# TESTING STATISTICAL HYPOTHESIS ON RANDOM TREES AND applications to the protein classification problem ${ }^{1}$ 

By Jorge R. Busch, Pablo A. Ferrari, Ana Georgina Flesia, ${ }^{2}$<br>Ricardo Fraiman, Sebastian P. Grynberg and<br>Florencia Leonardi ${ }^{3}$<br>Universidad de Buenos Aires, Universidade de São Paulo, Universidad Nacional de Córdoba, Universidad de San Andrés, Universidad de Buenos Aires and Universidade de São Paulo


#### Abstract

Efficient automatic protein classification is of central importance in genomic annotation. As an independent way to check the reliability of the classification, we propose a statistical approach to test if two sets of protein domain sequences coming from two families of the Pfam database are significantly different. We model protein sequences as realizations of Variable Length Markov Chains (VLMC) and we use the context trees as a signature of each protein family. Our approach is based on a Kolmogorov-Smirnov-type goodness-of-fit test proposed by Balding et al. [Limit theorems for sequences of random trees (2008), DOI: 10.1007/s11749-008-0092-z]. The test statistic is a supremum over the space of trees of a function of the two samples; its computation grows, in principle, exponentially fast with the maximal number of nodes of the potential trees. We show how to transform this problem into a max-flow over a related graph which can be solved using a Ford-Fulkerson algorithm in polynomial time on that number. We apply the test to 10 randomly chosen protein domain families from the seed of Pfam-A database (high quality, manually curated families). The test shows that the distributions of context trees coming from different families are significantly different. We emphasize that this is a novel mathematical approach to validate the automatic clustering of sequences in any context. We also study the performance of the test via simulations on Galton-Watson related processes.


1. Introduction. The primary structure of a protein is represented by a sequence of 20 different symbols called amino acids. Proteins can be composed of one or more functional regions, called domains; the identification of domains that occur within a protein can provide insights into its function. For this reason biologists classify protein domains into families and care about the reliability of the classification [Stein (2001)]. But in a protein domain database not only the quality

[^0]of the classification is important, the number of proteins encoded by the genomes that are assigned to the families is also important. This is usually referred to as proteome coverage. For this reason, usually in most databases there must be some balance between quality and quantity.

The Pfam database is a large collection of protein domain families [Finn et al. (2006)]. In its last release of July 2007, the Pfam database comprises 9318 annotated families (Pfam-A) as well as a lower quality, unannotated collection (Pfam-B). Each Pfam-A family consists of two parts: a manually curated set of protein domains called seed and a set of automatically detected protein domains using a profile hidden Markov model (profile HMM), whose parameters are estimated from the seed of the family.

To our knowledge, no independent method to validate the Pfam classification has been proposed, in spite of problems that the uncertainty in the alignment of sequences can lead to [Wong, Suchard and Huelsenbeck (2008)]. We make a step in this direction by presenting a statistical method to test if two samples from protein domains come from two different families. If some families are not significantly different, then the problem of classifying new proteins becomes risky.

We start by modeling protein sequences as Variable Length Markov Chains (VLMC), a model introduced by Rissanen (1983). A VLMC is a discrete time stochastic process with the property that the law of the process at any given time depends on a finite (but not of fixed length) portion of the process at precedent times [Bühlmann and Wyner (1999)]. As usual in the applications of VLMC, we assume that the process is a Markov chain of order at most $L$ (finite memory process). The minimum set of sequences needed to completely specify the distribution of the next symbol in the sequence is known as a context tree and it is denoted by $t$. Calling $p$ the conditional transition probabilities associated to the nodes of $t$, the pair $(t, p)$ completely determines the law of the VLMC.

VLMC have been successfully applied to model and classify protein sequences [Bejerano and Yona (2001)]. As in the case of profile HMM in the construction of the Pfam families, the VLMC approach of Bejerano and Yona takes, for each family, a set of already classified protein domains and estimates a VLMC model, that is, a pair $(t, p)$. Then, the estimated VLMC model is used to classify other protein sequences into the family. Instead, we treat the context trees of sequences of a given family as random samples of a distribution associated to the family; this distribution is used as a signature of the family [Galves et al. (2004), Leonardi et al. (2007)]. That is, we propose that the context trees of the sequences disregarding the associated probabilities are sufficient to test if two samples of sequences come from different Pfam families. We take two samples of protein sequences of size $n$ and $m$ respectively and for each sequence we construct the estimated context tree using the PST algorithm introduced by Ron, Singer and Tishby (1996) and implemented by Bejerano (2004), obtaining two samples of trees $\mathbf{t}=\left(t_{1}, \ldots, t_{n}\right)$, $\mathbf{t}^{\prime}=\left(t_{1}^{\prime}, \ldots, t_{m}^{\prime}\right)$. We assume that the samples are independent and that the trees in each sample are independent and identically distributed with laws $\pi$ and $\pi^{\prime}$,
respectively. We test $H_{0}: \pi=\pi^{\prime}$ against $H_{A}: \pi \neq \pi^{\prime}$ using the test proposed in Balding et al. (2008) (in what follows we will denote it by BFFS test). Rejection of the null hypothesis leads us to conclude that the protein families are distinct.

The BFFS test is a Kolmogorov-Smirnov-type goodness-of-fit test. A distance $d$ defined later in (8) is considered in the space of trees $\mathcal{T}$ and the statistic for the two-sample test is given by

$$
\begin{equation*}
W\left(\mathbf{t}, \mathbf{t}^{\prime}\right):=\sup _{t \in \mathcal{T}}\left|\bar{d}(t, \mathbf{t})-\bar{d}\left(t, \mathbf{t}^{\prime}\right)\right| \tag{1}
\end{equation*}
$$

where $\bar{d}(t, \mathbf{t})=\frac{1}{n} \sum_{i=1}^{n} d\left(t, t_{i}\right)$; that is, $W\left(\mathbf{t}, \mathbf{t}^{\prime}\right)$ is the supremum over $t$ in the space of trees $\mathcal{T}$ of the difference of the empiric mean distances of $t$ to each of the two samples $\mathbf{t}$ and $\mathbf{t}^{\prime}$. The null hypothesis is rejected for large values of $W\left(\mathbf{t}, \mathbf{t}^{\prime}\right)$. Since the law of $W$ under $H_{0}$ is not explicitly known, a simulation procedure is performed to find the $p$-values.

The computation of the test statistic $W\left(\mathbf{t}, \mathbf{t}^{\prime}\right)$ is a priori difficult; a naive search would involve an exponential complexity of the algorithm on the number of potential nodes. A major point of this paper is to show that the problem can be reexpressed as to find the maximal flow on a graph constructed as a function of the samples. The approach is inspired by the search for the Maximum a Posteriori in Bayesian image reconstruction using the Ising model, as proposed by Greig, Porteous and Seheult (1989) [see also Kolmogorov and Zabih (2004)], but requires the introduction of a penalty to guarantee that the solution is in $\mathcal{T}$. The max-flow problem can be solved in polynomial time on the maximal number of nodes of the tree, using Ford-Fulkerson type algorithms.

Statistical analysis of tree-like data has been performed in several papers. Banks and Constantine (1998) obtain trees by hierarchical clustering of authors of written texts, using search-related features. They assume a parametric model and use a metric in the space of trees to get a center point and a confidence band around it. Computation of the distribution's parameters, center point and spread are feasible when a distance of the same type as in the BFFS approach is used. Wang and Marron (2007) analyze a sample of blood vessels in the human brain, represented by trees. Each node represents a blood vessel, and the edges represent vessel connections. The statistical analysis of this data was based on a Fréchet approach, which in turn is based on a metric. The Fréchet mean of a data set is the point which minimizes the sum of the squared distances to the data points. In the specific application to blood vessels, both the structure of the trees (i.e., connectivity properties) and the attributes of the nodes (such as the locations and orientations of the blood vessels) were considered. This is a major difference with the BFFS approach, where only the structure of the trees enters in the test statistics.

