

The Distribution of Word Matches Between Markovian Sequences with Periodic Boundary Conditions

CONRAD J. BURDEN,¹ PAUL LEOPARDI,¹ and SYLVAIN FORÊT²

ABSTRACT

Word match counts have traditionally been proposed as an alignment-free measure of similarity for biological sequences. The D_2 statistic, which simply counts the number of exact word matches between two sequences, is a useful test bed for developing rigorous mathematical results, which can then be extended to more biologically useful measures. The distributional properties of the D_2 statistic under the null hypothesis of identically and independently distributed letters have been studied extensively, but no comprehensive study of the D_2 distribution for biologically more realistic higher-order Markovian sequences exists. Here we derive exact formulas for the mean and variance of the D_2 statistic for Markovian sequences of any order, and demonstrate through Monte Carlo simulations that the entire distribution is accurately characterized by a Pólya-Aeppli distribution for sequence lengths of biological interest. The approach is novel in that Markovian dependency is defined for sequences with periodic boundary conditions, and this enables exact analytic formulas for the mean and variance to be derived. We also carry out a preliminary comparison between the approximate D_2 distribution computed with the theoretical mean and variance under a Markovian hypothesis and an empirical D_2 distribution from the human genome.

Key words: Markov chains, sequence analysis, statistical models.

1. INTRODUCTION

THE D_2 STATISTIC is defined as the number of short word matches of a given pre-specified length k between two sequences of letters from a finite alphabet \mathcal{A} . This statistic was first analyzed in the precise form studied below by Lippert et al. (2002). It was motivated by more general statistics based on word counts proposed by Blaisdell (1986), and by a statistic defined as a sum over word lengths of weighted inner products of word counts, known as d^2 (Torney et al., 1990; Hide et al., 1994). Such statistics have been proposed as a measure of similarity between biological sequences in cases where the more commonly used alignment methods may not be appropriate. A review of word-based alignment-free sequence comparison measures in existence at or about the time of the Lippert et al. article [including angle metrics (Stuart et al., 2002a, b), which bear considerable similarity to D_2] can be found in Vinga and Almeida (2003).

¹Mathematical Sciences Institute, Australian National University, Canberra, ACT, Australia.

²Research School of Biology, Australian National University, Canberra, ACT, Australia.

In subsequent developments, a number of variants of the D_2 statistic have been studied and analyzed. A shortcoming of the D_2 statistic, first noted by Lippert et al. (2002), is that the signal of biological sequence similarity one is trying to detect, namely, simultaneous over-representation of certain words in both sequences, is masked by the natural variability of word counts in each of the two sequences. This is most likely to be a problem for longer sequences, although perhaps not for sequences of short to moderate length (Burden et al., 2012a). To address this problem, Reinert et al. (2009) introduced centered and standardized statistics, which were demonstrated to have higher power to detect sequence similarity (Wan et al., 2010). Other variations on the D_2 statistic include allowing word matches up to a certain number of mismatches (Burden et al., 2008) for detecting regulatory modules (Forêt et al., 2006, 2009a; Göke et al., 2012), and the introduction of weighting factors to acknowledge chemically similar amino acids when studying protein sequences (Jing et al., 2011; Burden et al., 2012b).

The distributional properties of the D_2 statistic under the null hypothesis of sequences composed of independently and identically distributed (i.i.d.) letters have been studied extensively. Rigorous results for limiting asymptotic distributions have been derived for D_2 by Lippert et al. (2002) and Kantorovitz et al. (2006) and for D_2 with mismatches by Burden et al. (2008). Exact analytic formulas exist for the mean (Waterman, 1995) and variance (Kantorovitz et al., 2006; Forêt et al., 2009b) of D_2 , and of the weighted (Jing et al., 2011) and centered (Burden et al., 2012a) versions of D_2 . Accurate approximations to distribution of D_2 and its variants in terms of gamma and Pólya-Aeppli (or compound Poisson) distributions have been demonstrated via Monte Carlo simulations by Forêt et al. (2009a,b), Jing et al. (2011), and Burden et al. (2012a), allowing for fast and practical calculations of approximate p values under the i.i.d. null hypothesis.

However, analysis of the k -mer spectra of the genomes of several species by Chor et al. (2009a) provides strong evidence that genomic sequences are more appropriately modeled as having a Markovian dependence, possibly up to fifth order. In the current work, we extend previous exact analytic results for the mean, variance, and an empirical distribution of D_2 for i.i.d. sequences to the case of Markovian sequences. A previous study of this problem, with some approximations, has been carried out by Kantorovitz et al. (2007) in the process of developing a method for detecting regulatory modules in genomic sequences.

The current study differs in that we consider sequences with periodic boundary conditions (PBCs), for which we introduce a new definition of Markovian sequences. For i.i.d. sequences, Forêt et al. (2009b) have found imposition of PBCs to be an approximation that works well for biologically realistic sequences. In practice, PBCs are imposed on the D_2 calculation for finite-length biological sequences simply by sewing the ends together and including in the word count words that overlap the join. The motivation for the restriction to periodic sequences is that it simplifies calculations of the mean and variance, enabling exact analytic formulas that are readily computable on a laptop computer to double precision accuracy for arbitrary sequence lengths. They also enable accurate practical approximations to the D_2 distribution under the null hypothesis of Markovian sequences for biologically realistic parameter values. The approximation does not model boundary effects, but it does capture accurately the more important effect of nonindependence of overlapping words (see terms V_1 to V_4 in the variance formula, Section 3.3).

The layout of the article is as follows: In Section 2, we define Markovian sequences of arbitrary order with PBCs in terms of an algorithm for generating such sequences. In Section 3, we define the D_2 statistic and derive exact analytic formulas for its mean and variance for Markovian sequences. In Section 4, the accuracy of the mean and variance formulas is checked numerically, and hypothesized asymptotic distributions are demonstrated to provide accurate representations of the complete D_2 distribution. These distributions are compared with empirical distributions of D_2 from the human genome. Conclusions are drawn in Section 5. Technical details of the derivation of $\text{Var}(D_2)$ are given in the Appendix, and computer codes for evaluating the mean and variance are given in the Supplementary Material (Supplementary Material is available online at www.liebertonline.com/cmb).

2. DEFINITIONS

Consider a sequence $\mathbf{x} = x_1, x_2, \dots$ of letters from an alphabet \mathcal{A} of size d . We say that \mathbf{x} has *periodic boundary conditions* (PBCs) and is of length m if $x_{i+m} = x_i$ for all $i = 1, 2, \dots$.

A sequence $\mathbf{X} = X_1, X_2 \dots$ of random letters has a θ -th order Markovian dependence if

$$\text{Prob}(X_{i+\theta} = b | (X_i, \dots, X_{i+\theta-1}) = (a_1, \dots, a_\theta)) = M(a_1, \dots, a_\theta; b), \quad (1)$$

for a specified $d^\theta \times d$ matrix M satisfying

$$0 \leq M(a_1, \dots, a_\theta; b) \leq 1; \quad \sum_{b \in \mathcal{A}} M(a_1, \dots, a_\theta; b) = 1, \quad (2)$$

for all $a_1, \dots, a_\theta, b \in \mathcal{A}$. As a shorthand notation, we will write a string of length θ with an arrow above:

$$\vec{x} = (x_1, \dots, x_\theta), \quad (3)$$

and write any substring of \mathbf{X} of length θ in a similar fashion, labeled by the index of the first element:

$$\vec{X}_i = (X_i, \dots, X_{i+\theta-1}), \quad (4)$$

Thus, Equation (1) is written more compactly as

$$\text{Prob}(X_{i+\theta} = b | \vec{X}_i = \vec{a}) = M(\vec{a}; b). \quad (5)$$

Following the notation of Reinert et al. (2005), define a $d^\theta \times d^\theta$ square matrix \mathbb{M} as

$$\mathbb{M}(\vec{a}, \vec{b}) = \begin{cases} M(\vec{a}; b_\theta) & \text{if } (a_2, \dots, a_\theta) = (b_1, \dots, b_{\theta-1}), \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

Then the Markovian dependency can be written as a first-order Markovian dependency as

$$\text{Prob}(\vec{X}_{i+1} = \vec{b} | \vec{X}_i = \vec{a}) = \mathbb{M}(\vec{a}, \vec{b}). \quad (7)$$

2.1. Markovian sequences with PBCs

Given an order θ transition matrix M , we first attempt to define a periodic random sequence $\mathbf{X} = X_1, X_2 \dots, X_n$ of length n via the following algorithm:

- Step 0:** Choose a probability distribution on the set of strings of length θ : $\text{Prob}(\vec{X}_1 = \vec{x}) = \pi(\vec{x})$, where $0 \leq \pi(\vec{x}) \leq 1$ and $\sum_{\vec{x} \in \mathcal{A}^\theta} \pi(\vec{x}) = 1$.
- Step 1:** Generate $\vec{X}_1 = X_1, \dots, X_\theta$ from this distribution.
- Step 2:** Generate $X_{\theta+1}, \dots, X_{\theta+n}$ using Equation (5).
- Step 3:** If $\vec{X}_{n+1} = \vec{X}_1$, accept the sequence $\mathbf{X} = X_1, X_2 \dots, X_n$, otherwise repeat from Step 1 until an accepted sequence is obtained.

Clearly this algorithm entails that

$$\text{Prob}(\mathbf{X} = \mathbf{x}) = \frac{\pi(\vec{x}_1) \mathbb{M}(\vec{x}_1, \vec{x}_2) \mathbb{M}(\vec{x}_2, \vec{x}_3) \dots \mathbb{M}(\vec{x}_n, \vec{x}_1)}{\sum_{\vec{u}_1, \dots, \vec{u}_n \in \mathcal{A}^\theta} \pi(\vec{u}_1) \mathbb{M}(\vec{u}_1, \vec{u}_2) \mathbb{M}(\vec{u}_2, \vec{u}_3) \dots \mathbb{M}(\vec{u}_n, \vec{u}_1)}, \quad (8)$$

where \mathbb{M} is the transition matrix of the equivalent-first order Markov chain defined by Equation (6). The idea behind PBCs is that there should be no privileged position along the sequence from which to begin numbering. Thus, we further impose a condition that the sequence should have no privileged starting point, that is, for each $i = 1, \dots, n$,

$$\text{Prob}(\mathbf{X} = x_{i+1}x_{i+2} \dots x_nx_1 \dots x_i) = \text{Prob}(\mathbf{X} = \mathbf{x}). \quad (9)$$

Equations (8) and (9) imply that $\pi(\vec{x}_{i+1}) = \pi(\vec{x}_1)$ for each i and for every sequence $\mathbf{x} \in \mathcal{A}^n$, which can only happen if

$$\pi(\vec{x}) = \frac{1}{d^\theta} \quad \forall \vec{x} \in \mathcal{A}^\theta. \quad (10)$$

This leads to the following definition:

Definition Given a transition matrix M of order θ , a random Markovian sequence with PBCs of length n is one generated by the algorithm of Section 2.1 with the initial distribution π in Step 0 equal to the uniform distribution Equation (10).

It follows from Equation (8) that for a random Markovian sequence \mathbf{X} of length n , the probability of the configuration $\mathbf{x} = (x_1 \dots, x_m)$ occurring is

$$\text{Prob}(\mathbf{X} = \mathbf{x}) = \frac{\mathbb{M}(\vec{x}_1, \vec{x}_2) \mathbb{M}(\vec{x}_2, \vec{x}_3) \dots \mathbb{M}(\vec{x}_m, \vec{x}_1)}{\text{tr}(\mathbb{M}^m)}. \quad (11)$$

The distribution Equation (11) has also been proposed by Percus and Percus (2006), who made an extensive study of the probability distribution of words on periodic sequences, which they refer to as rings. Our approach is novel in that it gives an algorithm that can be implemented in practice to generate an ensemble of such sequences.

