long-range contexts 111100 and 111101 that break the renewal condition of the state 0.

Figures [3](#) and [4](#) present the empirical distributions of PBFs computed in \log_{10} scale for each scenario and each model. In general, the strength of the evidence tends to be larger for $v = 1$ as more data is being used on the test

Table 2: AIBF and GIBF computed in \log_{10} scale for simulations for Model 2.

a	I	v	$T_i = 1000$		$T_i = 2500$		$T_i = 5000$	
			AIBF	GIBF	AIBF	GIBF	AIBF	GIBF
0	3	1	2.58	2.45	14.38	0.06	-7.51	-8.24
0	3	2	1.34	1.22	7.40	0.05	-2.94	-3.29
0	10	1	13.58	4.13	26.40	-7.77	70.50	-38.47
0	10	2	12.49	2.39	26.27	-2.94	65.34	-34.02
0	25	1	46.19	9.11	94.14	13.73	230.50	-95.80
0	25	2	46.49	5.54	93.36	-26.54	227.65	-111.42
1	3	1	-38.24	-40.92	-101.68	-114.08	-216.82	-233.08
1	3	2	-15.16	-18.32	-52.24	-54.20	-111.01	-113.50
1	10	1	-215.74	-222.35	-539.57	-573.34	-1062.60	-1167.87
1	10	2	-190.76	-201.17	-479.56	-509.43	-930.81	-1041.44
1	25	1	-566.23	-601.98	-1432.13	-1528.50	-2867.44	-3186.70
1	25	2	-535.09	-583.43	-1365.51	-1502.80	-2746.98	-3091.29

set, but on the other hand, $v = 2$ leads to more stable PBFs computed. The same behavior is observed as the number of independent sequences I increases, what is expected as adding more data has an impact in the scale of the marginal likelihood function, also rescaling the Bayes Factors.

With $T_i = 2500$, \log_{10} PBFs tend to be distributed around 0, with high variance, as can be observed in [Figure 4](#), resulting in a very unstable average, leading to correct results in some of the scenarios, and incorrect ones in others. As the sample size increases, those long-range contexts are more likely to be captured by the posterior distribution. With $T_i = 5000$, we have decisive evidence that 0 is not a renewal state, except for the scenario with $I = 3$ and $v = 1$, where the value of 0.13 provides very weak evidence, and, in general, the distribution of \log_{10} PBFs was highly concentrated with negative sign, except for a few outliers. AIBF was highly affected by outliers and resulted in incorrect conclusions with high evidence for some scenarios.

Note that, especially for the datasets with small samples ($T_i = 1000$), outliers with large values are observed, having great effect on the computed averages, although the conclusions are not affected. For large datasets ($T_i = 5000$), we have smaller variance in the computed PBFs among different training sets compared to scenarios with sequences of smaller sizes.

Therefore, we conclude that a decision based on GIBF leads to, at least, strong evidence in the scale from [Kass and Raftery \(1995\)](#) (greater than 1/2 in \log_{10} scale) for the correct hypothesis in all cases where we had $v = 2$ and $T_i = 2500$ or 5000. The AIBF was not robust to the presence of outliers in the set of PBFs, leading to incorrect conclusions in the scenarios where the correct detection of the renewal state is harder task and the sample sequences are shorter. The present of outliers also suggests other functions to summarize PBFs other than geometric and arithmetic averages may be useful for avoiding having results highly influenced for the results obtained for particular test samples, like