In Section 2 we define trees, describe VLMC and explain how to obtain the context trees from the observed protein domain sequences. In Section 3 we define the distance in the tree space and describe the BFFS test. In Section 4 we develop the algorithm to compute the BFFS test statistics. In Section 5 we describe the
one-sample test and discuss possible extensions of the approach. In Section 6.1 we perform pairwise comparisons of samples of trees corresponding to 10 Pfam families. Final remarks are in Section 7 and computing notes in Section 8. At the end of Section 3 and then in Appendix A.1, we discuss problems related to the power of the tests. Appendix A. 2 contains the proofs of selected results. In Appendix A. 3 we perform the test on several samples of Galton-Watson related trees obtained with Monte Carlo simulation.
2. Protein related random trees. A protein sequence can be modeled as a realization of a discrete time stochastic process having as state space the set $A$ of 20 amino acids. This is the basic idea in the modeling of protein domains by HMMs or VLMCs. In this section we introduce the basic concepts behind VLMC and show how the context tree associated to a protein sequence can be estimated using the Probabilistic Suffix Tree (PST) algorithm proposed by Ron, Singer and Tishby (1996) and implemented by Bejerano (2004).

Let $A$ be a finite alphabet and $V=\bigcup_{\ell=0}^{\infty} A^{\ell}$ the set of sequences of symbols in $A$. Denote $a_{\ell}^{j}$ the sequence $a_{\ell} a_{\ell+1} \cdots a_{j}$. Given a sequence $a_{1}^{j}$, any sequence $a_{\ell}^{j}$ with $1<\ell \leq j$ is called a suffix of $a_{1}^{j}$. Let $\mathcal{T}:=\left\{t \subset V: a_{1}^{j} \in t\right.$ implies $\left.a_{2}^{j} \in t\right\}$ be the space of rooted trees with nodes in $V$; the empty sequence is the root of the tree and it is called $\lambda$. The edges of $t$ are $\left\{\left(a_{1}^{j}, a_{2}^{j}\right): a_{1}^{j} \in t\right\}$. A node of an edge is a suffix of the other node and the difference in length of the two nodes is one. Hence, the tree $t$ is identified with its set of nodes.

Let $X=\left(X_{n}\right)_{n \in \mathbb{Z}}$ be a stationary stochastic process taking values in $A$. Define

$$
p\left(a \mid a_{-j}^{-1}\right):=P\left[X_{0}=a \mid X_{-j}^{-1}=a_{-j}^{-1}\right] .
$$

A finite sequences $a_{-k}^{-1} \in V$ is sufficient to determine the law of the next symbol if

$$
\begin{equation*}
p\left(\cdot \mid a_{-j}^{-k-1} a_{-k}^{-1}\right)=p\left(\cdot \mid a_{-k}^{-1}\right) \quad \text { for all } k<j \text { and all } a_{-j}^{-k-1} \in A^{j-k} \tag{2}
\end{equation*}
$$

where $a_{-j}^{-k-1} a_{-k}^{-1}$ denotes the concatenation of the sequences $a_{-j}^{-k-1}$ and $a_{-k}^{-1}$. We assume that the process is a Markov chain of order $L$, that is, that (2) holds for $k=L$ and all $a_{-L}^{-1} \in A^{L}$. A finite sequence $a_{-k}^{-1} \in A^{k}$ is called a context if it satisfies (2) and

$$
\begin{equation*}
p\left(\cdot \mid a_{-k}^{-1}\right) \neq p\left(\cdot \mid a_{-k+1}^{-1}\right) \tag{3}
\end{equation*}
$$

We say that $X$ is a VLMC if there are contexts of length less than $L$, that is, if there exists a $k<L$ and a sequence $a_{-k}^{-1}$ satisfying (2) and (3). The set $t$ of contexts and all their suffixes is called the context tree; each node in the context tree is labeled by a finite sequence over $A$. In this finite memory case, a VLMC is simply a parsimonious representation of a Markov chain of order $L$ that, in its strict sense, would have $(|A|-1)|A|^{L}$ parameters (a probability distribution associated to each sequence of fixed length $L$ ).
(a)

(b)


FIG. 1. An example of stationary conditional probability distributions $p$ over the alphabet $A=\{1,2\}$ (a) and the corresponding context tree $t$ (b). $p$ and $t$ completely specify a stationary VLMC process $X=\left(X_{n}\right)_{n \in \mathbb{Z}}$. We assume $0<\alpha<1, \alpha \neq 0.5$. Each node in the tree (a) is labeled by a sequence over the alphabet $A$ and has an associated probability distribution over $A$ (see text for more details). The contexts of the process are the sequences in bold face $\{111,211,122,222\}$ (note that 21 and 12 are not contexts in our definition). The context tree in this case is (b), representing the set $\{\lambda, 1,2,11,22,111,211,122,222\}$.

Under the assumption of bounded memory, the pair composed by the context tree and the set of transition probability distributions associated to the nodes of $t$ completely specify the law of the stationary process $X$.

Figure 1 summarizes an example of a context tree and transition probabilities for a stationary process $X$ over the alphabet $A=\{1,2\}$. For $0<\alpha<1$ the transition probabilities are given by

$$
p\left(1 \mid a_{-\infty}^{-1}\right)= \begin{cases}\alpha, & \text { if } a_{-3}^{-1}=111 \text { or } a_{-3}^{-1}=122  \tag{4}\\ 1-\alpha, & \text { if } a_{-3}^{-1}=211 \text { or } a_{-3}^{-1}=222 \\ 0.5, & \text { otherwise }\end{cases}
$$

see Figure $1(\mathrm{a})$. If $\alpha \neq 0.5$, the set of contexts is $\{111,122,211,222\}$ and the context tree is $t=\{\lambda, 1,2,11,22,111,211,122,222\}$; see Figure 1(b). If $\alpha=0.5$, the context tree is just $\lambda$, as the chain is a sequence of i.i.d. (0.5) random variables.

There are several approaches to estimate the context tree and transition probabilities of VLMCs from a finite realization of $X$. We mention the context algorithm proposed by Rissanen (1983); see also Bühlmann and Wyner (1999) and Galves et al. (2008). Recently, Csiszár and Talata (2006) proposed the use of the Bayesian Information Criterion (BIC) and the Minimum Description Length Principle (MDL). These algorithms provide consistent estimates of the parameters. Our work utilizes the PST algorithm and so we provide a brief review here.

Suppose $x_{1}, \ldots, x_{l}$ is a sample of a VLMC over $A$ specified by the pair $(t, p)$ (in our setting $x_{1}, \ldots, x_{l}$ represents a protein over the alphabet of 20 amino acids).

For any sequence $a_{1}^{j} \in A^{j}$ define the counters

$$
\begin{equation*}
N\left(a_{1}^{j}\right)=\sum_{i=0}^{l-j} \mathbf{1}\left\{x_{i+1}^{i+j}=a_{1}^{j}\right\} \tag{5}
\end{equation*}
$$

where the function 1 takes value 1 if $x_{i+1}^{i+j}=a_{1}^{j}$ and 0 otherwise. For any sequence $a_{1}^{j} \in A^{j}$ such that $N\left(a_{1}^{j}\right) \geq 1$ and any symbol $a \in A$, we define the empirical transition probabilities $\hat{p}\left(a \mid a_{1}^{j}\right)$ as

$$
\begin{equation*}
\hat{p}\left(a \mid a_{1}^{j}\right)=\frac{N\left(a_{1}^{j} a\right)}{\sum_{b \in A} N\left(a_{1}^{j} b\right)} . \tag{6}
\end{equation*}
$$

To estimate the context tree associated to the sequence, two parameters are fixed: $L$, the maximal depth of the estimated tree $\hat{t}$ and $r>1$, a threshold value. The PST algorithm defines the context tree estimator $\hat{t}$ as the tree containing all the sequences $a_{1}^{j}$, with $j \leq L, N\left(a_{1}^{j}\right) \geq 1$, such that there exists a symbol $a \in A$ satisfying

$$
\begin{equation*}
\left|\log \hat{p}\left(a \mid a_{2}^{j}\right)-\log \hat{p}\left(a \mid a_{1}^{j}\right)\right| \geq \log r \tag{7}
\end{equation*}
$$

That is, the node $a_{1}^{j}$ is a node of $\hat{t}$ if the conditional probabilities $\hat{p}\left(\cdot \mid a_{1}^{j}\right)$ and $\hat{p}\left(\cdot \mid a_{2}^{j}\right)$ are sufficiently far in the sense of (7); this is the empirical version of (2)-(3). To guarantee that $\hat{t}$ is a tree, include also all suffixes of included nodes; that is, $a_{1}^{j} \in \hat{t}$ implies $a_{2}^{j} \in \hat{t}$.