3. THE D_2 STATISTIC

3.1. Definition of D_2

Definition Given two random sequences \mathbf{X} and \mathbf{Y} with PBCs of length m and n , respectively, the D_2 statistic is defined as the number of k -word matches, including overlaps, between \mathbf{X} and \mathbf{Y} :

$$D_2 = \sum_{i=1}^m \sum_{j=1}^n I_{ij}, \quad (12)$$

where

$$I_{ij} = \begin{cases} 1 & \text{if } (X_i, \dots, X_{i+k-1}) = (Y_j, \dots, Y_{j+k-1}), \\ 0 & \text{otherwise} \end{cases} \quad (13)$$

is the word match indicator random variable for words length k positioned at site i in sequence \mathbf{X} and site j in sequence \mathbf{Y} .

Two Markovian sequences \mathbf{X} and \mathbf{Y} of order θ generated by the $d^\theta \times d$ matrix M define a random variable $D_2(k, M)$. By Equation (7), an equivalent specification of this situation is a pair of first-order Markovian sequences \mathbb{X} and \mathbb{Y} consisting of letters of an alphabet of size d^θ generated by the square matrix \mathbb{M} defined by Equation (6). The sparse structure of \mathbb{M} ensures that the set of possible sequence pairs (\mathbb{X}, \mathbb{Y}) is in one-to-one correspondence with the set of possible sequence pairs (\mathbf{X}, \mathbf{Y}) , and furthermore, for $k \geq \theta$, a word match of length k between \mathbf{X} and \mathbf{Y} is equivalent to a word match of length $k - \theta + 1$ between \mathbb{X} and \mathbb{Y} . It follows that the distributional properties of D_2 for Markovian sequences can be determined in terms of the properties of D_2 for an equivalent first-order system:

$$D_2(k, M) \equiv D_2(k - \theta + 1, \mathbb{M}), \quad k \geq \theta. \quad (14)$$

3.2. D_2 mean for arbitrary θ

Below we derive an exact formula for $E(D_2)$ for arbitrary-order Markovian sequences. In principle, the mean for any $k \geq \theta$ case can be derived in terms of an equivalent $\theta = 1$ case. However, here we give an *ab initio* proof for any θ , noting that, for $k \geq \theta$, the result is consistent with Equation (14).

Define the Hadamard product $\mathbb{A} \circ \mathbb{B}$ of two matrices \mathbb{A} and \mathbb{B} as the matrix whose (α, β) -th element is

$$(\mathbb{A} \circ \mathbb{B})_{\alpha\beta} = \mathbb{A}_{\alpha\beta} \mathbb{B}_{\alpha\beta}. \quad (15)$$

The mean of D_2 is

$$E(D_2(k, M)) = \begin{cases} \frac{mn}{\text{tr}(\mathbb{M}^m) \text{tr}(\mathbb{M}^n)} \text{tr} [(\mathbb{M}^{m-k+\theta} \circ \mathbb{M}^{n-k+\theta})(\mathbb{M} \circ \mathbb{M})^{k-\theta}] & \text{if } k \geq \theta, \\ \frac{mn}{\text{tr}(\mathbb{M}^m) \text{tr}(\mathbb{M}^n)} \sum_{u, v \in \mathcal{A}^{\theta-k}} \sum_{w \in \mathcal{A}^k} \mathbb{M}^m((wu), (wu)) \mathbb{M}^n((wv), (wv)) & \text{if } k < \theta, \end{cases} \quad (16)$$

where \mathbb{M} is defined by Equation (6), (wu) means the θ -tuple $(w_1 \dots w_k u_1 \dots u_{\theta-k})$, and similarly for (wv) .

Proof: We have that

$$E(D_2) = \sum_{i=1}^m \sum_{j=1}^n E(I_{ij}) = \sum_{i=1}^m \sum_{j=1}^n \text{Prob}(I_{ij} = 1), \quad (17)$$

where

$$\text{Prob}(I_{ij} = 1) = \sum_{w \in \mathcal{A}^k} \text{Prob}(X_i \dots X_{i+k-1} = w) \text{Prob}(Y_j \dots Y_{j+k-1} = w). \quad (18)$$

To calculate $\text{Prob}(X_i \dots X_{i+k-1} = w)$, we must consider separately the cases $k \geq \theta$ and $k < \theta$.

Consider first the case where $k \geq \theta$. The required probability is calculated by summing Equation (11) over all sequences \mathbf{x} subject to the restriction that $(x_i \dots x_{i+k-1}) = w$. The definition of the matrix \mathbb{M} , Equation (6), ensures that it is sufficient to restrict only those θ -tuples \vec{x}_i located within the word w , because contributions to the sum from any partially overlapping θ -tuples will be zero unless the overlap letters match those of w (see Fig. 1a). Thus

$$\text{Prob}(X_i \dots X_{i+k-1} = w) = \frac{\mathbb{M}^{m-k+\theta}(\vec{w}_{k-\theta+1}, \vec{w}_1) \mathbb{M}(\vec{w}_1, \vec{w}_2) \dots \mathbb{M}(\vec{w}_{k-\theta}, \vec{w}_{k-\theta+1})}{\text{tr}(\mathbb{M}^m)}, \quad (19)$$

where the θ -tuples $\vec{x}_1, \dots, \vec{x}_{i-1}, \vec{x}_{i+k-\theta+1}, \dots, \vec{x}_m$ have been summed over. Similarly we have

$$\text{Prob}(Y_j \dots Y_{j+k-1} = w) = \frac{\mathbb{M}^{n-k+\theta}(\vec{w}_{k-\theta+1}, \vec{w}_1) \mathbb{M}(\vec{w}_1, \vec{w}_2) \dots \mathbb{M}(\vec{w}_{k-\theta}, \vec{w}_{k-\theta+1})}{\text{tr}(\mathbb{M}^n)}. \quad (20)$$

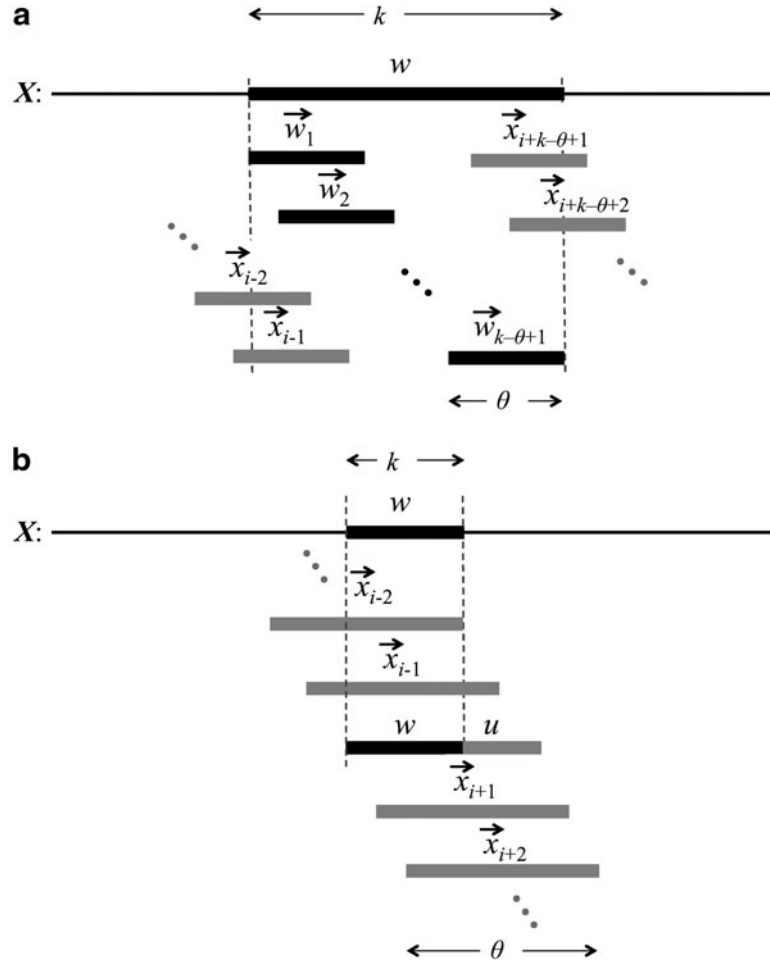


FIG. 1. Covering of the sequence \mathbf{X} with θ -mers for the calculation of $\text{Prob}(X_i \dots X_{i+k-1} = w)$ (a) in the case where $k \geq \theta$, and (b) in the case where $k < \theta$.

The definition Equation (6) of the matrix \mathbb{M} ensures that the sum over the k -word w in Equation (18) is equivalent to a sum over a set of independent θ -tuples $\vec{w}_1, \dots, \vec{w}_{k-\theta+1}$. Thus, substituting Equations (19) and (20) into Equation (18) gives

$$\text{Prob}(I_{ij} = 1) = \frac{\text{tr}[(\mathbb{M}^{m-k+\theta} \circ \mathbb{M}^{n-k+\theta})(\mathbb{M} \circ \mathbb{M})^{k-\theta}]}{\text{tr}(\mathbb{M}^m) \text{tr}(\mathbb{M}^n)}. \quad (21)$$

Equation (17) then gives the required result for the case $k \geq \theta$.

For the case $k < \theta$, the $\text{Prob}(X_i \dots X_{i+k-1} = w)$ is again calculated by summing Equation (11) over all sequences \mathbf{x} such that $(x_i \dots x_{i+k-1}) = w$. In this case, it is sufficient to restrict any one of the θ -tuples overlapping w to equal w on the overlap, and the structure of \mathbb{M} will ensure that only terms in which the other overlapping θ -tuples match w will contribute to the sum. Accordingly set $\vec{x}_i = (w_1 \dots w_k u_1 \dots u_{\theta-k})$, where the $u_1 \dots u_{\theta-k}$ are not fixed (see Fig. 1b). Then

$$\text{Prob}(X_i \dots X_{i+k-1} = w) = \frac{1}{\text{tr}(\mathbb{M}^m)} \sum_{u \in \mathcal{A}^{\theta-k}} \mathbb{M}^m((wu), (wu)), \quad (22)$$

and similarly

$$\text{Prob}(Y_j \dots Y_{j+k-1} = w) = \frac{1}{\text{tr}(\mathbb{M}^n)} \sum_{v \in \mathcal{A}^{\theta-k}} \mathbb{M}^n((wv), (wv)). \quad (23)$$

Substituting these two probabilities into Equations (18) and (17) gives the required result. ■

3.3. D_2 variance for $k \geq \theta$

For $k \geq \theta$, Equation (14) ensures that any $\theta > 1$ case can be reduced to an equivalent $\theta = 1$ case via the relation

$$\text{Var}(D_2(k, M)) = \text{Var}(D_2(k - \theta + 1, \mathbb{M})), \quad k \geq \theta, \quad (24)$$

where \mathbb{M} is a square first-order Markov matrix. Even for $\theta = 1$, the exact variance of D_2 for Markovian sequences with PBCs requires an extensive calculation. Here we give a summary of the $\theta = 1$ result, and leave the technical details of the derivation to the Appendix. The case $k < \theta$ remains intractable.