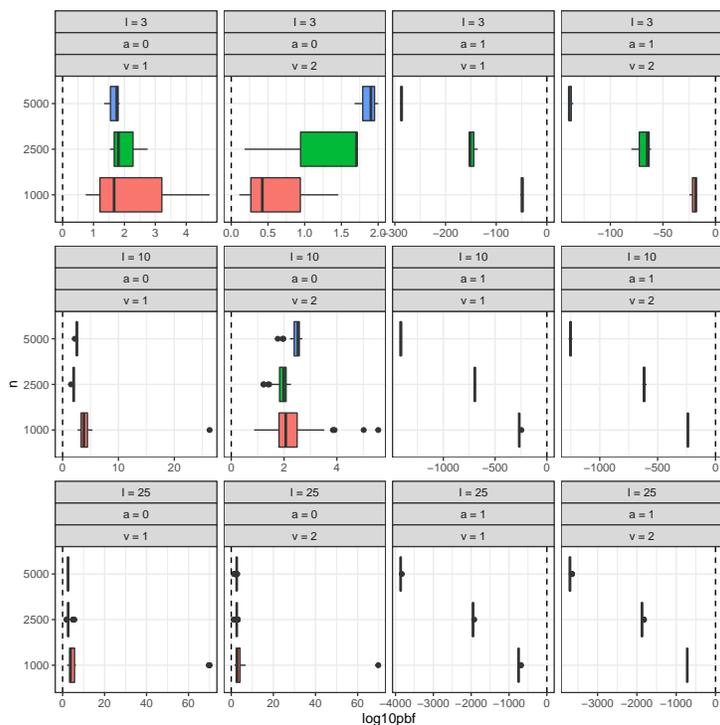


Figure 3: Computed PBFs for Model 1 of the simulation study. Only 0 is a renewal state for this model, therefore, the panels with $a = 0$ are expected to have positive PBFs in logarithmic scale and negative values for $a = 1$.

trimmed averages, removing the most extreme values from the IBF computation, or using the median PBF, which corresponds to a trimmed average trimming all but one value, as used in Charitidou et al. (2018). Context trees that break renewal state condition on long-ranged contexts require larger samples (T_i or v) in order to be captured and result in evidence against the renewal state hypothesis, while states that break renewal state condition in short contexts can be immediately identified, even with short observed sequences.

4.2 Application to rhythm analysis in Portuguese texts

It is known that Brazilian and European Portuguese (henceforth BP and EP) have different syntaxes. For example, Galves et al. (2005) inferred that the placement of clitic pronouns in BP and EP differ in two senses, one of them being: “EP clitics, but not BP clitics, are required to be in a non-initial position with respect to some boundary”. However, the question remains: are the choices of word placement related to different stress patterns preferences? This question was addressed by Galves et al. (2012) that found distinguishing rhythmic

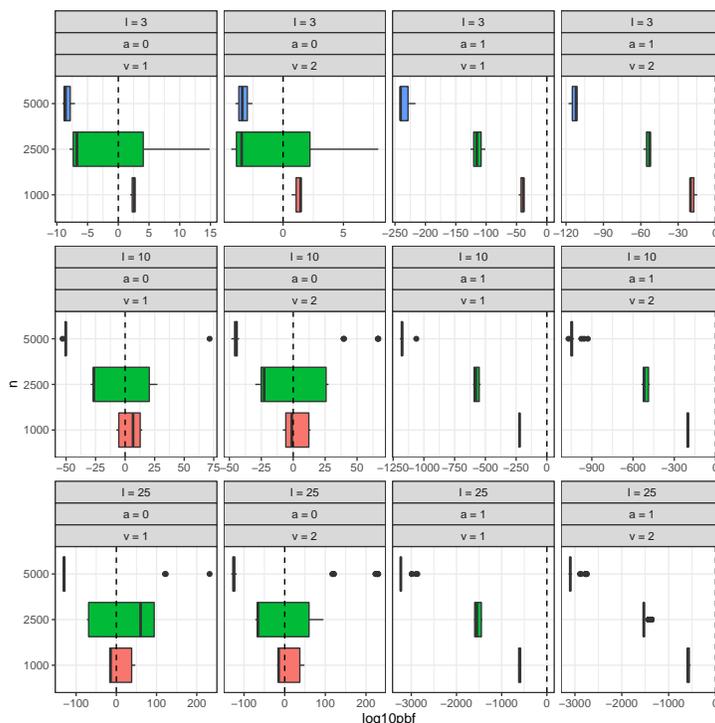


Figure 4: Computed PBFs for Model 2 of the simulation study. Neither 0 or 1 are renewal states for this model, therefore, all panels are expected to have negative PBFs in logarithmic scale. The structure of the tree used makes the renewal state hypothesis violation for $a = 0$ to not be captured (or does not produce strong evidence against the renewal state hypothesis) for lower sample sizes.

