

Computing exact *p*-values for DNA motifs (Part I)

Jing Zhang¹, Bo Jiang², Ming Li^{3,1}, John Tromp⁴, Xuegong Zhang², Michael Q. Zhang^{5,2}

¹State Key Laboratory of Intelligent Technology & System, Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China.

²Bioinformatics Division, TNLIST and Department of Automation, Tsinghua University, Beijing 100084, China.

³School of Computer Science, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada.

⁴CWI, P.O. Box 94079, 1090 GB Amsterdam, The Netherlands.

⁵Cold Spring Harbor Laboratory, Cold Spring Harbor, NY 11274, USA.

ABSTRACT

Motivation: Many heuristic algorithms have been designed to approximate *p*-values of DNA motifs described by position weight matrices, for evaluating their statistical significance. They often significantly deviate from the true *p*-value by orders of magnitude. Exact *p*-value computation is needed for ranking the motifs. Furthermore, surprisingly, the complexity of the problem is unknown.

Results: We show the problem to be NP-hard, and present MotifRank, software based on dynamic programming, to calculate exact *p*-values of motifs. We define the exact *p*-value on a general and more precise model. Asymptotically, MotifRank is faster than the best exact *p*-value computing algorithm, and is in fact practical. Our experiments clearly demonstrate that MotifRank significantly improves the accuracy of existing approximation algorithms.

Availability: MotifRank is available from <http://bio.dlg.cn>

Contact: mzhang@cshl.edu, mli@uwaterloo.ca

1 INTRODUCTION

Transcription factors play a prominent role in gene regulation; identifying and characterizing their binding sites is central to annotating genomic regulatory regions and understanding gene-regulatory networks. An important aspect of this is to determine the statistical significance of the occurrences of transcription factor binding site (TFBS), also called motifs, in a DNA sequence. Statistical measures used for evaluating overabundance of patterns in sequences have been studied extensively, among which the *z*-score and *p*-value are most popular. Between them, the *p*-value is more reliable in statistical significance evaluation [7].

In previous studies, statisticians and computational biologists have devoted considerable efforts to solving these

problems. Some are concerned with theoretical asymptotic word count distributions using either Gaussian [5, 17, 29], (Compound) Poisson [6, 10, 24] or Large Deviation approximations [7]. The problem of pattern autocorrelation has been well studied theoretically, beginning with its introduction by Guibas and Odlyzko [12]. Other works concerned with the exact distribution of word counts using generating functions [9, 14, 16, 11, 22] or other methods [3, 4, 23] soon followed. However these methods are too theoretical to be of practical use, especially when dealing with motifs described by a Position Weight Matrix (PWM). In practice, a variety of motif discovery tools evaluates results based on various approximate statistic criteria, such as Ahab [21], Clover [8], MotifScanner [28], MEME [1, 2], Consensus [13], and more recently the elegant approximate *p*-value calculation for multiple alignment [20]. These tools take less time to compute but give less accuracy in varying degrees or on different models.

Despite of all these efforts over many years, it is surprising that we do not even know the complexity of this problem. Is it really difficult? Neither approximation approaches nor the current inefficient exact *p*-value algorithms can be justified if there exists an efficient algorithm for computing exact *p*-values. This paper answers all these questions.

It turns out that the answers can be derived from a different field of bioinformatics: optimal spaced seeds for homology search introduced by [19]. By connecting the two fields, we not only obtain a proof that the *p*-value calculation problem is NP-hard, but also a practical algorithm, based on dynamic programming, to calculate the exact *p*-value of a PWM motif given in the most general form as a Position Weight Matrix (PWM). We use first-order Markov model to characterize the promoter region instead of i.i.d model. It is more general

and contains more information. Its time complexity is good enough to provide a real time calculation on typical biological data.

We have implemented the algorithm as the program MotifRank, to calculate exact p-values and to rank candidate motifs. The ability of MotifRank in discriminating known TFBSs from their backgrounds is validated on a simulated data set. Moreover, MotifRank compares favorably with MotifScanner, which estimates p-values, on a data set from *Saccharomyces cerevisiae* Promoter Database (SCPD) [30]. Our extensive tests show that, as expected, ranking of motifs according to exact p-values is better than using the approximate values given by MotifScanner.