The PST algorithm uses other parameters to smooth the estimated transition probabilities given by (6). This smoothing is useful to avoid null estimated probabilities that can damage the prediction step in the classification of new sequences. Since our interest is to estimate only the context tree, it is sufficient to consider the parameters $L$ and $r$. We refer the interested reader to Ron, Singer and Tishby (1996), Bejerano (2003) and Bejerano (2004) for a full explanation of the PST algorithm, its implementation and some basic examples.
3. The tree distance and the two-sample test. For the two sample problem, the BFFS test is based on the supremum (over the space of trees) of the difference between the empirical mean distance function of each sample to a given tree. More formally, for a node $v \in V$ and a tree $t$ in $\mathcal{T}$ denote $t(v)=\mathbf{1}\{v$ is a node of $t\}$. Let $\phi: V \rightarrow \mathbb{R}^{+}$be a nonnegative function and consider the distance in $\mathcal{T}$ defined by

$$
\begin{equation*}
d\left(t, t^{\prime}\right):=\sum_{v \in V} \phi(v)\left(t(v)-t^{\prime}(v)\right)^{2} \tag{8}
\end{equation*}
$$

Let $T$ be a random tree on $\mathcal{T}$ with law $\pi$ and $\mathbf{t}=\left(t_{1}, \ldots, t_{n}\right)$ a random sample of $T$ (independent random trees with the same law as $T$ ). Define the empiric expected
distance of a tree $t$ to the sample by

$$
\begin{equation*}
\bar{d}(t, \mathbf{t}):=\frac{1}{n} \sum_{i=1}^{n} d\left(t_{i}, t\right) \tag{9}
\end{equation*}
$$

Consider two samples $\mathbf{t}$ and $\mathbf{t}^{\prime}$ of random trees $T$ and $T^{\prime}$, with laws $\pi$ and $\pi^{\prime}$, with sample sizes $n$ and $m$ respectively. The two-sample problem is to test

$$
\begin{equation*}
H_{0}: \pi=\pi^{\prime}, \quad H_{A}: \pi \neq \pi^{\prime} \tag{10}
\end{equation*}
$$

BFFS show that, under $H_{0}$, the process

$$
\begin{equation*}
\left(\sqrt{\frac{n m}{n+m}}\left(\bar{d}(t, \mathbf{t})-\bar{d}\left(t, \mathbf{t}^{\prime}\right)\right), t \in \mathcal{T}\right) \tag{11}
\end{equation*}
$$

converges weakly as $\min (n, m) \rightarrow \infty$ to a Gaussian process and propose the statistic

$$
\begin{equation*}
W\left(\mathbf{t}, \mathbf{t}^{\prime}\right):=\sup _{t \in \mathcal{J}}\left|\bar{d}(t, \mathbf{t})-\bar{d}\left(t, \mathbf{t}^{\prime}\right)\right| . \tag{12}
\end{equation*}
$$

Under $H_{0}$ for large $n$ and $m, \sqrt{\frac{n m}{n+m}} W\left(\mathbf{t}, \mathbf{t}^{\prime}\right)$ has approximately the law of the supremum over $t$ of a Gaussian process indexed by $t \in \mathcal{T}$. Determining the quantiles $q_{\alpha}$ using the asymptotic law, the null hypothesis is rejected at level $\alpha$ when

$$
\left|W\left(\mathbf{t}, \mathbf{t}^{\prime}\right)\right|>q_{\alpha} .
$$

The quantiles are obtained using permutation-based randomization techniques [Manly (2007)]. See Section 6.1 for details.

About the power of the BFFS test. Strictly speaking, the null hypothesis of the BFFS test is

$$
H_{0}^{\prime}: \text { law of }(d(t, T), t \in \mathcal{T})=\text { law of }\left(d\left(t, T^{\prime}\right), t \in \mathcal{T}\right)
$$

Of course, rejection of $H_{0}^{\prime}$ implies rejection of $H_{0}$, so that we do not need to worry when rejecting. But the test could accept $H_{0}^{\prime}$ even when $H_{0}$ is false. We give in Appendix A. 1 examples of different $\pi \mathrm{s}$ giving rise to the same process ( $d(t, T), t \in \mathcal{T})$ and show some sufficient conditions on $\pi$ under which the law of ( $d(t, T), t \in \mathcal{T}$ ) determines $\pi$.
4. Graph computation of $W\left(\mathbf{t}, \mathbf{t}^{\prime}\right)$. To compute the test statistic $W\left(\mathbf{t}, \mathbf{t}^{\prime}\right)$, it is necessary to find the trees attaining the supremum (12). In this section we show that the problem can be reformulated in terms of finding the maximal flow on a graph constructed as a function of the sample.

Denote by $\bar{t}$ the empiric mean of the sample $\mathbf{t}$ :

$$
\begin{equation*}
\bar{t}(v):=\frac{1}{n} \sum_{i=1}^{n} t_{i}(v) . \tag{13}
\end{equation*}
$$

Since $\bar{t}(v)$ is not always an integer, $\bar{t}$ is not necessarily a tree, but $\bar{t}(a v) \leq \bar{t}(v)$ if $v$ and $a v$ are in $V$. Notice that

$$
\begin{equation*}
\bar{d}(t, \mathbf{t})-\bar{d}\left(t, \mathbf{t}^{\prime}\right)=2 \mathcal{L}(t)+\sum_{v \in V} \phi(v)\left(\bar{t}(v)-\overline{t^{\prime}}(v)\right), \tag{14}
\end{equation*}
$$

where

$$
\begin{equation*}
\mathcal{L}(t)=\sum_{v \in V} \phi(v)\left(\overline{t^{\prime}}(v)-\bar{t}(v)\right) t(v) . \tag{15}
\end{equation*}
$$

Since the last term in (14) does not depend on $t$, to maximize $\left|\bar{d}(t, \mathbf{t})-\bar{d}\left(t, \mathbf{t}^{\prime}\right)\right|$ on $\mathcal{T}$ is equivalent to minimize $\mathcal{L}(t)$ and $-\mathcal{L}(t)$ on $\mathcal{T}$.

Define the space of configurations

$$
y:=\{0,1\}^{V} .
$$

This set can be identified with the set of subsets of $V$, so that $\mathcal{T} \subset \mathcal{y}$. In order to penalize configurations in $\mathcal{y}$ that are not trees, we define the quadratic function $\mathcal{P}: \mathscr{Y} \rightarrow \mathbb{N}$ which counts the number of orphan nodes in a configuration:

$$
\begin{equation*}
\mathcal{P}(y)=\sum_{\{v, a v\} \subset V} y(a v)(1-y(v)), \quad y \in \mathcal{y}, \tag{16}
\end{equation*}
$$

where for a node $v=a_{1} \cdots a_{j}, a v=a a_{1} \cdots a_{j}$ is a son of $v$. It is clear that $\mathcal{P}(y) \geq 0$ and $\mathscr{P}(y)=0$ if and only if $y \in \mathcal{T}$.

The following proposition shows that to maximize $\left|\bar{d}(t, \mathbf{t})-\bar{d}\left(t, \mathbf{t}^{\prime}\right)\right|$ on $\mathcal{T}$ is equivalent to minimize $\mathcal{L}(y)+\beta \mathcal{P}(y)$ and $-\mathcal{L}(y)+\beta \mathcal{P}(y)$ on $\mathcal{y}$ for $\beta$ sufficiently large.

Proposition 4.1. Let $\beta>\sum_{v \in V} \phi(v)$. Then

$$
\begin{align*}
& \arg \max _{t \in \mathcal{T}}\left|\bar{d}(t, \boldsymbol{t})-\bar{d}\left(t, \boldsymbol{t}^{\prime}\right)\right| \\
& \quad \subset \arg \min _{y \in \mathcal{Y}}(\mathscr{L}(y)+\beta \mathcal{P}(y)) \cup \arg \min _{y \in \mathcal{Y}}(-\mathcal{L}(y)+\beta \mathcal{P}(y)) . \tag{17}
\end{align*}
$$

The proof of Proposition 4.1 is given in Appendix A.2. The proposition reduces the minimization problem on $\mathcal{T}$ to the task of minimizing the Hamiltonians $\mathcal{L}+\beta \mathcal{P}$ and $-\mathscr{L}+\beta \mathcal{P}$ on $\mathcal{y}$.