For the remainder of this section, we take M to be a square $d \times d$ first-order Markov matrix. We have

$$\text{Var}(D_2) = E(D_2^2) - E(D_2)^2. \quad (25)$$

The second term can be calculated from Equation (16). The first term is a sum of contributions obtained from Equation (12) by partitioning a sum over words beginning at positions i and i' in sequence \mathbf{X} and beginning at j and j' in sequence \mathbf{Y} ,

$$\begin{aligned} E(D_2^2) &= \sum_{i, i'=1}^m \sum_{j, j'=1}^n E(I_{ij} I_{i'j'}) \\ &= \sum_{i, i'=1}^m \sum_{j, j'=1}^n \text{Prob}(I_{ij} = 1, I_{i'j'} = 1) \\ &= V_0 + V_1 + V_2 + V_3 + V_4. \end{aligned} \quad (26)$$

The partitioning reflects the degree of overlap between words in each of the two sequences, and is illustrated in Figure 2. We assume $m, n \geq 2k$, which will almost certainly be the case in any biological application.

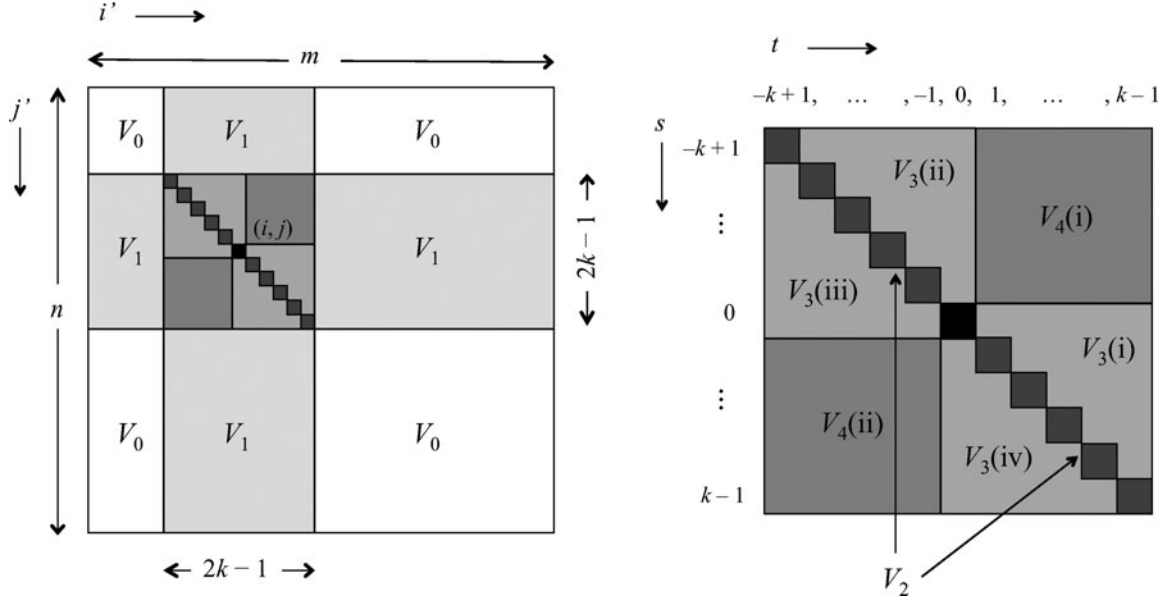


FIG. 2. Contributions to $\text{Var}(D_2)$ via the sum in Equation (26). The left-hand diagram shows the (i', j') -plane for a fixed value of (i, j) , shown as the black square. The right-hand diagram is an expanded view of the ‘‘accordion’’ region $-k + 1 \leq s, t \leq k - 1$, where $t = i' - i$ and $s = j' - j$ up to PBCs [see Equations (41) and (42)].

We will write a Hadamard product of q factors, $M \circ \dots \circ M$, using the shorthand notation $M^{\circ q}$. With this notation, the contributions to $E(D_2^2)$ are:

$$V_0 = \frac{mn}{\text{tr}(M^m)\text{tr}(M^n)} \times \sum_{r=0}^{m-2k} \sum_{s=0}^{n-2k} \text{tr}[(M^{r+1} \circ M^{s+1})(M \circ M)^{k-1}(M^{m-2k-r+1} \circ M^{n-2k-s+1})(M \circ M)^{k-1}], \quad (27)$$

$$V_1 = \frac{mn}{\text{tr}(M^m)\text{tr}(M^n)} \times \left\{ \sum_{s=0}^{n-2k} [\text{tr}\{[(M \circ M \circ M)^{k-1} \circ (M^{s+1})^T](M^{m-k+1} \circ M^{n-2k-s+1})\} + 2 \sum_{r=1}^{k-1} \text{tr}\{(M \circ M)^r [(M \circ M \circ M)^{k-r-1} \circ (M^{s+1})^T] \times (M \circ M)^r (M^{m-k-r+1} \circ M^{n-2k-s+1})\}] + \text{the same with } m \text{ and } n \text{ interchanged} \right\}, \quad (28)$$

$$V_2 = \frac{mn}{\text{tr}(M^m)\text{tr}(M^n)} \times \left\{ \text{tr}[(M^{m-k+1} \circ M^{n-k+1})(M \circ M)^{k-1}] + 2 \sum_{t=1}^{k-1} \text{tr}[(M^{m-k-t+1} \circ M^{n-k-t+1})(M \circ M)^{k+t-1}] \right\}, \quad (29)$$

$$V_3 = \frac{2mn}{\text{tr}(M^m)\text{tr}(M^n)} \sum_{t=1}^{k-1} \sum_{s=0}^{t-1} \text{tr}[(M \circ M)^s Q (M \circ M)^s \times (M^{m-k-t+1} \circ M^{n-k-s+1} + M^{n-k-t+1} \circ M^{m-k-s+1})], \quad (30)$$

where

$$\nu = \left\lfloor \frac{k-s}{t-s} \right\rfloor, \quad \rho = (k-s) \bmod (t-s), \quad (31)$$

and

$$Q = \begin{cases} (M^{\circ(2\nu+3)})^{\rho-1} \circ [(M^{\circ(2\nu+1)})^{t-s-\rho+1}]^T & \text{if } \rho > 0, \\ (M^{\circ(2\nu+1)})^{t-s-1} \circ (M^{\circ(2\nu-1)})^T & \text{if } \rho = 0. \end{cases} \quad (32)$$

Finally,

$$V_4 = \frac{2mn}{\text{tr}(M^m)\text{tr}(M^n)} \sum_{r,t=1}^{k-1} \text{tr}U, \quad (33)$$

where

$$\nu = \left\lfloor \frac{k}{r+t} \right\rfloor, \quad \zeta = k \bmod (r+t), \quad (34)$$

and

$$U = \left\{ \begin{array}{ll} \left\{ (M^{\circ(2\nu+1)})^{t-1} \circ (M^{m-k-t+1})^T \right\} M^{\circ 2\nu} & \\ \quad \times \left\{ (M^{\circ(2\nu+1)})^{r-1} \circ (M^{n-k-r+1})^T \right\} M^{\circ 2\nu} & \text{if } \zeta = 0, \\ \left\{ (M^{\circ(2\nu+1)})^{r-\zeta+1} \circ M^{m-k-t+1} \right\} (M^{\circ(2\nu+2)})^{\zeta-1} & \\ \quad \times \left\{ (M^{\circ(2\nu+1)})^{t-\zeta+1} \circ M^{n-k-r+1} \right\} (M^{\circ(2\nu+2)})^{\zeta-1} & \text{if } 0 < \zeta \leq r, t, \\ \left\{ (M^{\circ(2\nu+3)})^{\zeta-r-1} \circ (M^{m-k-t+1})^T \right\} (M^{\circ(2\nu+2)})^r & \\ \quad \times \left\{ (M^{\circ(2\nu+1)})^{t-\zeta+1} \circ M^{n-k-r+1} \right\} (M^{\circ(2\nu+2)})^r & \text{if } r < \zeta \leq t, \\ \text{as above with } m \text{ and } n \text{ interchanged} & \\ \quad \text{and } r \text{ and } t \text{ interchanged} & \text{if } t < \zeta \leq r, \\ \left\{ (M^{\circ(2\nu+3)})^{\zeta-r-1} \circ (M^{m-k-t+1})^T \right\} (M^{\circ(2\nu+2)})^{t+r-\zeta+1} & \\ \quad \times \left\{ (M^{\circ(2\nu+3)})^{\zeta-t-1} \circ (M^{n-k-r+1})^T \right\} (M^{\circ(2\nu+2)})^{t+r-\zeta+1} & \text{if } r, t < \zeta. \end{array} \right. \quad (35)$$

A full derivation of these contributions is given in the Appendix.

3.4. Computational advantages of PBCs

One could, in principle, define D_2 more conventionally without imposing PBCs by considering standard Markov chains and stopping the sums in Equation (12) at $n - k + 1$ and $m - k + 1$, respectively. This is the approach used by Kantorovitz et al. (2007). However, the PBCs confer two computational advantages that allow the mean and variance to be calculated to higher orders and without further approximation.

Firstly, the “no privileged starting point” condition implies that the summands in Equation (17) for the mean and Equation (26) for the variance are independent of the word positions i and j , which reduces the sums to multiplicative factors of m and n , respectively. The variance summand in particular depends only on the relative word positions $i' - i$ and $j' - j$. Kantorovitz et al. (2007) deal with this by assuming the first word occurrence in each random sequence to have a stationary distribution. This amounts to neglecting end effects, which introduces roughly the same order of approximation as PBCs.

Secondly, and more importantly, calculation of the variance via the Kantorovitz et al. (2007) approach entails multiple sums over sets of all possible words up to length $2k - 1$, whereas the PBCs reduce these sums to traces of powers of matrices that are readily computed. In particular, the terms V_2 to V_4 above can be computed very rapidly, whereas Kantorovitz et al. (2007) suggest that the equivalent terms be omitted to save computation.

3.5. Differing Markov models

It may be necessary in some biological situations to consider a situation in which the sequences \mathbf{X} and \mathbf{Y} are generated by differing transition matrices, say M_1 and M_2 , respectively. In this case, the formula for the mean easily generalizes to

$$E(D_2(k, M_1, M_2)) = \begin{cases} \frac{mn}{\text{tr}(M_1^m)\text{tr}(M_2^n)} \text{tr}[(M_1^{m-k+\theta} \circ M_2^{n-k+\theta})(M_1 \circ M_2)^{k-\theta}] & \text{if } k \geq \theta, \\ \frac{mn}{\text{tr}(M_1^m)\text{tr}(M_2^n)} \sum_{u,v \in \mathcal{A}^{\theta-k}} \sum_{w \in \mathcal{A}^k} M_1^m((wu), (wu)) M_2^n((wv), (wv)) & \text{if } k < \theta. \end{cases} \quad (36)$$

A detailed formula for the variance in this case is beyond the scope of this article, but the key points of difference arising between this case and the case of a single common Markov matrix are clear. To summarize, the terms V_0 , V_1 , and V_2 generalize relatively easily, but the terms V_3 and V_4 require more attention. For instance, the symmetry between the subcases of V_3 shown in Figure 2 is broken, and thus the factor $M^{\circ(2v+1)}$ in Equation (32) becomes either $M_1^{\circ v} \circ M_2^{\circ(v+1)}$ or $M_1^{\circ(v+1)} \circ M_2^{\circ v}$, thus doubling the number of terms.

4. NUMERICAL RESULTS

4.1. Computer implementation of the mean and variance

In the Supplementary Material, we provide an R implementation (R Core Development Team, 2012) of $E(D_2(k, M))$ for arbitrary k and of $\text{Var}(D_2(k, M))$ for $k \geq \theta$ using the formulas derived above. The $k > \theta$ means and variances are calculated by reducing the problem to the equivalent $\theta = 1$ calculation with effective $d^\theta \times d^\theta$ Markov matrix M and effective word length $k - \theta + 1$ [see Equation (24)].