2 PROBLEM DESCRIPTION

Given a set of genomic DNA sequences and a set of candidate motifs, the general problem of evaluating statistical significance of DNA motifs is to rank such motifs according to a underlying model. From a theoretical point of view, regulatory regions can be divided into two parts: the binding sites which play an important role in regulating gene expression, and the background which is not bound by transcription factors of interest. The key point for discriminating the signals from the background is to estimate if the motif is over-represented under the null hypothesis. We use a Markov chain to model the background and Position Weight Matrices (PWMs) with appropriate thresholds to describe motifs; we take p-value as our statistic to rank the potential biological relevance of different motifs. Assuming the observed number of occurrences of motif m is k , its p-value is defined to be the probability that m hits (matches) the Markov chain at least k times. The central problem here is how to calculate the exact p-value efficiently.

Without loss of generality, we will use first-order Markov chains in this paper to model the genetic sequences. Generalization to higher order is straightforward. For any string (or Markov chain) S , we use $S[i]$ to denote the i th position of S , and $S[i, j]$ to denote the substring of S from $S[i]$ to $S[j]$ inclusive, i.e. $S[i]S[i+1] \dots S[j]$. Consider a first-order Markov chain R of length n . In the basic problem, let us consider word motif $\{A, C, G, T\}$. A motif m is considered to hit a string if the string contains the motif as a substring. That is, there $\exists i, j$, such that $m = R[i, j]$. The number of hits is the number of such different i , regardless of overlaps. Now we give the formal definition of the basic problem.

Input: A word m over alphabet $\Omega = \{A, C, G, T\}$, an integer $k \geq 1$, the stationary probability and transition probability of a first-order Markov region R

Output: $\Pr(m \text{ hits region } R \text{ at least } k \text{ times})$

We next consider the multiple-word motif, which hits a string if any of its words does. This gives the following intermediate problem:

Input: A set of words $M = \{m_1, m_2, \dots, m_t\}$ over alphabet Ω , an integer k , the stationary probability and transition probability of a first-order Markov region R

Output: $\Pr(M \text{ hits region } R \text{ at least } k \text{ times})$

When we use a PWM to describe a motif, not only is the sequence region probabilistic, but the motif itself is as well. The position weight matrix representation of a motif m of length l is a $4 \times l$ PWM matrix P , in which $P_{i,j}$ is the frequency for the i th character at the j th position of m . For a string s of the same length l , we define the score of s as the product of corresponding frequency scores in P . (alternatively, we will sometimes use log odds ratios, the logarithm of the ratio of a letter's frequency at a position in the motif to its average frequency in the entire background). The meaning of "hit" here is different from the word motif model. We say a motif m hits a string s when there is a substring s' of length l whose score is at least some threshold c_0 . In biological applications, there are different methods to define c_0 and here it is regarded as an input parameter. Now, We define our third and central problem.

Input: A motif m with PWM P over alphabet $\Omega = \{A, C, G, T\}$, a threshold c_0 , an integer k , the stationary probability and transition probability of a first-order Markov region R

Output: $\Pr(m \text{ hits region } R \text{ at least } k \text{ times})$

We will give a linear time DP algorithm to calculate the p-value for a word motif in subsection 3.1. We will adapt the DP algorithm to solve the word set problem in subsection 3.2. The time and space complexity analysis will be given in subsection 4.2. Calculating the p-value of a matrix motif is the main problem of this paper. In subsection 4.1, we will prove that it is an NP-hard problem and give an efficient algorithm for it by reducing it to the second problem in subsection 3.3.

3 ALGORITHMS

3.1 Calculating the p-value for a word motif

We first give a simple algorithm for the case of $k = 1$ (i.e. the motif hits the region at least once) and then sketch an algorithm for the general k . The algorithms and ideas for computing the optimal spaced seeds [19, 15] can be readily borrowed and generalized to our situation.

The basic idea is to calculate a series of conditional probabilities instead of the target probability. The conditional probabilities can be calculated by DP and the number of such probabilities is polynomial to n . We will give the formal definition of the conditional probability first and then constrain the domain of the parameters of the conditional probability.