The Hamiltonian $\mathcal{L}+\beta \mathcal{P}$ is represented by the oriented graph $(\widetilde{V}, \widetilde{E})$ given by

$$
\begin{equation*}
\widetilde{V}:=V \cup\{s\} \cup\{b\}, \quad \widetilde{E}:=\{(v, w): v, w \in \tilde{V}\} \tag{18}
\end{equation*}
$$

where $s$ (source) and $b$ (sink) are two extra nodes. The graph has the following capacities associated to the (oriented) edges:

$$
c(v, w):= \begin{cases}\left(\phi(v)\left(\bar{t}(w)-\overline{t^{\prime}}(w)\right)\right)^{+}, & \text {if } v=s \text { and } w \in V  \tag{19}\\ \left(\phi(v)\left(\bar{t}(v)-\overline{t^{\prime}}(v)\right)\right)^{-}, & \text {if } v \in V \text { and } w=b, \\ \beta, & \text { if } v \in V, w=a v \in V, a \in A \\ 0, & \text { otherwise },\end{cases}
$$

where $x^{+}=\max \{x, 0\}, x^{-}=\max \{-x, 0\}$. That is, the edges linking a node of $V$ to its sons have capacity $\beta$, the edge linking a node of $V$ to the sink, and the edge linking the source to a node of $V$ have capacity $\left.\phi(v)\left(\bar{t}(v)-\overline{t^{\prime}}(v)\right)\right)^{ \pm}$according to the sign of $\left(\bar{t}(v)-\overline{t^{\prime}}(v)\right)$; the other edges have zero capacity.

A configuration $y \in \mathcal{Y}$ defines a cut of the graph

$$
\begin{equation*}
C(y):=\{(v, w) \in \widetilde{E}: c(v, w)>0, v \in(V \backslash y) \cup\{s\}, w \in y \cup\{b\}\} \tag{20}
\end{equation*}
$$

whose capacity $c(y)$ is

$$
\begin{equation*}
c(y):=\sum_{(v, w) \in C(y)} c(v, w) . \tag{21}
\end{equation*}
$$

The next result has been proven by Kolmogorov and Zabih (2004), as a generalization of an approach of Greig, Porteous and Scheult (1989). We give some details of the proof in Appendix A.2.

PROPOSITION 4.2. It holds that $c(y)=k+\mathcal{L}(y)+\beta \mathcal{P}(y)$ for all $y \in \mathcal{y}$, where $k$ does not depend on $y$.

Proposition 4.2 shows that to minimize $\mathcal{L}(y)+\beta \mathcal{P}(y)$ it is sufficient to find a minimum cut in its associated graph. This problem can be solved by means of the Ford-Fulkerson type of algorithm as proposed by Greig, Porteous and Scheult (1989). We use the variant and implementation of Kolmogorov and Zabih (2004).

The idea behind the Ford-Fulkerson algorithm is the following. Suppose that liquid is flowing from source to sink in the graph with nodes $\widetilde{V}$ and pipes $\widetilde{E}$ with capacities $c(\cdot, \cdot)$. Take a piece of chalk and draw an arrow in the direction of the flow over the pipes with positive capacity that are not totally filled. Draw an arrow in the direction opposite to the flow over the pipes carrying some liquid. Of course there may be pipes with arrows in both directions! Now try to walk from the source to the sink, always following your arrows. When you arrive at a dead end, return to the source. If you never arrive to the sink, the flow is maximal and the nodes $y$ that you have not visited define a cut $C(y)$ with minimal capacity. If you arrive to the sink, you can increase the total flow by $\epsilon$ by increasing it by $\epsilon$ in the pipes that you have walked forward from the source, and decreasing it by the same amount in the pipes that you walked backward. It turns out that the number of operations necessary to find the minimal cut is polynomial in the number of nodes.

## 5. Generalizations.

The one-sample test. Given a sample $\mathbf{t}$ of a random tree with law $\pi$, the onesample test is

$$
\begin{equation*}
H_{0}: \pi=\pi^{\prime}, \quad H_{A}: \pi \neq \pi^{\prime} \tag{22}
\end{equation*}
$$

for a given probability $\pi^{\prime}$ on $\mathcal{T}$. The BFFS statistic in this case is given by

$$
\begin{equation*}
W(\mathbf{t}):=\sup _{t \in \mathcal{T}}|\bar{d}(t, \mathbf{t})-\pi d(t)|, \tag{23}
\end{equation*}
$$

where

$$
\begin{equation*}
\pi d(t):=\sum_{t^{\prime} \in \mathcal{T}} \pi\left(t^{\prime}\right) d\left(t^{\prime}, t\right) \tag{24}
\end{equation*}
$$

is the expected distance between $t$ and a random tree with law $\pi$. The graph computation of $W(\mathbf{t})$ is done exactly as in Section 4, but in the definition (15) of $\mathcal{L}$ the mean occupation value $\overline{t^{\prime}}(v)$ is substituted by

$$
\begin{equation*}
\mu_{\pi^{\prime}}(v):=\sum_{t} \pi^{\prime}(t) t(v) \tag{25}
\end{equation*}
$$

the average occupation number of node $v$ under $\pi^{\prime}$.
More general trees. The BFFS test works for finite or infinite rooted trees contained in a full tree $V . V$ must satisfy that the number of children per parent is uniformly bounded by $m$ (say). This condition is necessary for the Proposition 4.1, which transforms the problem of minimizing the difference of the distances on the minimization of a Hamiltonian. The nodes in $V$ can be coded with finite sequences of letters of the alphabet $A=\{1, \ldots, m\}$ in such a way each node is coded with the sequence coding of his father plus a letter of $A$. In our case the letter is added at the beginning of the sequence so that the sequence corresponding to a parent is a suffix of those corresponding to its children. Alternatively, the letter can be added at the end; in this case the parents sequences are prefixes of the children. Any other labeling would work in the same way, as only the structure of the tree (and not the labeling of the nodes) is relevant in the construction of the test. The structure of the tree is the set of nodes and the set of edges; with our coding the set of edges is deduced from the node coding: $E=\left\{\left(a_{1}^{j}, a_{2}^{j}\right)\right\}$, otherwise $E$ must be explicitly defined.

If $V$ is infinite, $V$ is truncated to nodes with at most $L$ symbols and calling $\pi_{L}$ the law of the truncated tree, the null hypothesis is $H_{0}: \pi_{L}=\pi_{L}^{\prime}$.

## 6. Numerical results.

6.1. Testing protein related populations of trees. In this subsection we present some results obtained by applying the two-sample test over protein domain families of the Pfam-A database. As mentioned in the Introduction, our framework is the following:

- Each family $\mathcal{F}$ of protein domains induce a (different, hopefully) probability distribution $\pi$ on the space of trees $\mathcal{T}$.
- Given two families $\mathcal{F}$ and $\mathcal{F}^{\prime}$, we consider their associated signatures, that is, the probability laws $\pi$ and $\pi^{\prime}$ on the space $\mathcal{T}$.
- For each family $\mathcal{F}_{j}$ we take a sample of protein sequences of size $n_{j}$, and for each sequence in the sample we construct the PST context tree estimator, as described in Section 2. We obtain a sample of size $n_{j}$ of i.i.d. random elements on $\mathcal{T}$ with distribution $\pi_{j}$.
- Finally, for each pair of families $\mathcal{F}_{j}, \mathcal{F}_{j^{\prime}}$ we test if both distributions $\pi_{j}$ and $\pi_{j^{\prime}}$ are the same.

To test the approach, we randomly choose families $\mathcal{F}_{1}, \ldots, \mathcal{F}_{10}$ whose names start with letter A and such that their average lengths are larger than 150 amino acids (this last condition was to ensure some precision in the context tree estimation step). In order to guarantee the quality of the samples, we only choose sequences in the seed of each family. The chosen families are ABC-2membrane, ABC-membrane, Amidase, Amidino-trans, AMME-CR1, AOX, ArgK, ASC, Asp-Arg-Hidrox and Asp-Al-Ex.

We randomly select $n_{j}=50$ sequences from each family $\mathcal{F}_{j}$ and compute the associated PST context tree estimator of each sequence using the PST algorithm with parameters $r=1.05$ [as Bejerano (2001)] and $L=4$. In this way we obtain a sample of 50 trees per family.