The computationally most expensive parts of the computation of $\text{Var}(D_2)$ are the sums over r and s occurring in Equation (27) and the first line of Equation (28). These sums are implemented efficiently for large sequence lengths m and n by storing powers of M out to convergence and by making use of the fact that the summand is essentially constant over parts of the domain of summation for which these matrix powers have converged. Although the programs are not yet fully optimized, they calculate $\text{Var}(D_2)$ in about 30 sec on a standard laptop computer for an alphabet of size $d = 4$, Markovian order $\theta = 3$, word lengths up to $k = 20$, and arbitrarily large sequence lengths m and n . The variance program slows for higher order Markov models as the size of M grows exponentially with θ . Considerable gains are possible for the case $k = \theta$, as the terms V_2 , V_3 , and V_4 in the equivalent $\theta = 1$ calculation are automatically zero, and double sum in the term V_0 can be computed more efficiently by using the identity

$$\begin{aligned} \sum_{r=0}^{m-2} \sum_{s=0}^{n-2} \text{tr}[(M^{r+1} \circ M^{s+1})(M^{m-r-1} \circ M^{n-s-1})] = \\ \sum_{i,j=1}^{d^\theta} \left(\sum_{r=0}^{m-2} (M^{r+1})_{ij} (M^{m-r-1})_{ji} \right) \left(\sum_{s=0}^{n-2} (M^{s+1})_{ij} (M^{n-s-1})_{ji} \right). \end{aligned} \quad (37)$$

Also included in the Supplementary Material is a test program that generates the complete distribution of D_2 for short sequences for a randomly chosen Markovian model created by choosing each matrix element from a uniform distribution on the interval $[0, 1]$ and then normalizing each row sum to 1. Using this program, we have confirmed the accuracy of the above mean and variance formulas to 13 significant figures for sequences up to length $m = n = 10$ for various values of the alphabet size d , Markov order θ , and word length k . Two examples of the exact D_2 distribution for short sequences are shown in Figure 3.

For the case of sequences composed of i.i.d. letters, certain rigorous results are known for the asymptotic distribution of D_2 as the sequence lengths $m, n \rightarrow \infty$. For $m = n$, it has been shown that the limiting distribution is normal in the regime $k < 1/2 \log_b n + \text{const.}$ (Burden et al., 2008) and Pólya-Aeppli in the regime $k > 2 \log_b n + \text{const.}$ (Lippert et al., 2002). Here $b = 1 / \sum_{a \in \mathcal{A}} p_a^2$ where p_a is the probability of occurrence of letter a . A Pólya-Aeppli random variable is the sum of a Poisson number of geometric random variables, and is therefore an example of a compound Poisson random variable. It often arises in the study of random word counts as a Poisson number of clumps of overlapping words, each clump containing a geometric number of k -words (Reinert and Schbath, 1998; Reinert et al., 2005). In earlier work on i.i.d. sequences, Burden et al. (2012a) have found in general that, for simulations of D_2 for moderate to long sequences, the gamma distribution provides a good interpolation between the normal and Pólya-Aeppli regimes. Although the asymptotic results for D_2 are not proved for Markovian sequences, it is a

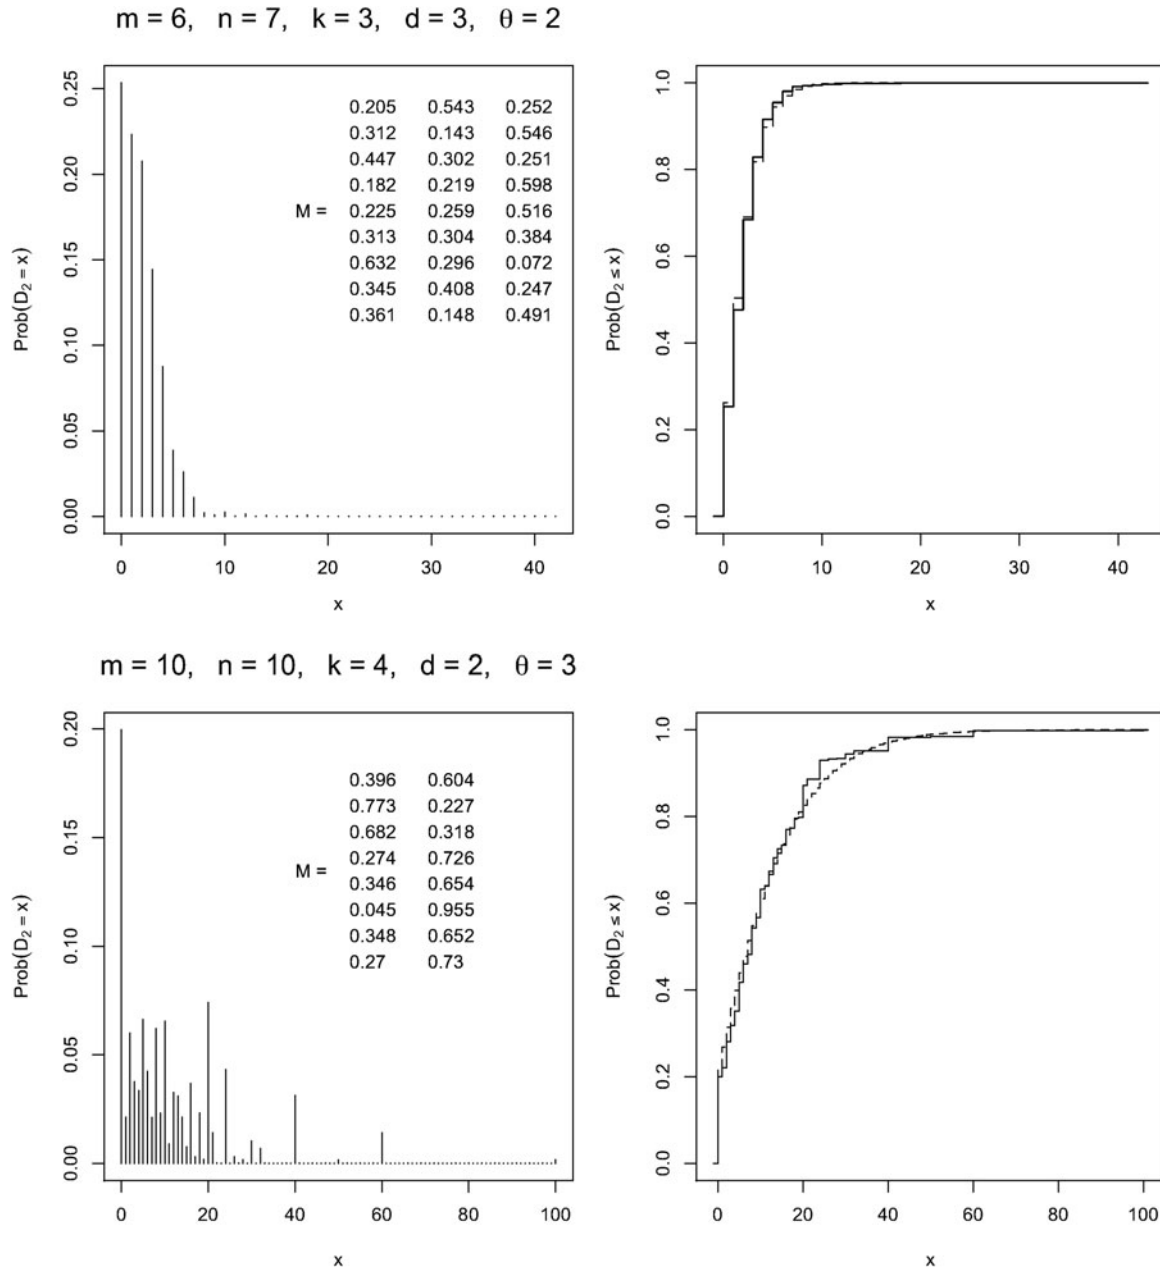


FIG. 3. The exact distribution of the D_2 statistic for short sequences of length m, n and words of length k from a Markov model of order θ and alphabet of size d . The Markov matrix M has been generated randomly in each case, and the exact distribution has been calculated by enumerating all d^{m+n} possible sequence pairs. Also shown (dashed curve) is the cumulative distribution of the Pólya-Aeppli distribution with mean and variance set to the theoretical values using the formulas of Section 3.

reasonable experiment to compare our numerical simulations with these distributions as they may potentially provide an accurate estimate of p values in biological applications.

One would not expect the asymptotic distributions to be an accurate fit to the exact distributions for the short sequences considered in Figure 3. Nevertheless, we have added the Pólya-Aeppli distribution function with the mean and variance adjusted to their theoretical values to the plots, and find it to be a surprisingly close fit. Disagreement arises in the tail of the distribution because, for combinatoric reasons, certain values of D_2 within the range 0 to mn do not occur, whereas the Pólya-Aeppli has support over the whole range (and also out to ∞ , albeit with very low probability).

4.2. Comparison with simulated distributions

For sequences of realistic biological length composed of the four-letter nucleotide alphabet, it is necessary to resort to Monte Carlo simulations to investigate the D_2 probability distribution.

We used a combination of R scripts and the SAFT program [Sequence Alignment-Free Tool, under development (Forêt, 2012)] to further verify the formulas for the mean and variance, and to compare the empirical distribution of the D_2 statistic with the conjectured asymptotic normal, Pólya-Aeppli, and gamma distributions. For this purpose, as well as using randomly generated Markov matrices, we used matrices obtained from DNA sequences occurring in nature. The Supplementary Material to Chor et al. (2009a) contains maximum likelihood estimates of Markov matrices for a number of species and for different regions within the human genome. As an example, we used the Markov matrices for human chromosome 1, with Markov orders 0, 1, 2, and 3 (Chor et al., 2009b). For each of these matrices, we used an R script that implements the algorithm of Section 2.1, using the built-in random number generator of R, via the function `sample.int()`, to generate 20,000 sequences of length 1,000, arranged as 10,000 pairs of cyclic sequences. The SAFT program calculated the D_2 statistic for each of these 10,000 pairs. We then used a second R script, based on the code in the Supplementary Material, to compare the mean and variance of the empirical distribution of the D_2 statistic with the theoretical values given by Equations (16) and (25) to (35), to compare the empirical cumulative distribution of the D_2 statistic with known distributions, and to plot results. Some simulations were also carried out for sequences of length 100 and 400.

As the purpose of these simulations is to verify the accuracy of the mean and variance formulas and to test the validity of certain functional approximations to the distribution function, the short to moderate sequence lengths are chosen to be in a range in which all terms in the variance formula are observed to make a noticeable contribution. The variance term V_0 is $O(m^2n^2)$ in the sequence lengths, V_1 is $O(mn(m+n))$, and the remaining terms are $O(mn)$, and so for longer sequences the term V_0 dominates. It happens that the sequence lengths in these simulations reflect typical sizes of *cis*-regulatory modules (Kantorovitz et al., 2007), but the theory will be applicable to biological sequences of any length.