patterns for BP and EP based on written journalistic texts. The data consists of 40 BP texts and 40 EP texts randomly extracted from an encoded corpus of newspaper articles from the 1994 and 1995 editions of *Folha de São Paulo* (Brazil) and *O Público* (Portugal). Texts were encoded according on rhythmic features resulting in discrete sequences with around 2500 symbols each and are available at <http://dx.doi.org/10.1214/11-AOAS511SUPP>. After a pre-processing of the texts (removing foreign words, rewriting of symbols, dates, compound words, etc) the syllables were encoded by assigning one of four symbols according to whether or not (i) the syllable is stressed; (ii) the syllable is the beginning of a prosodic word (a lexical word (noun, verbs,...) together with the functional non-stressed words (articles, prepositions, ...) which precede or succeed it). This classification can be represented by 4 symbols. Additionally an extra symbol was assigned to encode the end of each sentence. The alphabet $\mathcal{A} = \{0, 1, 2, 3, 4\}$ was obtained as follows.

- 0 = non-stressed, non-prosodic word initial syllable;
- 1 = stressed, non-prosodic word initial syllable;
- 2 = non-stressed, prosodic word initial syllable;
- 3 = stressed, prosodic word initial syllable;
- 4 = end of each sentence.

For example, the sentence *O sol brilha forte agora*. (The sun shines bright now.) is coded as

Sentence	O	sol	bri	lha	for	te	a	go	ra	.
Code	2	1	3	0	3	0	2	1	0	4

The Smallest Maximizer Criteria proposed by Galves et al. (2012) to select the best tree for BP and EP, uses the fact that the symbol 4 appears as a renewal state to perform Bootstrap sampling. Moreover, they conclude that “the main difference between the two languages is that whereas in BP both 2 (unstressed boundary of a phonological word) and 3 (stressed boundary of a phonological word) are contexts, in EP only 3 is a context.” These are exactly the type of questions to be addresses by the renewal state detection algorithm.

Due to the encoding used, the grammar of the language, and the general structure of written texts, some transitions are not possible. For example, two end of sentences (symbol 4) cannot happen consecutively, therefore, a transition from 4 to 4 is not allowed. Furthermore, there is one, and only one, stressed syllable in each prosodic word. Table 3 summarizes the allowed and prohibited one-step transitions.

Table 3: Allowed transitions for each encoded symbol.

From/To	0	1	2	3	4
0	yes	yes	yes	yes	yes
1	yes	no	yes	yes	yes
2	yes	yes	no	no	no
3	yes	no	yes	yes	yes
4	no	no	yes	yes	no

These prohibited transition conditions are included in the model with proper modifications to the prior distribution, assigning zero probability to some context trees and forcing the probabilities related to prohibited transitions to be zero. The modifications are:

1. If a transition from k to k' is prohibited and a context \mathbf{s} has k as its last symbol, we force $p_{\mathbf{s},k} = 0$ in our prior distribution. The remaining probabilities associated with allowed transitions are then a priori distributed as a Dirichlet distribution with lower dimension.

For example, for a context 102, the only allowed transitions from 2 are to 0 and 1, we have $p_{102,2} = p_{102,3} = p_{102,4} = 0$ with prior probability 1,

and the free probabilities, $(p_{102,0}, p_{102,1})$, distributed as a 2-dimensional Dirichlet distribution with hyper-parameters $(\alpha_{102,0}, \alpha_{102,1})$.

2. If $\mathbf{s} \in \tau$ includes a prohibited transition, then $n_{\mathbf{s},k}(\tilde{\mathbf{z}}) = 0$ as there will not be any occurrences of such sequence in the sample. As a consequence, these suffixes have no contribution in $q(\tau, \tilde{\mathbf{z}})$ as the term related to \mathbf{s} of the product in (5) gets cancelled.
3. We define \mathcal{T}_L^* the space of context trees that do not have prohibited transitions in inner nodes. For example, since the transition 44 is prohibited, a tree that contains a suffix 044 cannot be in \mathcal{T}_L^* because the transition from 4 to 4 is not in a leaf (final node), whereas a tree in \mathcal{T}_L can contain the suffix 44.

Note that, allowing final nodes to contain a prohibited transition is necessary to keep the consistency of our definition based on **full** m -ary trees, as full tree contains the suffix 34 (allowed) if, and only if, it contains another suffix ending in 44 (prohibited), what is not a problem because this prohibited suffix will not contribute to the marginal likelihood.