We consider the distance function $\phi(v)=\theta^{\operatorname{gen}(v)}$, where the function gen $(v)$ is defined as the length of the sequence labeling node $v$. That is, if node $v$ corresponds to sequence $a_{1}^{k}$, then $\operatorname{gen}(v)=k$. For each $\theta \in\{0.001,0.01,0.35\}$ we run the BFFS test for each pair of families using the corresponding samples of trees. We also run the tests under the null hypothesis collecting two independent samples from the same family. For each pair of samples of trees we estimate the $(1-\alpha)$-quantile under the null hypothesis using Monte Carlo randomization [Manly (2007)], that is, we permute the pooled sample a thousand times and compute the test statistic for each of the replicates using half of the permuted pooled sample for each population. The estimated quantile is therefore the empirical $(1-\alpha)$-quantile for the vector of size 1000 built up in this way.

All 45 tests are rejected at level $\alpha=0.001$ for the three values of $\theta$. Despite the conservative level used in the tests, the hypothesis of equal distribution is rejected in all cases when the samples came from different Pfam-A families, confirming the discriminative power of the context trees associated to the sequences. In the case of the same family, but with independent samples of trees, for $\theta=0.35$ we observe $p$-values ranging from 0.15 to 0.87 , compatible with the uniform distribution (the law of the $p$-values under the null hypothesis). Similar results were obtained with the other two values of $\theta$.
6.2. Simulation results. We also challenge our method in a small Monte Carlo simulation for Galton-Watson processes, for three different models, parameters and sample sizes. The results are reported in Appendix A.3.
7. Final remarks. We perform the BFFS method to test if two samples of context trees come from different distributions, and we propose a feasible way to compute its statistic, allowing the treatment of reasonably big trees. The test rejects the null hypothesis in the case of high quality, manually curated Pfam families, and it does not reject on random subsets of the same family. This supports the use of the test as a method to distinguish different groups of protein domains when a specific task, as, for example, sequence annotation, does not give conclusive results.

Our results strongly indicate that the context trees associated to protein domain sequences are sufficient to discriminate between different families in the Pfam database. In this sense we have benefited from ideas coming from the analysis of sequences related to linguistics. Galves and collaborators have proposed with success the use of context trees to discriminate languages from codified written text [Galves et al. (2004)]. More recently, they have used similar ideas in a preliminary work to study the phylogeny of protein sequences [Leonardi et al. (2007)].

We emphasize that the test is not restricted to the analysis of samples of context trees. Any space of trees satisfying the assumptions of Section 5 will be suitable for using our approach. On the other hand, for particular distributions like Galton-Watson processes, more simple tailor-made tests can be developed. Our simulations show that the BFFS test is able to distinguish between distributions determined by the node-marginal distributions, which is a large family of distributions for applications. This class includes tree laws with a Markovian hypothesis, as shown in Proposition A.2. We discuss this item in detail in Appendix A.1.
8. Computing notes. The code to compute the test statistic is available from Jorge R. Busch (jbusch@fi.uba.ar) upon request. Calculation reported here used Scilab INRIA (http://www.scilab.org/) and C++ code from Bejerano (2004) and Kolmogorov and Zabih (2004).

The computational burden for our algorithm allows us to work with trees with up to $3^{20}$ nodes. Each $p$-value involves 1001 test statistic calculations, with a sample size $n=50$ for each family, taking at most 15 minutes to be complete, with a Pentium Core 2 duo with 2Gb of RAM memory.

## APPENDIX

## A.1. Mean distances and Markovian hypotheses.

Do the mean distances determine a measure? Recall the mean distance $\pi d$ is defined in (24). Our test is universally consistent within the class of distributions for which $(\pi d(t), t \in \mathcal{T})$ determine the probability $\pi$. That is, whatever is the law of the families, the test will asymptotically detect the difference. We show in Lemma A. 1 that $\pi d$ determines the law of the marginals $(T(v), v \in V)$. Proposition A. 2 says that, under Markov type hypotheses, the marginal distributions determine the measure.

For a random tree $T$ with law $\pi$ recall $\mu_{\pi}(v)$ is the mean occupancy node $v$ defined in (25) and define $\sigma_{\pi}^{2}(v)$, the variance of $T(v)$, by

$$
\begin{equation*}
\sigma_{\pi}^{2}(v)=\mu_{\pi}(v)\left(1-\mu_{\pi}(v)\right) \tag{A.1}
\end{equation*}
$$

Lemma A.1. Let $\pi$ and $\pi^{\prime}$ be measures on $\mathcal{T}$. Then, $\pi d=\pi^{\prime} d$ if and only if $\mu_{\pi}=\mu_{\pi^{\prime}}$.

The proof is given later in this section.
Different measures may have the same mean distances. For instance, consider $A=\{1,2\}$ and $\pi, \pi^{\prime}$ defined by

$$
\begin{aligned}
\pi(\varnothing)=\frac{1}{2} ; & \pi(\{\lambda\})=\pi(\{\lambda, 1\})=\pi(\{\lambda, 2\})=\pi(\{\lambda, 1,2\})=\frac{1}{8} \\
\pi^{\prime}(\varnothing)=\frac{1}{2} ; & \pi^{\prime}(\{\lambda\})=\pi^{\prime}(\{\lambda, 1,2\})=\frac{1}{4} .
\end{aligned}
$$

Then, $\mu_{\pi}(v)=\mu_{\pi^{\prime}}(v)$ for all $v \in V$.
Lemma A. 1 implies that in general the functions $\pi d$ and $\pi^{\prime} d$ do not help to solve the discrimination problem. But in some cases these functions do discriminate. To show that, we need some extra notation. For a set $I$ of nodes denote $T_{I}$ the restriction of $T$ to $I$ and $T_{I}=1$ means that $T(v)=1$ for all $v \in I$, while $T_{I}=0$ means that $T(v)=0$ for all $v \in I$.

Let $v$ be a node, and $a, b \in A$. We shall call $v$ father of $a v, a v$ son of $v$, and $a v$ brother of $b v$. Let $f: V \backslash\{\lambda\} \rightarrow V$ be a function such that, for each $v \neq \lambda, f(v)$ is father or brother of $v$, and $f^{n}(v)=\lambda$ for some $n=n(v) \in \mathbb{N}$. Notice that, in this case, $f^{-1}(v)$ is empty or formed up by brothers and sons of $v$. We call such a function a tree-shift. Consider, for instance, the function $f$ that assigns to a node its father.

Let $T$ be a random tree with law $\pi$. We say that $T$ satisfies a Markov hypothesis if there exists a tree-shift $f$ such that if $v=f(w)$, then
(a) $0<\mu_{\pi}(w)<\mu_{\pi}(v)<1$,
(b) $\mathbb{P}(T(w)=1 \mid T(v)=0)=0$,
(c) let $I, J \subset V$ be such that if $v \in I, w \notin I \cup J$ and $\mathbb{P}\left(T_{I}=1, T_{J}=0\right)>0$, then

$$
\begin{equation*}
\mathbb{P}\left(T(w)=1 \mid T_{I}=1, T_{J}=0\right)=\mathbb{P}(T(w)=1 \mid T(v)=1) \tag{A.2}
\end{equation*}
$$

Proposition A.2. Under the Markov hypotheses, the marginals ( $\mu_{\pi}(v), v \in$ $V)$ determine the probabilities $(\pi(t), t \in \mathcal{T})$.

The proof is given later in this section.

Examples of measures satisfying the Markov hypotheses. The alphabet for the following examples is $A=\{1, \ldots, m\}$.

1. Let $f$ be defined by $f(a v)=v, a \in A$. Let $k(v)$ be such that $f^{k(v)-1}(v)=1$, for $v \in V$. If $\mu_{\pi}(v)=p^{k(v)}(0<p<1)$, we obtain a tree with $\pi(\{\lambda\})=p$, and when $T(v)=1, T(a v)$ is Bernoulli with parameter $p$, for $a \in A$. We call such tree distributions pseudo Galton-Watson processes.
2. Let $f$ be defined by $f(1 v)=v$, and $f((a+1) v)=a v$ for $1 \leq a<m$. That is, $v=f(w)$ if $w$ is the eldest brother and $v$ is the father of $w$, or if $v$ is the nearest older brother of $v$. Let now $p_{0}, \ldots, p_{m}$ be given probabilities, with $p_{0}>0$ and $p_{0}+\cdots+p_{m}=1$. If

$$
\begin{aligned}
\mu_{\pi}(1 v) & =\left(p_{1}+\cdots+p_{m}\right) \mu_{\pi}(v), \\
\mu_{\pi}((a+1) v) & =\frac{p_{a+1}+\cdots+p_{m}}{p_{a}+\cdots+p_{m}} \mu_{\pi}(a v) \quad(1 \leq a \leq m-1),
\end{aligned}
$$

we obtain the classical Galton-Watson process, with parameter probabilities $p_{0}, \ldots, p_{m}$.