Table 1 presents the results for the mean and variance for Markov orders 0 to 3. For the mean, the row labeled “Theoretical” is calculated from the corresponding Markov matrix using formula (16), the row labeled “Empirical” is estimated from the 10,000 values of D_2 obtained via SAFT, and the rows labeled “Lower 95%” and “Upper 95%” are obtained from the confidence interval returned by the R function `t.test()` that implements Student’s t test. For the variance σ^2 , the row labeled “Theoretical” is calculated from the corresponding Markov matrix using formulas (25) to (35), the row labeled “Empirical” is estimated from the 10,000 values of D_2 obtained via SAFT, and the rows labeled “Lower 95%” and “Upper 95%” are obtained via the χ^2 distribution, using the R quantile function `qchisq` and the inequality given by Snedecor and Cochran (1980; Section 5.10.2, p. 74),


$$(N-1)s^2/\chi_{0.025}^2 \leq \sigma^2 \leq (N-1)s^2/\chi_{0.975}^2,$$

TABLE 1. MEAN AND VARIANCE OF D_2 CALCULATED FROM THEORETICAL FORMULAS

		Order			
		0	1	2	3
Mean	Lower 95%	18.84	24.70	27.66	28.84
	Theoretical	18.92	24.73	27.79	28.97
	Empirical	18.95	24.83	27.80	29.00
	Upper 95%	19.07	24.96	27.95	29.15
Variance	Lower 95%	32.89	43.23	53.06	59.01
	Theoretical	33.24	44.69	55.56	60.53
	Empirical	33.81	44.44	54.54	60.65
	Upper 95%	34.77	45.70	56.09	62.37

Mean and variance of D_2 were calculated from the theoretical formulas derived in Section 3, and estimated from synthetically generated data (10,000 sequence pairs) for Markov models of order $\theta = 0, 1, 2$, and 3 using Markov matrices estimated from human chromosome 1. Word length $k = 8$, alphabet size $d = 4$, sequence lengths $m = n = 1,000$.

FIG. 4. Comparison of empirical cumulative distribution function for simulated D_2 using Markov matrices for human chromosome 1 from Chor et al. (2009a,b), for orders $\theta = 0, 1, 2,$ and 3 . 10,000 pairs per order, word length $k = 8$, alphabet size $d = 4$, sequence lengths $m = n = 1,000$.




where $N = 10,000$ in this case, and s^2 is the sample variance. In these and in a number of other simulations we have performed (data not shown), we find that in roughly the expected proportion of times the mean and variance calculated from the formulas of Section 3 lie within the 95% confidence intervals computed from the ensemble.

As a general rule, and as can be seen from Table 1, we observe that both the mean and variance of D_2 increase markedly as the Markov order increases for fixed word length k and sequence lengths m and n . The difference between the empirical cumulative distribution functions for the different Markov orders for the parameters of Table 1 is further illustrated in Figure 4.

We compared the empirical distribution of D_2 for each Markov order with conjectured asymptotic distributions based on the theoretical mean and variance calculated via Equations (16) and (25) to (35). For Markov order 3, this is illustrated by Figure 5. Here the cumulative gamma and normal distributions are plotted using the built-in R functions `pgamma()` and `pnorm()`, respectively, and the cumulative Pólya-Aeppli distribution is plotted using the function `pPolyaAeppli()` included in the Supplementary Materials. We observe that, for these parameter values, the three conjectured distributions do not differ greatly from one another, although the Pólya-Aeppli clearly gives the best fit, particularly in the important tail of the

FIG. 5. Comparison of Pólya-Aeppli, gamma, and normal cumulative distributions with empirical cumulative distribution function for simulated D_2 using Markov order 3 matrix for human chromosome 1 from Chor et al. (2009a,b). 10,000 pairs, word length $k = 8$, sequence lengths $m = n = 1,000$.



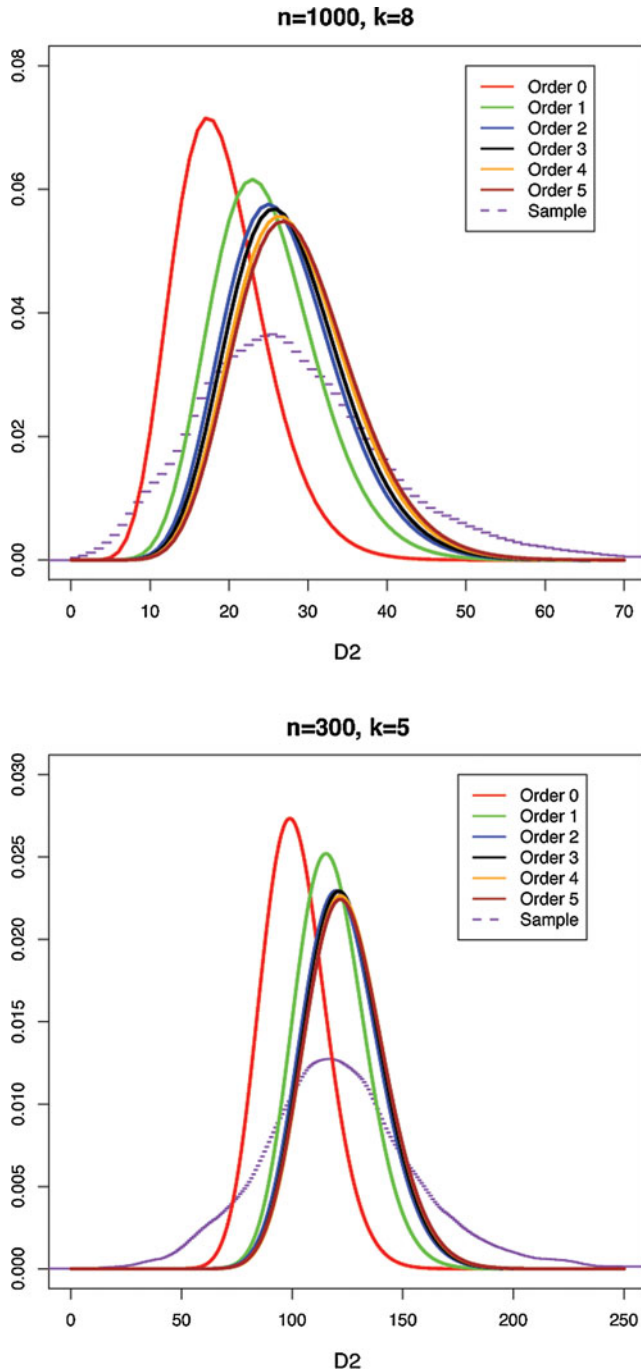


FIG. 6. Density of the empirical distribution of D_2 from human chromosome 1 sample data from Ensembl compared with gamma distributions with calculated mean and variance, based on Markov models of various orders θ . 10,000 sample pairs, word length $k = 8$, and sequences lengths $m = n = 1,000$ (*upper plot*), word length $k = 5$ and sequences lengths $m = n = 300$ (*lower plot*).

distribution relevant to estimating p values. This trend is also observed for the other Markov orders and sequence lengths simulated including the simulations in Figure 4. In general, the Pólya-Aeppli behavior that is expected to apply asymptotically for large sequence lengths is reached within the accuracy expected of our Monte Carlo simulations at sequence lengths of some hundreds of letters. For parameters leading to large values of $E(D_2)$, the continuous normal and gamma distributions are more readily computable, although slightly less accurate, than the Pólya-Aeppli, and of these two the gamma is invariably observed to give a better fit.

4.3. Comparison with chromosomal DNA

Ultimately, one hopes to use D_2 or similarly defined statistics as an alignment-free tool to assess the relatedness of biological sequences. To this end, it is helpful to know to what extent genomic sequences can

be modeled as Markovian sequences for the purpose of defining a null-hypothesis distribution for the D_2 statistic. With this in mind, we have performed some exploratory comparisons between the D_2 distributions obtained via simulating the Markov processes using maximum likelihood estimates of Markov matrices and the D_2 distribution obtained by sampling original DNA data, for example, the DNA sequence from human chromosome 1 from Ensembl (Wellcome Trust Sanger Institute and European Bioinformatics Institute, (2012)). For consistency with the range of parameters used in the simulations of the previous section, the comparisons were done for sequence lengths $m = n = 300$ and 1,000.

Figure 6 illustrates the comparison between D_2 distributions approximated by gamma distributions with exact means and variances calculated under Markovian hypotheses of various orders, and the empirical density of the D_2 distribution obtained from sampling human chromosome 1. The Markov transition matrices used for calculating the mean and variance at each order were estimated using maximum likelihood from the same subset of human chromosome 1 as that used for obtaining the empirical D_2 density (see below). The gamma representation was demonstrated to provide a very accurate approximation to the D_2 distribution (as estimated from Monte Carlo simulations) for $m = n = 1,000$ and Markov order $\theta = 0, 1, 2,$ and 3 in the previous Section (see Fig. 5). Here we also assume the gamma approximation to the D_2 distribution for θ and up to 5 to avoid further Monte Carlo simulations, as the computational demands of the algorithm of Section 2.1 for generating sequences with PBCs become prohibitive for higher Markov orders.

To obtain the empirical density, we took the soft-masked DNA sequence for human chromosome 1 from Ensembl, and took uniform random samples of subsequences of length 300 or 1,000, according to Knuth’s Algorithm S (Knuth, 1981, Section 3.4.2), but avoiding all ambiguous and masked regions. Ensembl’s masking removes repetitive regions including tandem repeats. This data source and procedure for estimating Markov transition matrices correspond to those described by Chor et al. (2009a), except that the Markov matrices have been estimated from Ensembl’s “soft-masked” sequences with the repeat regions (i.e., the lowercase letters) ignored, whereas Chor et al. include the repeat regions. We find that, as expected, including the repeat regions leads to a skewed empirical D_2 distribution with an extremely heavy right-hand tail corresponding to repetitive regions.

The sample mean and variance from the soft-masked DNA sequence, together with the theoretical values, are shown in Table 2. In general, agreement between the Markovian model and the empirical distribution improves as the Markovian order increases. For higher orders, the Markovian mean overshoots slightly. The Markovian variance, on the other hand, severely underestimates the empirical variance at any order. This is consistent with earlier observations by Csürös et al. (2007) that genomic word count distributions tend to have heavier tails than that predicted by Markovian models, or, to put it another way, certain k -mers are “under-” or “over-represented” within genomes.

Note also that the Markovian plots in Figure 6 suggest that $\theta = k$ may be in some sense a limiting case. Recall that the formula for the mean takes a different form for $\theta > k$ [see Equation (16)], and that the formula derived for the variance is only valid for $\theta \leq k$ and remains intractable for $\theta > k$. We suspect that this is related to the fact that, for sufficiently long sequences, θ -mer frequencies are determined by the stationary eigenvector of the Markov matrix, and that the statistics of k -mers for $k < \theta$ is implicit with the statistics of θ -mers.

TABLE 2. EMPIRICAL ESTIMATES OF MEAN AND VARIANCE OF D_2 FROM HUMAN CHROMOSOME 1

	<i>Theoretical values</i>						<i>Sample estimate</i>
	$\theta = 0$	1	2	3	4	5	
<i>m = n = 1,000, k = 8</i>							
Mean	19.08	24.74	26.90	27.55	28.30	28.74	27.66
Variance	33.58	44.62	50.99	52.19	54.37	56.01	181.1
Std. Dev.	5.795	6.680	7.141	7.224	7.373	7.484	13.46
<i>m = n = 300, k = 5</i>							
Mean	101.1	117.4	122.4	123.5	124.3	124.3	120.7
Variance	216.6	254.4	307.5	307.8	315.6	321.9	1,258.
Std. Dev.	14.71	15.95	17.53	17.55	17.77	17.94	35.47

Empirical estimates of the mean and variance of D_2 from human chromosome 1 sample data from Ensembl (right-hand column) were compared with the theoretical mean and variance based on Markov models of various orders using estimated Markov matrices for human chromosome 1. The variance for $k = \theta = 5$ was calculated by implementing Equation (37).