For a string w over alphabet $\Omega = \{A, C, G, T\}$ and $|w| \leq n - i$, the conditional probability is defined to be

$$f(i, w) = \Pr(m \text{ hits } R[i, n] | w \text{ is a prefix of } R[i, n]) \quad (1)$$

In another word, the conditional probability is the hit probability of motif word m in a subregion of R under the condition that the prefix of the subregion is w . We can see that the target hit probability of m in region R equals $f(1, \varepsilon)$. We will try to compute $f(i, w)$ in terms of other $f(i', w')$ computed earlier and limit the set of w we need to consider in the process.

For any $f(i, w)$, we decompose it according to the character following w in region $R[i, n]$, and then the following relation holds:

$$f(i, w) = \sum_{c \in \Omega} (f(i, wc) \times \Pr(R[i + |w|] = c | w \text{ is a prefix of } R[i, n]))$$

$\Pr(R[i + |w|] = c | w \text{ is a prefix of } R[i, n])$ expresses the probability that the character following w in region $R[i, n]$ is c . It can be easily derived from the transition matrix of the first-order Markov chain when $|w| > 0$. The case $|w| = 0$ requires additional information about the last character before the region. Similar to equation (1), for any $i \leq n$, we define

$$F(i, q) = \Pr(m \text{ hits } R[i, n] | R[i - 1] = q). \quad (2)$$

We have, for $i > 0$,

$$F(i, q) = \sum_{c \in \Omega} f(i, c) \times \Pr(R[i] = c | R[i - 1] = q) \quad (3)$$

We next determine the domain of w . Clearly, if w is not a prefix of m , $\Pr(m = R[i, i + |m| - 1] | w \text{ is a prefix of } R[i, n]) = 0$, which means that m cannot start at $R[i]$. We can see that, when $w[j, |w| - 1]$ is not a prefix of m , m can not start at $R[j]$ which has a prefix $w[j, |w| - 1]$. So we have to find the longest suffix of w which is also a prefix of m . This leads us to the following definition.

DEFINITION 1. Let $P(m)$ be the set of all prefixes of a word m . For any string w , let $S_m(w)$ denote the longest suffix of w which is in $P(m)$.

For example, if $m = ACCAC$ and $w = CCAC$, then $S_m(w) = AC$ which is a suffix of w and a prefix of m . It's the longest string that satisfies the two requirements simultaneously.

The following observation helps to constrain the domain of w in $P(m)$. For w does not belong to $P(m)$,

$$f(i, w) = \begin{cases} F(i', w_{-1}) & \text{if } S_m(w) = \epsilon \\ f(i', S_m(w)) & \text{otherwise} \end{cases} \quad (4)$$

where $i' = i + |w| - |S_m(w)|$ and w_{-1} denotes the last letter of w . Thus, we only have to compute w which is a prefix of m . Algorithm 1 shows that $f(i, w)$ and $F(i, q)$ can be computed by DP in polynomial time.

We show how to calculate $f(i, w)$ in Algorithm 2. $F(i, q)$ can be calculated similarly.

Algorithm 1 Dynamic Programming for $k = 1$

Input: A motif word m over $\Omega = \{A, C, G, T\}$, the stationary probability $I_{1 \times 4}$ and transition matrix $T_{4 \times 4}$ of a first-order Markov region R .

Output: Probability that m hits R at least once.

```

1: Calculate map  $g : P(m) \times \Omega \mapsto P(m), \forall w \in P(m)$ 
   and  $c \in \Omega, g(w, c) = S_m(wc)$  using suffix tree
2: for  $i = n \rightarrow 1$  do
3:   for  $w \in P(m)$  from longest to shortest do
4:     if  $|w| = 0$  then
5:       calculate  $F(i, q), \forall q \in \Omega$ 
6:     else
7:       calculate  $f(i, w)$ 
8:     end if
9:   end for
10: end for
11: output  $\sum_{q \in \Omega} I(q)F(1, q)$ 

```

Algorithm 2 Calculate $f(i, w)$

```

1: if  $|m| > n + 1 - i$  then
2:    $f(i, w) = 0$ 
3: else if  $w = m$  then
4:    $f(i, w) = 1$ 
5: else
6:   for  $c \in \Omega$  do
7:      $r_c = T[w_{-1}][c]$  which is the transition probability
       from  $w_{-1}$  to  $c$ , where  $w_{-1}$  denotes the last character
       of  $w$ 
8:      $i' = i + |wc| - |S_m(wc)|$ 
9:     if  $|S_m(wc)| = 0$  then
10:       $f_c = F(i', c)$ 
11:     else
12:       $f_c = f(i', S_m(wc))$ 
13:     end if
14:   end for
15:    $f(i, w) = \sum_{c \in \Omega} f_c \times r_c$ 
16: end if