If $\pi$ is a distribution on $\mathcal{T}$, and $T$ is a random tree with distribution $\pi$, then by a simple computation,

$$
\begin{align*}
\pi d(t) & =\sum_{v \in V} \phi(v)\left(\mu_{\pi}(v)-t(v)\right)^{2}+\sum_{v \in V} \phi(v) \sigma_{\pi}^{2}(v)  \tag{A.3}\\
& =\sum_{v \in V} \phi(v) \mu_{\pi}(v)(1-2 t(v))+\sum_{v \in V} \phi(v) t(v) \tag{A.4}
\end{align*}
$$

Proof of Lemma A.1. Notice first that from (A.4) it follows that

$$
\begin{equation*}
\pi d(t)-\pi^{\prime} d(t)=\sum_{v \in V} \phi(v)\left(\mu_{\pi}(v)-\mu_{\pi^{\prime}}(v)\right)(1-2 t(v)) \tag{A.5}
\end{equation*}
$$

which implies that if $\mu_{\pi}(v)=\mu_{\pi^{\prime}}(v)$ for all $v \in V$, then $\pi d(t)=\pi^{\prime} d(t)$. This proves sufficiency.

To prove necessity, we proceed by induction. When $\pi d(t)=\pi^{\prime} d(t)$ for all $t \in \mathcal{T}$, from (A.5) we obtain

$$
\begin{align*}
0 & =\sum_{v \in V} \phi(v)\left(\mu_{\pi}(v)-\mu_{\pi^{\prime}}(v)\right)(1-2 t(v)) \\
& =-\sum_{v \in t} \phi(v)\left(\mu_{\pi}(v)-\mu_{\pi^{\prime}}(v)\right)+\sum_{v \notin t} \phi(v)\left(\mu_{\pi}(v)-\mu_{\pi^{\prime}}(v)\right) \tag{A.6}
\end{align*}
$$

for all $t \in \mathcal{T}$. Letting $t=\varnothing$, the empty tree, and $t=\{\lambda\}$ in (A.6), we obtain

$$
\begin{equation*}
0=\sum_{v \in V} \phi(v)\left(\mu_{\pi}(v)-\mu_{\pi^{\prime}}(v)\right) \tag{A.7}
\end{equation*}
$$

$$
\begin{equation*}
0=-\phi(\lambda)\left(\mu_{\pi}(\lambda)-\mu_{\pi^{\prime}}(\lambda)\right)+\sum_{v \neq \lambda} \phi(v)\left(\mu_{\pi}(v)-\mu_{\pi^{\prime}}(v)\right) \tag{A.8}
\end{equation*}
$$

Substracting (A.7) and (A.8) and using that $\phi(v)>0$, we get $\mu_{\pi}(\lambda)=\mu_{\pi^{\prime}}(\lambda)$.
Inductive step. Let $t \in \mathcal{T}$, and $h \in V \backslash\{t\}$ such that $t \cup\{h\} \in \mathcal{T}$. We show that if $\mu_{\pi}(v)=\mu_{\pi^{\prime}}(v)$ for all $v \in t$, then $\mu_{\pi}(h)=\mu_{\pi^{\prime}}(h)$. First, we obtain from (A.6)
(A.9) $0=\sum_{v \notin t} \phi(v)\left(\mu_{\pi}(v)-\mu_{\pi^{\prime}}(v)\right)$,
(A.10) $0=-\phi(h)\left(\mu_{\pi}(h)-\mu_{\pi^{\prime}}(h)\right)+\sum_{v \notin \nsubseteq \cup h\}} \phi(v)\left(\mu_{\pi}(v)-\mu_{\pi^{\prime}}(v)\right)$, and it follows that $\mu_{\pi}(h)=\mu_{\pi^{\prime}}(h)$.

It is easy to prove the following lemma
Lemma A.3. If $T$ is a random tree satisfying the Markov hypotheses with tree-shift $f$, then, given $T(v)=1$, the variables $T(w): w \in f^{-1}(v)$ are independent. Furthermore, if $v=f(w)$,

$$
\begin{equation*}
\mathbb{P}(T(w)=1 \mid T(v)=1)=\frac{\mu_{\pi}(w)}{\mu_{\pi}(v)} \tag{A.11}
\end{equation*}
$$

Proof of Proposition A.2. First, notice that

$$
\begin{equation*}
\pi(\varnothing)=1-\mu_{\pi}(\lambda) \tag{A.12}
\end{equation*}
$$

From Lemma A.3,

$$
\begin{equation*}
\pi(\{\lambda\})=\mu_{\pi}(\lambda) \prod_{h \in f^{-1}(\lambda)}\left(1-\frac{\mu_{\pi}(h)}{\mu_{\pi}(\lambda)}\right) . \tag{A.13}
\end{equation*}
$$

Let $t \in \mathcal{T} \backslash\{\varnothing, V\}$ and $h$ be a node such that $h \notin t$ and $v=f(h) \in t$. We shall show that

$$
\begin{equation*}
\pi(t \cup\{h\})=\pi(t) \frac{\mu_{\pi}(h)}{\mu_{\pi}(v)-\mu_{\pi}(h)} . \tag{A.14}
\end{equation*}
$$

First, we have

$$
\begin{align*}
\pi(t \cup\{h\}) & =\mathbb{P}\left(T_{t}=1, T(h)=1, T_{(t \cup\{h\})^{c}}=0\right) \\
& =\mathbb{P}\left(T(h)=1 \mid T_{t}=1, T_{(t \cup\{h\})^{c}}=0\right) \mathbb{P}\left(T_{t}=1, T_{(t \cup\{h\})^{c}}=0\right)  \tag{A.15}\\
& =\mathbb{P}(T(h)=1 \mid T(v)=1) \mathbb{P}\left(T_{t}=1, T_{(t \cup\{h\})^{c}}=0\right) .
\end{align*}
$$

On the other hand,

$$
\begin{align*}
\pi(t) & =\mathbb{P}\left(T_{t}=1, T(h)=0, T_{(t \cup\{h\})^{c}}=0\right) \\
& =\mathbb{P}(T(h)=0 \mid T(v)=1) \mathbb{P}\left(T_{t}=1, T_{(t \cup\{h\})^{c}}=0\right)  \tag{A.16}\\
& =(1-\mathbb{P}(T(h)=1 \mid T(v)=1)) \mathbb{P}\left(T_{t}=1, T_{(t \cup\{h\})^{c}}=0\right) .
\end{align*}
$$

From (A.15) and (A.16) it follows that

$$
\begin{equation*}
\pi(t \cup\{h\})=\pi(t) \frac{\mathbb{P}(T(h)=1 \mid T(f(h))=1)}{1-\mathbb{P}(T(h)=1 \mid T(f(h))=1)} \tag{A.17}
\end{equation*}
$$

This shows (A.14). Our main statement follows now by induction from (A.13) and (A.14), noticing that any finite tree may be constructed from $\{\lambda\}$ in this way.
A.2. Redefining the minimization problem. In this subsection we prove Propositions 4.1 and 4.2. Call $\Delta(v)=\left(\bar{t}(v)-\overline{t^{\prime}}(v)\right)$ so that

$$
\begin{equation*}
\mathcal{L}(t)=\sum_{v \in V} \phi(v) \Delta(v) t(v) \tag{A.18}
\end{equation*}
$$

Recall the space of configurations $\mathcal{y}=\{0,1\}^{V}$. Trees minimizing $\mathcal{L}$ will also minimize $\mathcal{L}+\beta \mathcal{P}$ for all positive $\beta$ : If $t \in \mathcal{T} \cap \arg \min _{y \in y} \mathcal{L}(y)$, then $t \in \arg \min _{y \in y}(\mathcal{L}(y)+\beta \mathcal{P}(y))$ for all $\beta>0$. On the other hand, if $\beta$ is big enough, we expect the configurations minimizing $\mathcal{L}+\beta \mathcal{P}$ to be trees. Since on the set of trees the form $\mathcal{P}$ vanishes, the minimizing trees should also minimize $\mathcal{L}$. This is proven in the following lemma.