5. DISCUSSION

The primary purpose of this article is to demonstrate that it is possible to construct accurate representations of the distribution of the D_2 statistic under the null hypothesis of periodic Markovian sequences without the need to resort to computationally expensive Monte Carlo simulations or to asymptotic approximations valid only when $\log n \gg k$. Our method consists of deriving exact formulas for the mean and variance of D_2 that are readily computable for any sequence lengths, to which we fit functional forms based on asymptotic distributions typically observed for word count statistics. We have demonstrated that, for sequences of moderate length of up to only a few hundred letters, and for which $\log n \approx k$, the Pólya-Aeppli distribution with parameters determined by the exact formulas for the mean and variance developed herein accurately represents the true D_2 distribution for Markovian sequences of any order (see Fig. 5). For comparatively longer sequences with higher $E(D_2)$, for which evaluating the Pólya-Aeppli distribution may be slow, the gamma distribution provides an acceptable approximation that is more accurate than the normal distribution.

It is known that the D_2 statistic itself, if used directly as a measure of sequence similarity, may perform poorly as the signal of over-representation of the same words in the query and target sequences is masked by the natural variability of word counts in each of the two sequences (Lippert et al., 2002). Variations on the theme of the D_2 statistic, such as the weighted, centered statistic D_2^* studied by Reinert et al. (2009), have been developed to circumvent this problem. Burden et al. (2012a,b) have extended calculations of the exact mean and variance for i.i.d. sequences to weighted and centered versions of D_2 , and it is expected that the analogous calculation for Markovian sequences will be entirely feasible.

The secondary purpose of this article is a preliminary comparison of the approximate D_2 distribution computed with the theoretical mean and variance under a Markovian hypothesis with an empirical genomic D_2 distribution. As a test example, we have considered the empirical distribution of the D_2 statistic between randomly chosen segments of a single human chromosome, avoiding highly repetitive parts of the chromosome such as stretches of tandem repeats. In general, we find that the empirical distribution has much heavier tails than the D_2 distribution for a Markovian sequence of any order up to $\theta = 5$ (see Fig. 6). We interpret this as a signal that the chromosome, taken as a whole, contains a number of strongly over- and under-represented k -mers, relative to a Markovian sequence. Thus, one is tempted to conclude that a Markov model will tend to overestimate significance and give an inflated false-positive rate when attempting to detect relatedness of genomic sequences.

However, this test is preliminary, and takes no account of the structure of the genome. In particular, we have not restricted ourselves to non-protein-coding segments. As current opinion is that even the non-coding part of the human genome may be up to 80% functional (ENCODE Project Consortium, 2012), the possibility exists that the over- and under-represented words are restricted to segments of genome with specific, possibly yet unknown, functions. Thus, the potential exists, for instance, to use D_2 as an exploratory probe to detect functional regions within the noncoding part of the genome: Using a randomly generated Markovian probe sequence (a random probe of length $m = 10,000$, say, would contain almost all 6-mers), one could calculate D_2 between the probe and a moving window running along the genome. This exercise would expose whether, for instance, the genome consists of a sea of “null hypothesis” Markovian sequence containing islands of repeated motifs, or whether the genome is uniformly peppered with a particular set overexpressed words. The ability to easily calculate the null D_2 distribution as a function of sequence and word lengths enables the experiment to be performed readily at different resolutions. Furthermore, the property of D_2 that it is dominated by the natural variability in either of the two sequences being compared becomes an advantage. If a subset of words is over-represented within the moving window at a specific location in the genome, provided that subset contains some words also present in the probe sequence, its over-representation within the window will manifest as an extreme D_2 .

6. APPENDIX: CONTRIBUTIONS TO VAR (D_2)

We derive the contributions V_0 to V_4 to $Var(D_2)$ when $\theta = 1$ given in Section 3.3. These contributions are the partial sums $\sum_{i, i', j, j'} \text{Prob}(I_{ij} = 1, I_{i'j'} = 1)$ contributing to Equation (26) where, for given (i, j) , the indices (i', j')

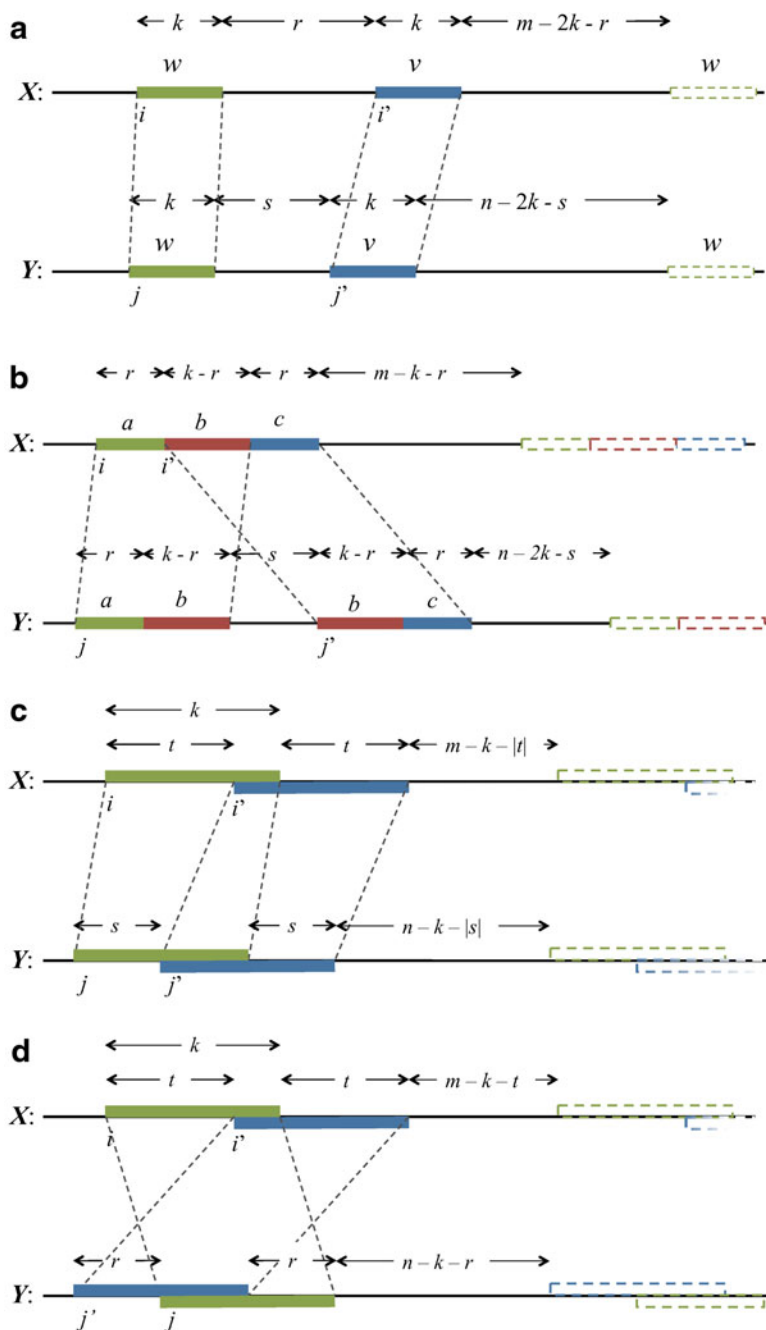


FIG. 7. Arrangements of word matches contributing to (a) non-overlapping words, V_0 , (b) crabgrass V_1 , (c) accordions V_2, V_3 , and (d) accordion V_4 . Images of words due to periodic boundary conditions are shown as a dashed outline.

range over the regions shown in Figure 2. The event “ $I_{ij} = 1, I_{i'j'} = 1$ ” means that the k -words beginning at sites i and i' in sequence \mathbf{X} match the k -words beginning at sites j and j' in sequence \mathbf{Y} , respectively.

Nonoverlapping words in both sequences: V_0

Taking into account the PBCs, these are the contributions from the cases for which both $k \leq |i' - i| \leq m - k$ and $k \leq |j' - j| \leq n - k$ occur simultaneously. Consider the situation

$$\begin{aligned} i' &= (i + k + r) \bmod m, & r &= 0, \dots, m - 2k, \\ j' &= (j + k + s) \bmod n, & s &= 0, \dots, n - 2k, \end{aligned}$$

shown in Figure 7a. As the two sequences are independent, applying Equation (11) gives

$$\begin{aligned}
\text{Prob}(I_{ij} = 1, I_{i'j'} = 1) &= \sum_{w, v \in \mathcal{A}^k} \text{Prob}(X_i \dots X_{i+k-1} = w) \text{Prob}(X'_i \dots X'_{i'+k-1} = v) \\
&\quad \times \text{Prob}(Y_j \dots Y_{j+k-1} = w) \text{Prob}(Y'_j \dots Y'_{j'+k-1} = v) \\
&= \frac{1}{\text{tr}(M^m) \text{tr}(M^n)} \times \sum_{w, v \in \mathcal{A}^k} (M^{m-2k-r+1})_{v_k w_1} M_{w_1 w_2} \dots M_{w_{k-1} w_k} (M^{r+1})_{w_k v_1} M_{v_1 v_2} \dots M_{v_{k-1} v_k} \\
&\quad \times (M^{n-2k-s+1})_{v_k w_1} M_{w_1 w_2} \dots M_{w_{k-1} w_k} (M^{s+1})_{w_k v_1} M_{v_1 v_2} \dots M_{v_{k-1} v_k} \\
&= \frac{\text{tr}[(M^{r+1} \circ M^{s+1})(M \circ M)^{k-1} (M^{m-2k-r+1} \circ M^{n-2k-s+1})(M \circ M)^{k-1}]}{\text{tr}(M^m) \text{tr}(M^n)}. \tag{38}
\end{aligned}$$

Summing over r and s , and including a factor of mn to account for the sum over i and j then gives Equation (27).

Overlaps in one sequence only: V_1

These are cases for which either $k \leq |j' - j| \leq n - k$ and $|i' - i| < k$ or $> m - k$ (overlaps in \mathbf{X} but not in \mathbf{Y}), or $k \leq |i' - i| \leq m - k$ and $|j' - j| < k$ or $> n - k$ (overlaps in \mathbf{Y} but not in \mathbf{X}). This region is referred to as the ‘‘crabgrass’’ by Waterman (1995). Figure 7b shows the case of overlaps in \mathbf{X} but not \mathbf{Y} , where we have set

$$\begin{aligned}
r &= \begin{cases} i' - i & \text{if } |i' - i| < k, \\ i' - i - m & \text{if } |i' - i| > m - k, \end{cases} \\
j' &= (j + k + s) \bmod n,
\end{aligned}$$

for $r = -k + 1, \dots, 0, \dots, k - 1$ and $s = 0, \dots, n - 2k$. We split the common word beginning at i and j into a piece a of length r and a piece b of length $k - r$, and split the common word beginning at i' and j' into the piece b and a piece c of length r .