```

We generalize algorithms 1 and 2 to arbitrary k by defining a series of probabilities $\{f^{(1)}(i, w), f^{(2)}(i, w), \dots, f^{(k)}(i, w)\}$ where

$$f^{(j)}(i, w) = \Pr(m \text{ hits } R[i, n] \text{ at least } j \text{ times} | w \text{ is a prefix of } R[i, n])$$

for $1 \leq j \leq k$. $f^{(k)}(1, \varepsilon)$ is exactly the p-value we want to calculate. $F^{(j)}(i, q)$ is defined similarly.

The recursion formulae here are

- When $w = m$, $f^{(j)}(i, w) = \sum_{c \in \Omega} (f^{(j-1)}(i, wc) \times \Pr(R[i + |w|] = c | w \text{ is a prefix of } R[i, n]))$

- When $w \neq m$, $f^{(j)}(i, w) = \sum_{c \in \Omega} (f^{(j)}(i, wc) \times \Pr(R[i + |w|] = c | w \text{ is a prefix of } R[i, n]))$

In another word, we use the conditional probabilities in a smaller subregion with fewer hits and a prefix constructed from w to calculate $f^{(j)}(i, w)$.

Algorithm 3 Computing p-value of a motif word

- 1: Calculate the map g using suffix tree
 - 2: **for** $i = n \rightarrow 1$ **do**
 - 3: **for** $j = 1 \rightarrow k$ **do**
 - 4: **for** $w \in P(m)$ from longest to shortest **do**
 - 5: Calculate $f^{(j)}(i, w)$
 - 6: **end for**
 - 7: **end for**
 - 8: **end for**
-

The details of calculating $f^{(j)}(i, w)$ are the same as Algorithm 1 and 2 with appropriate modification of the recursion formulae. The complete time complexity analysis will be given in subsection 4.2.

3.2 Calculating p-value for multiple-word motif

The basic idea is similar to Algorithm 1. We just extend all the definitions in subsection 3.1 from a single word to a set of words. Two main definitions are conditional probability and prefix set.

The conditional probability for multiple-word motif M is

$$\mathbb{f}^{(j)}(i, w) = \Pr(M \text{ hits region } R[i, n] \text{ at least } j \text{ times} | b \text{ is a prefix of } R[i, n])$$

The definition of $\mathbb{F}^{(j)}(i, q)$ is similar and we omit it here. The prefix set of M is the union of the prefix sets of all the motif words in M , that is $P(M) = \bigcup_{s=1}^t P(m_s)$. The recursion formulae for multiple-word motif are

- if w matches any word in M ,
 $\mathbb{f}^{(j)}(i, w) = \sum_{c \in \Omega} (\mathbb{f}^{(j-1)}(i, wc) \times \Pr(R[i + |w|] = c | w \text{ is a prefix of } R[i, n]))$
- otherwise,
 $\mathbb{f}^{(j)}(i, w) = \sum_{c \in \Omega} (\mathbb{f}^{(j)}(i, wc) \times \Pr(R[i + |w|] = c | w \text{ is a prefix of } R[i, n]))$

Replacing $P(m)$, $f^{(j)}(i, w)$ and $F^{(j)}(i, q)$ by $P(M)$, $\mathbb{f}^{(j)}(i, q)$ and $\mathbb{F}^{(j)}(i, q)$ in algorithms in subsection 3.1, we can get the DP algorithm to calculate p-value for multiple-word motif. We will analyze time and space complexity of the algorithm in subsection 4.2.

3.3 Computing one PWM motif hit probability

We will prove that calculating p-value for a PWM motif over a Markov region is NP-hard in subsection 4.1. This explains the prevalence of approximation and exponential algorithms

in the field. We present a DP algorithm by constructing a multiple-words motif from a matrix motif with threshold c_0 and then use the algorithms described in Section 3.2.