Lemma A.4. If $\beta>\sum_{v \in V} \phi(v)$, then

$$
\begin{equation*}
\arg \min _{t \in \mathcal{T}} \mathcal{L}(t)=\arg \min _{y \in \mathcal{Y}}(\mathscr{L}(y)+\beta \mathcal{P}(y)) . \tag{A.19}
\end{equation*}
$$

Proof. Observe that minimizing configurations in $\mathcal{y}$ satisfy
(A.20) $y^{\prime} \in \arg \min _{y \in \mathcal{Y}} \mathcal{L}(y) \quad$ if and only if $\quad y^{\prime}(v)= \begin{cases}0, & \text { if } \Delta(v)>0, \\ 1, & \text { if } \Delta(v)<0 .\end{cases}$

Let $\mathcal{L}_{\text {min }}$ and $\mathcal{L}_{\text {max }}$ be the values of the minimum and maximum of $\mathcal{L}$ over $\mathcal{y}$, that is,

$$
\begin{equation*}
\mathcal{L}_{\min }=\sum_{v: \Delta(v)<0} \phi(v) \Delta(v), \quad \mathcal{L}_{\max }=\sum_{v: \Delta(v)>0} \phi(v) \Delta(v) . \tag{A.21}
\end{equation*}
$$

Notice that

$$
\begin{equation*}
\mathcal{L}_{\max }-\mathcal{L}_{\min }=\sum_{v \in V} \phi(v)|\Delta(v)| \leq \sum_{v \in V} \phi(v) . \tag{A.22}
\end{equation*}
$$

For all $y \in \mathcal{y}$ it holds

$$
\begin{equation*}
\mathcal{L}(y)+\beta \mathcal{P}(y) \geq \mathcal{L}_{\min }+\beta \sum_{v: y(v)=0} \sum_{a \in A} y(a v)(1-y(v)) . \tag{A.23}
\end{equation*}
$$

If $y$ is not a tree, there exists $v \in V$ and $a \in A$ such that $y(v)=0$ and $y(a v)=1$, hence,

$$
\begin{equation*}
\mathscr{L}(y)+\beta \mathcal{P}(y) \geq \mathscr{L}_{\min }+\beta>\mathcal{L}_{\min }+\sum_{v \in V} \phi(v) \geq \mathscr{L}_{\max } \tag{A.24}
\end{equation*}
$$

by (A.22). On the other hand, $\mathcal{L}(t)+\beta \mathcal{P}(t)=\mathscr{L}(t) \leq \mathcal{L}_{\max }$ for any tree $t$. Hence, for these values of $\beta$, if $y$ is not a tree, then $\mathcal{L}(y)+\beta \mathcal{P}(y)>\max _{t \in \mathcal{T}}(\mathcal{L}(t)+$ $\beta \mathcal{P}(t)$ ) and the result follows.

Proof of Proposition 4.1. It follows from (14), (15) and the above lemma applied to $\mathcal{L}$ and $-\mathcal{L}$.

Proof of Proposition 4.2. The proof follows from a simple algebra. More generally, the Hamiltonian $\mathscr{H}: \mathcal{y} \rightarrow \mathbb{R}$ given by

$$
\begin{equation*}
\mathscr{H}(y)=\sum_{v \in V} H^{v}(y(v))+\sum_{v, w \in V} H^{v, w}(y(v), y(w)) \tag{A.25}
\end{equation*}
$$

is said regular if the quadratic terms satisfy

$$
\begin{equation*}
H^{v, w}(0,0)+H^{v, w}(1,1) \leq H^{v, w}(0,1)+H^{v, w}(1,0) . \tag{A.26}
\end{equation*}
$$

Theorem 4.1 of Kolmogorov and Zabih (2004) says that regular Hamiltonians are graph representable, meaning that it is possible to associate capacities to the graph ( $\widetilde{V}, \widetilde{E}$ ) defined in (18) in such a way that

$$
\begin{equation*}
C(y)=k+\mathscr{H}(y), \quad y \in \mathcal{y}, \tag{A.27}
\end{equation*}
$$

where $C(y)$ is the capacity of the cut defined by $y$ [see (21)] and $k$ is a constant.
The Hamiltonian $\mathcal{L}+\beta \mathcal{P}$ is regular because $H^{v, a v}(t(v), t(a v))=\beta t(a v)(1-$ $t(v)$ ) satisfy (A.26). The graph (18) with capacities (19) is the Kolmogorov and Zabih representation of the Hamiltonian $\mathcal{L}+\beta \mathcal{P}$.
A.3. Simulation results. In this subsection we perform the test for the twosample problem (10) using samples $\mathbf{t}$ and $\mathbf{t}^{\prime}$ with distribution $\pi$ and $\pi^{\prime}$ on $\mathcal{T}$, trees with maximal depth $L=8$. We compute the BFFS statistic $W$ given in (12) using the approach of Section 4. We use the distance (8) with $\phi(v)=\theta^{\operatorname{gen}(v)}$ for various values of $\theta$. Since in this case we know $\pi$ and $\pi^{\prime}$, we estimate the quantiles directly by Monte Carlo simulation, as follows:

## Quantile

1. Generate two samples of size $n$ both from $\frac{1}{2} \pi+\frac{1}{2} \pi^{\prime}$, a fair mixture of the laws. Label them sample 1 and sample 2. Compute the test statistic $W$ using the samples.
2. Repeat the above procedure a fixed number of times $N$.
3. Order the computed statistics values increasingly and define the quantile $q(1-$ $\alpha)$ as the statistic in place $(1-\alpha) N$.
4. Calculate the quantiles for several values of $\alpha$ and sample sizes $n$.

Power

1. Generate sample 1 from $\pi$, sample 2 from $\pi^{\prime}$ and compute $W$ using them.

TABLE 1
Model 1. Power of the tests with $p=0.5$ and $p^{\prime}=0.6,0.7,0.8$, sample size $n=31,51,125$

| $\alpha / p^{\prime}$ | $n=31$ |  |  | $n=51$ |  |  | $n=125$ |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | 0.6 | 0.7 | 0.8 | 0.6 | 0.7 | 0.8 | 0.6 | 0.7 | 0.8 |
| 0.01 | 0.101 | 0.504 | 0.923 | 0.122 | 0.744 | 0.999 | 0.466 | 0.994 | 0.9999 |
| 0.05 | 0.242 | 0.708 | 0.986 | 0.313 | 0.889 | 0.9999 | 0.661 | 0.999 | 0.9999 |
| 0.1 | 0.325 | 0.806 | 0.998 | 0.426 | 0.935 | 0.9999 | 0.743 | 0.9999 | 0.9999 |

2. Compare the obtained value against the quantile, and reject the null hypothesis with level $\alpha$ if $W>q(1-\alpha)$.
3. Repeat the last two steps a fixed number of times and compute the percentage of rejections for each value of $\alpha$ as a measure of the power of the test.

Model 1: Binomial. Let $\pi$ be the law of a Galton-Watson process with offspring distribution $\operatorname{Binomial}(2, p)$ and $\pi^{\prime}$ is the same with parameter $p^{\prime}$. We use $p=0.5$ and $p^{\prime}=0.6,0.7,0.8$. Table 1 shows the percentage of rejection over 1000 tests of level $\alpha=0.10,0.05,0.01$ for sample sizes $n=31,51,125$. The results show the consistency of the BFFS test for alternatives with any value of $p \neq 0.5$. For this simple model, small sample sizes are enough to get high power.

Model 2: Mixture of Binomials. Let $\pi$ be the law of a GaltonWatson process with offspring distribution a mixture of Binomials. Independently at each node with probability $q$ use a $\operatorname{Binomial}\left(2, p_{1}\right)$, otherwise a $\operatorname{Binomial}\left(2, p_{2}\right)$. For $\pi^{\prime}$ we use $q^{\prime}, p_{1}^{\prime}$ and $p_{2}^{\prime}$. But we take $q=q^{\prime}=0.5$ in the examples.