Then

$$\begin{aligned}
V_1 &= mn \left\{ \sum_{s=0}^{n-2k} \sum_{r=-k+1}^{k-1} \sum_{a \in \mathcal{A}^r} \sum_{b \in \mathcal{A}^{k-r}} \sum_{c \in \mathcal{A}^r} \text{Prob}(\text{configuration in Fig. 7b}) \right. \\
&\quad \left. + \text{a similar sum with the roles of } \mathbf{X} \text{ and } \mathbf{Y} \text{ interchanged} \right\}, \tag{39}
\end{aligned}$$

where the sums over r and s arise from sums over i' and j' for fixed i and j , and the factor of mn arises from the outer sum over i and j . Using Equation (11),

$$\begin{aligned}
&\sum_{a \in \mathcal{A}^r} \sum_{b \in \mathcal{A}^{k-r}} \sum_{c \in \mathcal{A}^r} \text{Prob}(\text{configuration in Fig. 7(b)}) \\
&= \frac{1}{\text{tr}(M^m) \text{tr}(M^n)} \times \sum_{a, b, c} M_{a_1 a_2} \dots M_{a_{r-1} a_r} M_{a_r b_1} M_{b_1 b_2} \dots M_{b_{k-r-1} b_{k-r}} M_{b_{k-r} c_1} \\
&\quad \times M_{c_1 c_2} \dots M_{c_{r-1} c_r} (M^{m-k-r+1})_{c_r a_1} \times M_{a_1 a_2} \dots M_{a_{r-1} a_r} M_{a_r b_1} M_{b_1 b_2} \dots M_{b_{k-r-1} b_{k-r}} (M^{s+1})_{b_{k-r} b_1} \\
&\quad \times M_{b_1 b_2} \dots M_{b_{k-r-1} b_{k-r}} M_{b_{k-r} c_1} M_{c_1 c_2} \dots M_{c_{r-1} c_r} (M^{n-2k-s+1})_{c_r a_1} \\
&= \sum_{a_1, b_1, b_{k-r}, c_r} [(M \circ M)^r]_{a_1 b_1} [(M \circ M \circ M)^{k-r-1}]_{b_1 b_{k-r}} [(M \circ M)^r]_{b_{k-r} c_r} \\
&\quad \times (M^{s+1})_{b_{k-r} b_1} (M^{m-k-r+1})_{c_r a_1} (M^{n-2k-s+1})_{c_r a_1} \\
&= \text{tr}\{(M \circ M)^r [(M \circ M \circ M)^{k-r-1} \circ (M^{s+1})^T] \times (M \circ M)^r (M^{m-k-r+1} \circ M^{n-2k-s+1})\},
\end{aligned} \tag{40}$$

where the superscript T indicates the matrix transpose. Equations (39) and (40) combine to give the crabgrass contribution Equation (28).

Overlaps in both sequences

The set of configurations for which the words at positions i , i' , j , and j' overlap in both sequences simultaneously is referred to as the “accordion” by Waterman (1995). For convenience, we define the following overlap distances (illustrated in Fig. 7c):

$$t = \begin{cases} i' - i & \text{if } |i' - i| < k, \\ i' - i - m & \text{if } |i' - i| > m - k, \end{cases} \quad (41)$$

in sequence \mathbf{X} and

$$s = \begin{cases} j' - j & \text{if } |j' - j| < k, \\ j' - j - n & \text{if } |j' - j| > n - k, \end{cases} \quad (42)$$

in sequence \mathbf{Y} . These definitions ensure that $-k + 1 \leq t, s \leq k - 1$. The remaining three contributions are from the accordion.

Diagonal part of the accordion: V_2

This is the contribution from those cases with $s = t$, in which case Figure 7c becomes a match between the $(k + |t|)$ -letter word at position i in \mathbf{X} and the $(k + |t|)$ -letter word at position j in \mathbf{Y} . Noting that the probability of this match is independent of i and j , we have

$$V_2 = mn \sum_{t=-k+1}^{k-1} \text{Prob}((X_i \dots X_{i+k+|t|-1}) = (Y_j \dots Y_{j+k+|t|-1})), \quad (43)$$

where, by analogy with Equation (21),

$$\text{Prob}((X_i \dots X_{i+k+|t|-1}) = (Y_j \dots Y_{j+k+|t|-1})) = \frac{\text{tr}[(M^{m-k-t+1} \circ M^{n-k-t+1})(M \circ M)^{k+t-1}]}{\text{tr}(M^m)\text{tr}(M^n)}. \quad (44)$$

Combining Equations (43) and (44) gives Equation (29).

Off-diagonal part of the accordion: subcases contributing to V_3

The off-diagonal part of the accordion is divided into a number of subcases. Consider first the contribution from the four subcases making up the region V_3 in Figure 2:

- 3(i): $0 \leq s < t \leq k - 1$;
- 3(ii): $-k + 1 \leq s < t \leq 0$;
- 3(iii): $-k + 1 \leq t < s \leq 0$; and
- 3(iv): $0 \leq t < s \leq k - 1$.

By symmetry, each subcase makes an equivalent contribution to the variance. Subcase 3(i) is shown in Figure 8, and the required contribution takes the form

$$V_3 = 2mn \sum_{t=1}^{k-1} \sum_{s=0}^{t-1} \sum_{a, b \in \mathcal{A}^s} \sum_{c \in \mathcal{A}^\rho} \sum_{d \in \mathcal{A}^\sigma} [\text{Prob}(\text{configuration shown in Fig. 8}) \\ + \text{Prob}(\text{same configuration with } m \text{ and } n \text{ interchanged})]. \quad (45)$$

To calculate the probability of the configuration, the overlapping words have been divided into repeating independent elements. Elements a and b are the nonoverlapping parts of length s at either end of the words at j and j' in \mathbf{Y} . The nonoverlapping part of the words at i and i' in \mathbf{X} are segmented into elements (acd) and (dcb) shown in the upper part of Figure 8. The segment (cd) repeats an integer number v times within the overlapping part in sequence \mathbf{Y} , with a segment c of length ρ left over. We set the length of element d equal to σ . Thus




FIG. 8. Arrangements of word matches contributing to subcase 3(i) when $\rho = (k - s) \bmod (t - s) > 0$ (upper figure) and $\rho = 0$ (lower figure).

$$\nu = \left\lfloor \frac{k-s}{t-s} \right\rfloor, \quad \rho = (k-s) \bmod (t-s), \quad \text{and} \quad \sigma = t-s-\rho. \quad (46)$$

When $\rho = 0$ the element c does not occur (lower part of Fig. 8).

Using arguments similar to those for the crabgrass contribution, we have, for $\rho > 0$,

$$\begin{aligned} & \sum_{a, b \in \mathcal{A}^s} \sum_{c \in \mathcal{A}^\rho} \sum_{d \in \mathcal{A}^\sigma} \text{Prob}(\text{configuration shown in Fig. 8}) = \frac{1}{\text{tr}(M^m) \text{tr}(M^n)} \\ & \times \sum_{a_1, a_s, b_1, b_s, c_1, c_\rho, d_1, d_\sigma \in \mathcal{A}} [(M \circ M)^{s-1}]_{a_1 a_s} (M \circ M)_{a_s c_1} [(M^{\circ(2\nu+3)})^{\rho-1}]_{c_1 c_\rho} \\ & \times (M^{\circ(2\nu+1)})_{c_\rho d_1} [(M^{\circ(2\nu+1)})^{\sigma-1}]_{d_1 d_\sigma} (M^{\circ(2\nu+1)})_{d_\sigma c_1} \\ & \times (M \circ M)_{c_\rho b_1} [(M \circ M)^{s-1}]_{b_1 b_s} (M^{m-k-t+1} \circ M^{n-k-s+1})_{b_s a_1} \\ & = \frac{1}{\text{tr}(M^m) \text{tr}(M^n)} \text{tr} [(M \circ M)^s \{ (M^{\circ(2\nu+3)})^{\rho-1} \circ [(M^{\circ(2\nu+1)})^{\sigma+1}]^T \} \\ & \times (M \circ M)^s (M^{m-k-t+1} \circ M^{n-k-s+1})], \end{aligned} \quad (47)$$

whereas for $\rho = 0$ we have

$$\begin{aligned} & \sum_{a, b \in \mathcal{A}^s} \sum_{c \in \mathcal{A}^\rho} \sum_{d \in \mathcal{A}^\sigma} \text{Prob}(\text{configuration shown in Fig. 8}) = \frac{1}{\text{tr}(M^m) \text{tr}(M^n)} \\ & \times \sum_{a_1, a_s, b_1, b_s, d_1, d_\sigma \in \mathcal{A}} [(M \circ M)^{s-1}]_{a_1 a_s} (M \circ M)_{a_s d_1} [(M^{\circ(2\nu+1)})^{t-s-1}]_{d_1 d_\sigma} (M^{\circ(2\nu-1)})_{d_\sigma d_1} \\ & \times (M \circ M)_{d_\sigma b_1} [(M \circ M)^{s-1}]_{b_1 b_s} (M^{m-k-t+1} \circ M^{n-k-s+1})_{b_s a_1} \\ & = \frac{1}{\text{tr}(M^m) \text{tr}(M^n)} \text{tr} [(M \circ M)^s \{ (M^{\circ(2\nu+1)})^{t-s-1} \circ (M^{\circ(2\nu-1)})^T \} \\ & \times (M \circ M)^s (M^{m-k-t+1} \circ M^{n-k-s+1})]. \end{aligned} \quad (48)$$

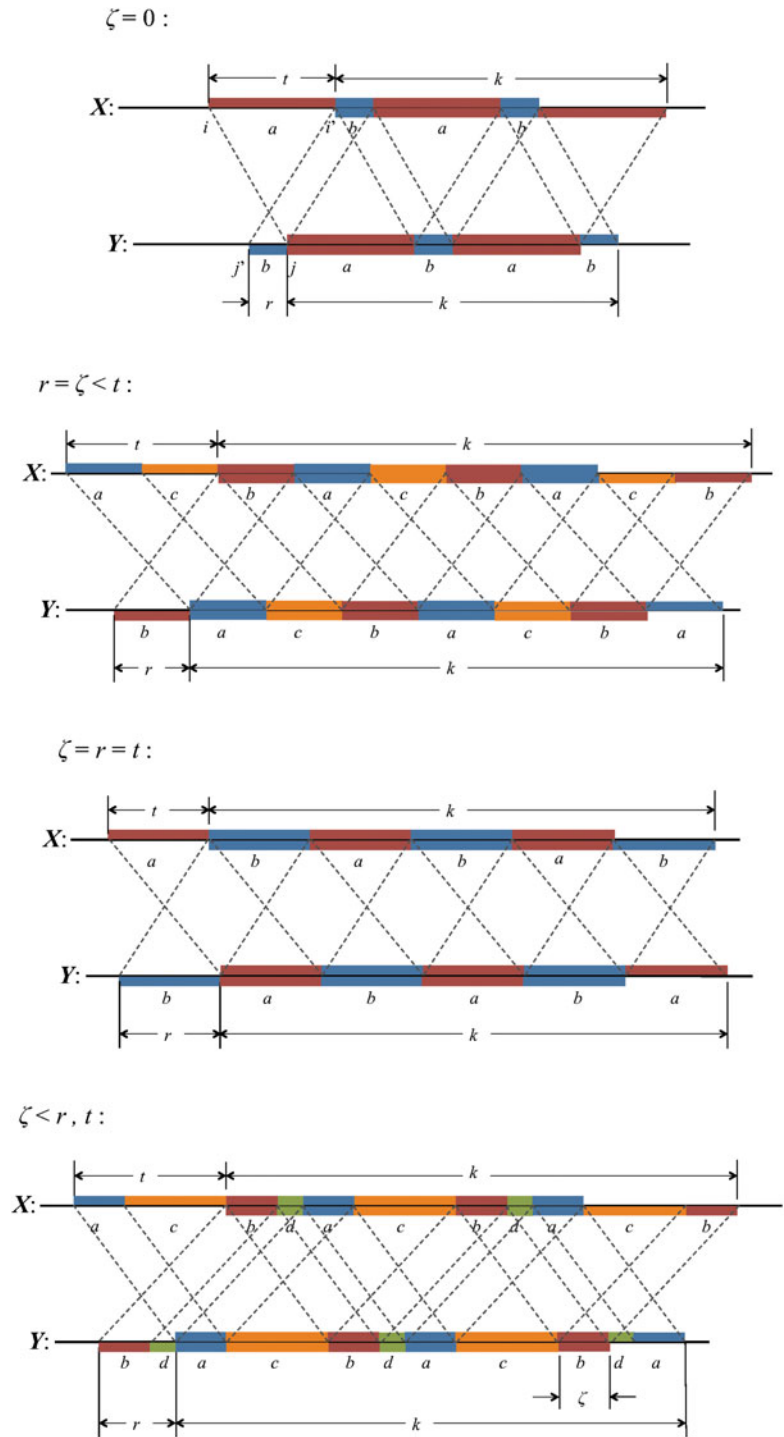


FIG. 9. Arrangements of word matches contributing to V_4 (first four cases).