For each set of sequences in BP and EP, we compute Intrinsic Bayes Factors as evidence for the five hypotheses that 0, 1, 2, 3, and 4 are renewal states. We took $L = 5$ which should be enough to cover all relevant context trees based on the results from the original paper. For prior distributions, we used the same uniform distributions as in the simulation experiment, but restricted to the trees that do not include prohibited transitions on inner nodes, i.e.,

$$h_a(\tau) = \mathbb{1}(\tau \in \mathcal{T}_5^a \cap \mathcal{T}_5^*) \quad \text{and} \quad h_{\bar{a}}(\tau) = \mathbb{1}(\tau \in \bar{\mathcal{T}}_5^a \cap \mathcal{T}_5^*).$$

We also set $\alpha_{\mathbf{s}k} = 0.001$ for every $\mathbf{s} \in \tau$ and $k \in \mathcal{A}$. The algorithm ran for $n_{\text{iter}} = 10^6$ iterations for computing each tree posterior distribution under each hypothesis, using $v = 2$ sequences (around 5000 symbols in each training sample) for each Partial Bayes Factor, resulting in a total of $\binom{40}{2} = 780$ posterior distributions for each hypothesis and PBFs to average.

Due to the numerical instabilities caused by outliers in the set of estimated PBFs as identified in the simulation study, especially in the AIBF, when some particular texts $\mathbf{z}^{(i)}$ are used as the training sample, we also computed a trimmed version of AIBF (and GIBF), which consists of computing the arithmetic (and geometric) average excluding the 10% lowest and 10% highest PBFs, this strategy was also used in [Berger and Pericchi \(1996\)](#). The empirical distributions for the estimated Partial Bayes Factors for all renewal state hypotheses after the 10% trimming are shown in [Figure 5](#).

From the results presented in [Table 4](#), we can see that decisive evidence was obtained when evaluating the renewal hypothesis for states 2, 3 and 4 for the BP dataset and 3 and 4 for the EP dataset which are consistent with the results from [Galves et al. \(2012\)](#).

Finally, for completeness of the Bayesian analysis of this example, we ran the Metropolis Hastings algorithm for the entire BP and EP datasets, considering the same uniform prior distribution on every tree in \mathcal{T}_5^* (no renewal

Table 4: AIBF and GIBF for the BP and EP datasets.

	a	AIBF		GIBF	
		Untrimmed	Trimmed	Untrimmed	Trimmed
BP	0	-10321.12	-10591.73	-10720.48	-10727.80
BP	1	5.88	-7.92	-7.56	-8.55
BP	2	3.49	2.00	1.64	1.64
BP	3	19.41	17.69	14.17	14.10
BP	4	7.96	6.61	6.68	6.61
EP	0	-11487.31	-11641.82	-11744.90	-11747.73
EP	1	7.87	-10.51	-11.67	-11.19
EP	2	1.52	-2.05	-2.56	-2.60
EP	3	18.37	13.49	13.40	13.37
EP	4	12.52	6.66	6.64	6.57

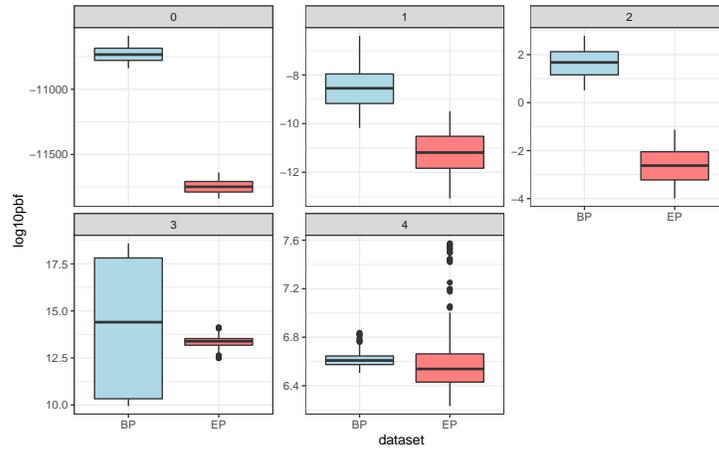


Figure 5: Boxplots presenting the distribution of estimated Partial Bayes Factors $\widehat{\text{PBF}}$ for the $\binom{40}{2}$ pairs of sequences used as training sets, in each of the BP and EP datasets after a trimming of 10% of the highest and lowest PBFs. Each pane corresponds to a hypothesis of a different state being evaluated as a renewal state.

state hypothesis considered, but controlling for prohibited transitions) and with $\alpha_{\mathbf{s},k} = 0.001$ for each valid pair (\mathbf{s}, k) . The two context trees with highest posterior probabilities for BP and EP are presented in Figure 6 and Figure 7, respectively. Leaves corresponding to contexts that contain prohibited transitions, or other contexts with no occurrences are omitted from the trees in the figures for interpretability purposes.