Given a threshold c_0 and a motif m represented by PWM P , we say a string of length $|m|$ is *compatible* with P if its motif score is at least c_0 . The main idea of the algorithm is to rank the letters in each position of the motif in decreasing order of score.

In position i of motif m , call the highest score letter $m_{i,0}$, the 2nd most $m_{i,1}$, the 3rd most $m_{i,2}$, and the lowest letter $m_{i,3}$. So each position independently ranks the letters from 0 through 3. We enumerate sequences of length $l = |m|$ in rank-lexicographic order, skipping incompatible strings whenever possible. In particular, if a string $s = m_{1,s_1} m_{2,s_2} \dots m_{l,s_l}$ is compatible with m , then its successor, the next string to be considered, is simply the rank-lexicographically next string $m_{1,s_1} \dots m_{i-1,s_{i-1}} m_{i,s_i+1} m_{i+1,0} \dots m_{l,0}$, where i is the largest position with letter rank less than 3. If, on the other hand, s is not compatible with m , then, letting j be the largest position with nonzero letter rank, we skip all 4^{l-j} sequences of the form $m_{1,s_1} \dots m_{j,s_j} m_{j+1,*} \dots m_{l,*}$. These sequences can obviously score no better than s and are therefore also incompatible. The successor of s is then $m_{1,s_1} \dots m_{i-1,s_{i-1}} m_{i,s_i+1} m_{i+1,0} \dots m_{l,s_0}$, with i the largest position before j with letter rank less than 3. Since no $l = |m|$ consecutive successors can be incompatible, we can enumerate the set of compatible strings in time $O(K \times |m|)$, where K is the size of the set of all strings compatible with m . If we consider the double stranded structure of DNA, we have to add the sequences reserve-complementary to those found.

Then we can use the algorithm for multiple-word motifs to calculate the p-value. Although it is an exponential time algorithm, in real biological applications c_0 is often large enough for MotifRank to be practical.

4 COMPLEXITY

4.1 The complexity of computing one PWM motif hit probability

In this subsection, we will prove that calculating p-value for a PWM motif over a Markov region is NP-hard by reducing the spaced seed problem to it. It indicates there does not exist a polynomial time complexity algorithm to calculate one PWM motif hit probability unless $P = NP$.

THEOREM 1. *Calculating the probability that a given PWM motif m over a constant size alphabet hits a first-order Markov region at least once at a given threshold is NP-hard.*

PROOF. Consider the simple case of a binary alphabet and an i.i.d. region. We reduce the *spaced seed* problem to this case. We describe the spaced seed problem here briefly. A spaced seed is a string over $\{0, 1\}$. There is an i.i.d. sequence R , each bit with probability p (similarity level) to be 0, and $1 - p$ to 1. A spaced seed s hits a string r when r has a

substring r' of the same length as s that is ‘implied’ by s , i.e. wherever s has a 1, r' also has a 1. The spaced seed sensitivity problem is to calculate the hit probability of a spaced seed on a binary i.i.d. sequence R of a given similarity level. It has been recently shown to be NP-hard for similarity $1/2$ in [18].

We reduce the spaced seed problem to our problem. Given a spaced seed s of length l with z 0s, we construct a motif m of length l over alphabet $\{0, 1\}$ as follows: for $1 \leq i \leq l$, if $s_i = 1$, we let $\Pr(m_i = 1) = 1$ and $\Pr(m_i = 0) = 0$, otherwise $\Pr(m_i = 1) = \Pr(m_i = 0) = 1/2$. Then a string r of length l is implied by s iff r has a nonzero motif score with respect to m . Thus, with a sufficiently small threshold like 2^{-z} , the probability of m passing the threshold in a random region is precisely the hit probability of spaced seed s . \square

4.2 Complexity of Algorithms in Section 3

We will give the time and space complexity analysis for the algorithms described in section 3. Lemma 1 and Theorem 2 are about Algorithm 3 for calculating the p-value for a word motif given in subsection 3.1. Theorem 3 is about the algorithm for calculating the p-value for multiple-word motif given in subsection 3.3.

The line 1 of Algorithm 3 is to calculate the mapping $g : P(m) \times \Omega \mapsto P(m)$, which means to find the longest suffix of the string composed of a prefix string in the prefix set $P(m)$ appending a character c and the suffix must be a prefix in $P(m)$, too. We have the following lemma about the time to calculate $S_m(w)$ for arbitrary string w .