Combining Equations (45) to (48) gives Equations (30) to (32).

Off-diagonal part of the accordion: subcases contributing to V_4

These are contributions from the subcases

- 4(i): $1 \leq t \leq k - 1, -k + 1 \leq s \leq -1$; and
- 4(ii): $1 \leq s \leq k - 1, -k + 1 \leq t \leq -1$,

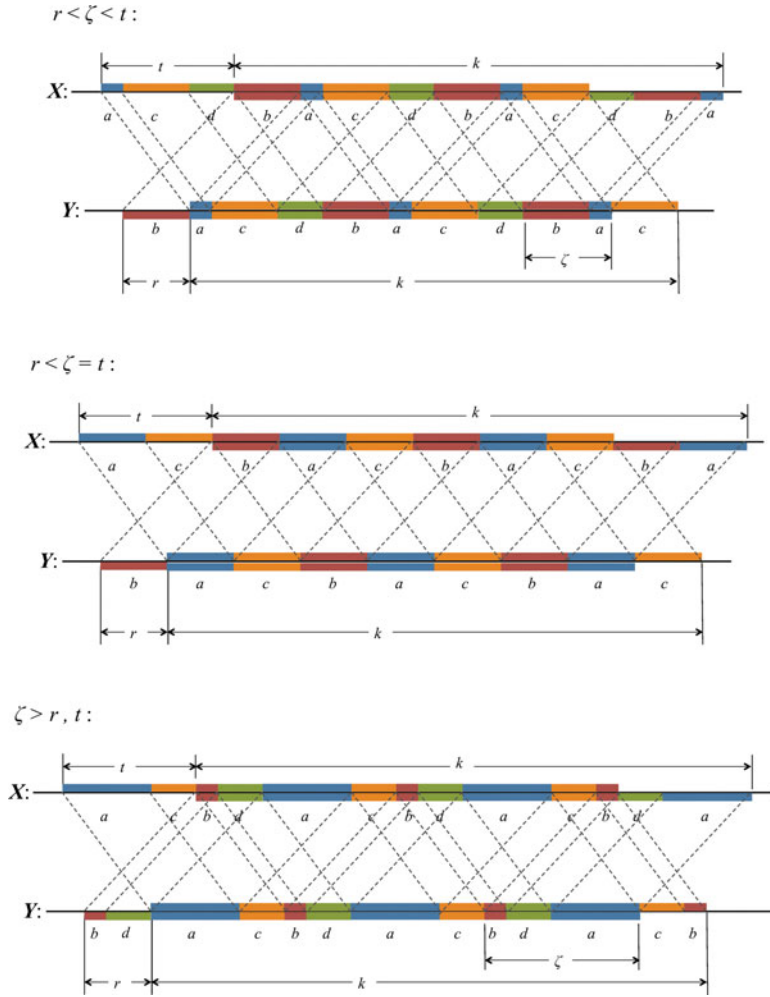


FIG. 10. Arrangements of word matches contributing to V_4 (final three cases).

labeled V_4 in Figure 2. In these cases, either t or s is negative. By symmetry, each of these two subcases makes an equivalent contribution to V_4 , so we consider subcase 4(i) and for convenience set $r = -s$ (see Fig. 7d). Then

$$V_4 = 2mn \sum_{r,t=1}^{k-1} \text{Prob}[(X_i \dots X_{i+k-1}) = (Y_j \dots Y_{j+k-1}), \\ (X_{i+t} \dots X_{i+t+k-1}) = (Y_{j-r} \dots Y_{j-r+k-1})], \quad (49)$$

where the factor mn arises from a sum over i and j , and we make use of the fact that for periodic Markovian sequences the summand is independent of i and j .

It is convenient to define

$$\nu = \left\lfloor \frac{k}{r+t} \right\rfloor, \quad \zeta = k \bmod (r+t). \quad (50)$$

Here ν is the integer number of times the complete repeat unit $(Y_j \dots Y_{j+r+t-1})$ fits inside the k -word $(Y_j \dots Y_{j+k-1})$, and ζ is the number of letters remaining (see Figs. 9 and 10). Calculation of the probability occurring in Equation (49) then proceeds in a similar fashion to that for V_3 by dividing the overlapping words into independent nonoverlapping elements. It turns out that the configuration of elements depends on the relationship between ζ , r , and t . The complete set of configurations is

enumerated in Figures 9 and 10, with the repeated elements labeled a , b , etc. The calculation is lengthy and repetitive but straightforward, and yields Equations (33) and (35) after recombining cases that give the same algebraic formula.

ACKNOWLEDGMENT

This work was supported in part by ARC Discovery grant DP120101422.

DISCLOSURE STATEMENT

The authors declare that no competing financial interests exist.

REFERENCES

- Blaisdell, B. 1986. A measure of the similarity of sets of sequences not requiring sequence alignment. *Proc. Natl. Acad. Sci. U.S.A.* 83, 5155–5159.
- Burden, C.J., Kantorovitz, M.R., and Wilson, S.R. 2008. Approximate word matches between two random sequences. *Ann. Appl. Probab.* 18, 1–21.
- Burden, C.J., Jing, J., Forêt, S., and Wilson, S.R. 2012. Application of k -word match statistics to the clustering of proteins with repeated domains. In Colubi, A., Fokianos, K., Kontoghiorghes, E., and González-Rodríguez, G., eds. *Proceedings of COMPSTAT 2012, 20th International Conference on Computational Statistics*. 131–142.
- Burden, C.J., Jing, J., and Wilson, S.R. 2012. Alignment-free sequence comparison for biologically realistic sequences of moderate length. *Stat. Appl. Genet. Mol. Biol.* 11, Article 3.
- Chor, B., Horn, D., Goldman, N., et al. 2009a. Genomic DNA k -mer spectra: models and modalities. *Genome Biol.* 10, R108.
- Chor, B., Horn, D., Goldman, N., et al. 2009b. k -mer analysis of multiple genomes. Available at www.ebi.ac.uk/goldman-srv/ChorEtAlSpectra/Spectra/HumanChromosomes/chr1/.
- Csűrös, M., Noé, L., and Kucherov, G. 2007. Reconsidering the significance of genomic word frequencies. *Trends Genet.* 23, 543–546.
- ENCODE Project Consortium, Bernstein, B.E., Birney, E., Dunham, I. et al. 2012. An integrated encyclopedia of DNA elements in the human genome. *Nature* 489, 57–74.
- Forêt, S. 2012. Sequence alignment-free tool. Available at <https://github.com/sylvainforet/saft>.
- Forêt, S., Kantorovitz, M.R., and Burden, C.J. 2006. Asymptotic behaviour and optimal word size for exact and approximate word matches between random sequences. *BMC Bioinformatics* 7 Suppl 5, S21.
- Forêt, S., Wilson, S.R., and Burden, C.J. 2009a. Characterizing the D_2 statistic: word matches in biological sequences. *Stat. Appl. Genet. Mol. Biol.* 8, Article 43.
- Forêt, S., Wilson, S.R., and Burden, C.J. 2009b. Empirical distribution of k -word matches in biological sequences. *Pattern Recognit.* 42, 539–548.
- Göke, J., Schulz, M., Lasserre, J., and Vingron, M. 2012. Estimation of pairwise sequence similarity of mammalian enhancers with word neighbourhood counts. *Bioinformatics* 28, 656–663.
- Hide, W., Burke, J., and Davison, D.B. 1994. Biological evaluation of d_2 , an algorithm for high-performance sequence comparison. *J. Comput. Biol.* 1, 199–215.
- Jing, J., Wilson, S.R., and Burden, C.J. 2011. Weighted k -word matches: a sequence comparison tool for proteins. *ANZIAM J.* 52 (CTAC2010), 172–189.
- Kantorovitz, M.R., Booth, H.S., Burden, C.J., and Wilson, S.R. 2006. Asymptotic behavior of k -word matches between two uniformly distributed sequences. *J. Appl. Probab.* 44, 788–805.
- Kantorovitz, M.R., Robinson, G.E., and Sinha, S. 2007. A statistical method for alignment-free comparison of regulatory sequences. *Bioinformatics* 23, i249–i255.
- Knuth, D.E. 1981. *The Art of Computer Programming*, Volume 2: Seminumerical Algorithms, 2nd ed. Addison-Wesley, Reading, MA.
- Lippert, R.A., Huang, H., and Waterman, M.S. 2002. Distributional regimes for the number of k -word matches between two random sequences. *Proc. Natl. Acad. Sci. U.S.A.* 99, 13980–13989.
- Percus, J., and Percus, O. 2006. The statistics of words on rings. *Commun. Pure Applied Math.* 59, 145–160.
- R Core Development Team. 2012. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. Available at www.R-project.org.

- Reinert, G., and Schbath, S. 1998. Compound Poisson and Poisson process approximations for occurrences of multiple words in Markov chains. *J. Comput. Biol.* 5, 223–253.
- Reinert, G., Schbath, S., and Waterman, M. 2005. Statistics on words with applications to biological sequences. In Lothaire, M., ed., *Applied Combinatorics on Words*, Chapter 6. Cambridge University Press, Cambridge.
- Reinert, G., Chew, D., Sun, F., and Waterman, M.S. 2009. Alignment-free sequence comparison (I): statistics and power. *J. Comput. Biol.* 16, 1615–1634.
- Snedecor, G.W., and Cochran, W.G. 1980. *Statistical Methods*, 7th ed. Iowa State University Press, Ames, IA.
- Stuart, G., Moffett, K., and Baker, S. 2002. Integrated gene and species phylogenies from unaligned whole genome protein sequences. *Bioinformatics* 18, 100–108.
- Stuart, G., Moffett, K. and Leader, J. 2002. A comprehensive vertebrate phylogeny using vector representations of protein sequences from whole genomes. *Mol. Biol. Evol.* 19, 554–562.
- Torney, D., Burks, C., Davison, D., and Sirotkin, K. 1990. Computation of d^2 . A measure of sequence dissimilarity, 109–125. In Bell, G., and Mrarr, T., eds. *Computers and DNA*, Santa Fe Institute Studies in the Sciences of Complexity. Addison-Wesley, New York.
- Vinga, S., and Almeida, J. 2003. Alignment-free sequence comparison—a review. *Bioinformatics* 19, 513–523.
- Wan, L., Reinert, G., Sun, F., and Waterman, M.S. 2010. Alignment-free sequence comparison (II): theoretical power of comparison statistics. *J. Comput. Biol.* 17, 1467–1490.
- Waterman, M.S. 1995. *Introduction to Computational Biology*. Chapman and Hall, London.
- Wellcome Trust Sanger Institute and European Bioinformatics Institute. 2012. Ensembl Genome Browser. Homo Sapiens DNA. Available at ftp.ensembl.org/pub/release-68/fasta/homo_sapiens/dna/, file `Homo_sapiens.GRCh37.68.dna_sm.chromosome.1.fa.gz`.

Address correspondence to:
Conrad J. Burden
Mathematical Sciences Institute
Australian National University
Canberra, ACT 0200
Australia

E-mail: conrad.burden@anu.edu.au