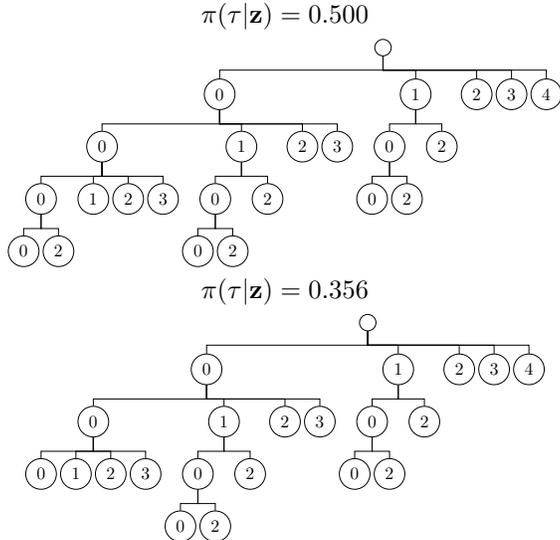


Figure 6: Highest posterior probability context trees for the BP dataset.

5 Conclusions

We propose a Bayesian approach to Variable-Length Markov Chain models with random context trees that allows us to evaluate evidence in favor of renewal hypothesis based on a priori distributions that assign positive probability to each tree on a subset of trees that depends on the renewal state being considered. The main novelties of this work are:

1. The use of Bayes Factor to test the renewal hypothesis for VLMC models;
2. The use of Intrinsic Bayes Factor to evaluate this evidence and overcame the problem of intractable normalizing constant from the prior distribution;
3. The proposal of a Metropolis-Hastings algorithm for sampling context trees that can be performed under a tree prior distribution proportional to any arbitrary function h . This freedom allows not only to incorporate experts prior information about the possible context trees, but also to exclude forbidden trees just by assigning zero probability to them.

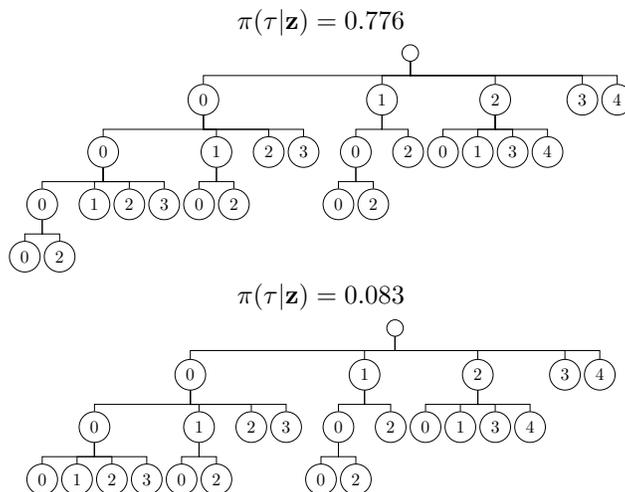


Figure 7: Highest posterior probability context trees for the EP dataset.

To show the strength of our method, we analyzed artificial datasets generated from two binary models and one example coming from the field of Linguistics. The analysis of the artificial datasets suggests that evidence becomes stronger as the number of replicates increases and/or the size of the sequences increases. However, it is possible to obtain good results with as few as 3 replicates of each chain. In the linguistics example we could observe that trimmed GIBF is more robust to possible outliers in the sample.

An R package containing functions for all the computations used in this work is available at github.com/Freguglia/ibfvlmc and R scripts used to reproduce the simulation study are available upon request.

Acknowledgments

This work was funded by Fundação de Amparo à Pesquisa do Estado de São Paulo - FAPESP grants 2017/25469-2 and 2017/10555-0, CNPq grant 304148/2020-2 and by Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001. We thank Helio Migon and Alexandra Schmidt for fruitful discussions.