MODEL 2.1. Take $p_{1}=0.45, p_{2}=0.5$ and $p_{1}^{\prime}=0.1, p_{2}^{\prime}=0.85$. Figure 2 shows histograms of the test statistic values obtained in 1000 iterations, for sample sizes 31 (left) and 131 (right) respectively. We have plotted the test statistic under null hypothesis on red and the one under the alternative on blue. The distance parameter was $\theta=0.35$. The expected mean at each node is the same because $p_{1}+p_{2}=p_{1}^{\prime}+p_{2}^{\prime}$; the distributions under the null and alternative hypothesis are close to each other but the BFFS test needs only a moderate sample size to give high power to the test. In Table 2 we evaluate the power of the test for three different values of the test level ( $\alpha$, as $0.01,0.05$ and 0.1 ), in the same way as in Table 1. The left plot on Figure 3 shows 5 curves of 1000 $p$-values each, for sample sizes $N=31,51,71,101,131$, and the distance parameter $\theta=0.35$. The right one shows the same results when the distance parameter is changed to $\theta=0.49$. Increasing the parameter $\theta$ decreases the power of the test.


FIG. 2. Model 2.1. Histogram of the test statistic under null hypothesis (black plot) and alternative (white plot), with sample sizes 31 and 131.

Model 2.2. The parameters $p_{1}=0.3, p_{2}=0.65, p_{1}^{\prime}=0.45$ and $p_{2}^{\prime}=0.5$ give a different scenario, since the distributions of the populations are very close. Figure 4 shows slow changes in the empirical distribution as the sample size grows. At the left is the histogram of the 1000 values of the test statistic under the null (red plot) and alternative hypothesis (blue plot) for sample size $N=50$. For the right histogram the sample size is $N=500$. In Figure 5 we consider larger sample sizes for Model 2.2. We fix $\theta=0.35$ and plot the $p$-values computed over 1000 replications for sample sizes $N=50,250$ and 500 (left plot) and the percentage of rejection as a function of the sample size, each curve computed with a different $\alpha$ level. Besides sample fluctuations, the percentage of rejection increases with sample size, but at a quite slow rate.

TABLE 2
Model 2.1. Power of the tests of level $\alpha$ for 1000 replications, with $\alpha=0.05,0.01,0.1$.
Sample sizes are 31, 51, 71, 101, 131. Parameter $\theta=0.35$

| $\boldsymbol{\alpha}$ | $\mathbf{3 1}$ | $\mathbf{5 1}$ | $\mathbf{7 1}$ | $\mathbf{1 0 1}$ | $\mathbf{1 3 1}$ |
| :--- | :---: | :---: | :---: | :---: | :---: |
| 0.01 | 0.045 | 0.108 | 0.325 | 0.594 | 0.819 |
| 0.05 | 0.288 | 0.363 | 0.747 | 0.887 | 0.973 |
| 0.1 | 0.438 | 0.642 | 0.857 | 0.976 | 0.990 |



Fig. 3. Model 2.1. p-values calculated 1000 times; sample sizes 31, 51, 71 and 131. Left plot: parameter $\theta=0.35$. Right plot: parameter $\theta=0.49$. Increasing the parameter $\theta$ decreases the power of the test.


Fig. 4. Model 2.2. Histogram of the test statistic under null hypothesis (black plot) and alternative (white plot); left plot size $=50$ and right plot size $=500$. Parameter $\theta=0.35$.


FIG. 5. Model 2.2. Left plot: p-values computed 1000 times, each curve related to a different sample size 50, 250 and 500. Right plot: percentage of rejection as a function of the sample size, each curve computed with a different $\alpha$ level, 0.01 in dotted line, 0.05 in dashed line and 0.1 in interpolated line. Parameter $\theta=0.35$.

Acknowledgments. We thank Antonio Galves for illuminating discussions about the discriminative power of context trees.

We thank the referees, an associate editor and the editor of the journal for their careful reading of the manuscript and several comments which helped to improve the paper.

## REFERENCES

Balding, D., Ferrari, P., Fraiman, R. and Sued, M. (2008). Limit theorems for sequences of random trees. Test (online). DOI: 10.1007/s11749-008-0092-z.
Banks, D. and Constantine, G. (1998). Metric models for random graphs. J. Classification 15 199-223. MR1665974
Bejerano, G. (2003). Automata learning and stochastic modeling for biosequence analysis. Ph.D. thesis, Hebrew Univ.
Bejerano, G. (2004). Algorithms for variable length Markov chain modeling. Bioinformatics 20 788-789.
Bejerano, G. and Yona, G. (2001). Variations on probabilistic suffix trees: Statistical modeling and prediction of protein families. Bioinformatics 17 23-43.
BÜhlmann, P. and Wyner, A. J. (1999). Variable length Markov chains. Ann. Statist. 27 480-513. MR1714720
Csiszár, I. and Talata, Z. (2006). Context tree estimation for not necessarily finite memory processes, via BIC and MDL. IEEE Trans. Inform. Theory 52 1007-1016. MR2238067

Finn, R. D., Mistry, J., Schuster-Böckler, B., Griffiths-Jones, S., Hollich, V., Lassmann, T., Moxon, S., Marshall, M., Khanna, A., Durbin, R., Eddy, S. R., Sonnhammer, E. L. L. and Bateman, A. (2006). Pfam: Clans, web tools and services. Nu cleic Acids Res. 34 D247-D51.
Galves, A., Galves, C., Garcia, N. and Peixoto, C. (2004). Correlates of rhythm in written texts of Brazilian and European Portuguese. Preprint. Available as Technical Report 08/08, IMECC/UNICAMP.
Galves, A., Maume-Deschamps, V. and Schmitt, B. (2008). Exponential inequalities for VMLC empirical trees. ESAIM Probab. Stat. 12 219-229. MR2374639
Greig, D., Porteous, B. and Seheult, A. (1989). Exact maximum a posteriori estimation for binary images. J. Roy. Statist. Soc. Ser. B 51 271-279.
Kolmogorov, V. and ZABIH, R. (2004). What energy functions can be minimized via graphs cuts? IEEE Trans. Pattern Analysis and Machine Intelligence 26 147-159.
Leonardi, F., Matioli, S. R., Armelin, H. A. and Galves, A. (2007). Detecting phylogenetic relations out from sparse context trees. Available at ArXiv:math/0804.4279.
Manly, B. F. J. (2007). Randomization, Bootstrap and Monte Carlo Methods in Biology, 3rd ed. Chapman \& Hall/CRC, New York.
Rissanen, J. (1983). A universal data compression system. IEEE Trans. Inform. Theory 29 656664. MR0730903

Ron, D., Singer, Y. and Tishby, N. (1996). The power of amnesia: Learning probabilistic automata with variable memory length. Machine Learning 25 117-149.
Stein, L. (2001). Genome annotation: From sequence to biology. Nat. Rev. Genet. 2 493-505.
Wang, H., and Marron, J. S. (2007). Object oriented data analysis: Sets of trees. Ann. Statist. 35 1849-1873. MR2363955
Wong, K. M., Suchard, M. A. and Huelsenbeck, J. P. (2008). Alignment uncertainty and genomic analysis. Science 319 473-476. MR2381044
J. R. BUSCH
S. P. Grynberg

Departamento de Matemáticas
Facultad de ingeniería
Universidad de Buenos Aires
Argentina
E-MAIL: jbusch@fi.uba.ar sebgryn@fi.uba.ar
A. G. Flesia

Centro de investigación y Estudios
de Matemática de Córdoba, CONICET
AND
FaMAF-UNC, Ing. Medina Allende S/N
Ciudad Universitaria
CP 5000, CÓRDOBA
Argentina
E-MAIL: flesia@mate.uncor.edu
P. A. Ferrari
F. Leonardi

Instituto de Matemática e Estatística Universidade de São Paulo
Brazil
E-MAIL: pablo@ime.usp.br florencia@usp.br
R. Fraiman

Departamento de Matemática y Ciencias
Universidad de San Andrés
Buenos Aires
Argentina
AND
CMAT, UNIVERSIDAD DE LA REPÚBLICA
URUGUAY
E-MAIL: fraiman@udesa.edu.ar


[^0]:    Received May 2008; revised October 2008.
    ${ }^{1}$ Supported in part by FAPESP, CNPq, Instituto do Milênio for the Global Advance of Mathematics, CAPES-Secyt agreement and PROSUL.
    ${ }^{2}$ Supported in part by PICT 2005-31659 and PID Secyt 69/08.
    ${ }^{3}$ Supported by a FAPESP fellowship, proc. 06/56980-0.
    Key words and phrases. Protein classification, hypothesis testing, random trees, variable length Markov chains.