References

- Aitkin, M. (1991). Posterior bayes factors. *Journal of the Royal Statistical Society: Series B (Methodological)*, 53(1):111–128.
- Aitkin, M. (1993). Posterior bayes factor analysis for an exponential regression model. *Statistics and Computing*, 3(1):17–22.

- Aitkin, M., Finch, S., Mendell, N., and Thode, H. (1996). A new test for the presence of a normal mixture distribution based on the posterior bayes factor. Statistics and Computing, 6(2):121–125.
- Balding, D., Ferrari, P. A., Fraiman, R., and Sued, M. (2009). Limit theorems for sequences of random trees. Test, 18(2):302–315.
- Berger, J. O. and Pericchi, L. R. (1996). The intrinsic bayes factor for model selection and prediction. Journal of the American Statistical Association, 91(433):109–122.
- Bühlmann, P. and Wyner, A. J. (1999). Variable length markov chains. The Annals of Statistics, 27(2):480–513.
- Busch, J. R., Ferrari, P. A., Flesia, A. G., Fraiman, R., Grynberg, S. P., and 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.
- Cabras, S., Castellanos, M. E., and Perra, S. (2015). A new minimal training sample scheme for intrinsic bayes factors in censored data. Computational Statistics & Data Analysis, 81:52–63.
- Charitidou, E., Fouskakis, D., and Ntzoufras, I. (2018). Objective bayesian transformation and variable selection using default bayes factors. Statistics and Computing, 28(3):579–594.
- Chib, S. and Greenberg, E. (1995). Understanding the metropolis-hastings algorithm. The american statistician, 49(4):327–335.
- Csiszár, I. and Talata, Z. (2006). Context tree estimation for not necessarily finite memory processes, via bic and mdl. IEEE Transactions on Information theory, 52(3):1007–1016.
- Dimitrakakis, C. (2010). Bayesian variable order markov models. In Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, pages 161–168. JMLR Workshop and Conference Proceedings.
- Galves, A., Galves, C., Garcia, J. E., Garcia, N. L., and Leonardi, F. (2012). Context tree selection and linguistic rhythm retrieval from written texts. The Annals of Applied Statistics, 6(1):186–209.
- Galves, A. and Löcherbach, E. (2008). Stochastic chains with memory of variable length. Festschrift for Jorma Rissanen, Grünwald et al. (eds), TICSP Series 38:117–133.
- Galves, C., Moraes, M. A. T., and Ribeiro, I. (2005). Syntax and morphology in the placement of clitics in european and brazilian portuguese. Journal of Portuguese Linguistics, 4(2).

- Hastings, W. K. (1970). Monte carlo sampling methods using markov chains and their applications.
- Kass, R. E. and Raftery, A. E. (1995). Bayes factors. Journal of the american statistical association, 90(430):773–795.
- Kontoyiannis, I., Mertzanis, L., Panotopoulou, A., Papageorgiou, I., and Skoularidou, M. (2020). Bayesian context trees: Modelling and exact inference for discrete time series. arXiv preprint arXiv:2007.14900.
- Madigan, D., York, J., and Allard, D. (1995). Bayesian graphical models for discrete data. International Statistical Review/Revue Internationale de Statistique, pages 215–232.
- O’Hagan, A. (1995). Fractional bayes factors for model comparison. Journal of the Royal Statistical Society: Series B (Methodological), 57(1):99–118.
- Rissanen, J. (1983). A universal data compression system. IEEE Transactions on information theory, 29(5):656–664.
- Ron, D., Singer, Y., and Tishby, N. (1996). The power of amnesia: Learning probabilistic automata with variable memory length. Machine learning, 25(2):117–149.
- Villa, C. and Walker, S. G. (2021). An objective bayes factor with improper priors. Computational Statistics & Data Analysis, page 107404.
- Willems, F. M., Shtarkov, Y. M., and Tjalkens, T. J. (1995). The context-tree weighting method: Basic properties. IEEE transactions on information theory, 41(3):653–664.
- Xiong, J., Jääskinen, V., and Corander, J. (2016). Recursive learning for sparse markov models. Bayesian analysis, 11(1):247–263.