LEMMA 1. *Given a word m and a string w over a constant alphabet, $S_m(w)$ can be calculated in $O(m)$.*

PROOF. Let \overleftarrow{m} denote the reverse of string m . Construct suffix tree for \overleftarrow{m} and search \overleftarrow{w} on it. The result path is $S_m(w)$. The time complexity of constructing a suffix tree is $O(m)$ and searching is $O(m)$, so the total time complexity is $O(m)$. \square

We have the following theorem about the time and space complexity of Algorithm 3.

THEOREM 2. *The probability that a word motif m over alphabet $\Omega = \{A, C, G, T\}$ hits a first-order Markov region R for at least k times can be calculated in polynomial time. The time complexity is $O(k \times n \times |m|)$ and the space complexity is $O(k \times |m|^2)$, where n is the length of R .*

PROOF. According to Lemma 1, line 1 takes time $O(|m|^2)$. Consider how many $f^{(j)}(i, w)$ will be calculated in Algorithm 3 (this will also bound the number of $F^{(j)}(i, q)$ computed). The three nested loops in Algorithm 3 produce $O(k \times n \times |m|)$ iterations. Each iteration calculates a $f^{(j)}(i, w)$ in constant time. So the time to calculate all f is $O(k \times n \times |m|)$. The space complexity is $O(k \times |m|^2)$ since when calculating $f^{(j)}(i, w)$, we only have to maintain $f^{(j)}(j, w)$ for $i \leq j \leq i + m$ of size at most $O(k \times |m|^2)$. \square

The algorithm for multiple-word motif is an extension to Algorithm 3, so the time and space analysis are similar to Theorem 2. We give the results directly.

THEOREM 3. *The probability that a set of word motifs $M = \{m_1, m_2, \dots, m_t\}$ over alphabet $\Omega = \{A, C, G, T\}$ hits a first-order Markov region R at least k times can be calculated in polynomial time. Let $\|M\| = \sum_{i=1}^t |m_i|$. The time complexity is $O(k \times n \times \|M\|)$ and the space complexity is $O(k \times \|M\| \times (\max_i |m_i|))$.*

Now we analyze the time complexity of the algorithms described in subsection 3.3. The algorithm to extract multiple-words from PWM motif has time complexity $O(K \times |m|)$ as already shown in the algorithm. Adding the time to calculate the p-value for multiple-words, the total time complexity of our exact p-value calculating algorithm is $O(K \times |m| + k \times n \times \|M\|)$, in which $M = \{m_1, m_2, \dots, m_K\}$ is the set of compatible strings, $\|M\| = \sum_{i=1}^K |m_i|$ and k is the hit time of M in the region of length n .

5 EXPERIMENTS AND RESULTS

5.1 Experiment Methods

We apply our algorithm given in subsection 3.3 to the statistical significance evaluation of PWM motifs and rank them according to their exact p-values.

Consider a first-order Markov model whose elements in transition matrix Π are estimated from a promoter region. The initial probability of the Markov model is chosen to be the stationary distribution μ defined by $\mu = \mu\Pi$.

There are many strategies for choosing the threshold of a PWM motif on a specific sequence. Here we adopt a simple way by allowing a user-defined multiplying factor, which was introduced in [27]. We then multiply the maximum of scores of all $|m|$ long strings on PWM motif m by the factor to define the threshold c_0 . This factor allows the user to define the type of motif to search for. It controls the balance between specificity and sensitivity. When the factor is reduced, more degeneracy in the extraction of motif words will be allowed, so sensitivity will increase but specificity will decrease. Besides, we remove the words which are not substrings of the promoter sequence from the word set extracted from m according to c_0 . Furthermore, since DNA is double stranded, we need to look for binding sites on both strands.

We use the algorithm introduced in subsection 3.3 to calculate the exact p-values for motifs in promoter regions.

5.2 Comparison with MotifScanner on SCPD Data

The *Saccharomyces Cerevisiae* (yeast) Promoter Database SCPD contains 130 promoter regions of different genes with known binding sites coming from 21 PWM motifs. We extract upstream (of ATG) sequences, of length 1001 basepairs, for

Number of reported motifs per region	Number of promoter regions	MotifRank		MotifScanner	
		Average rank	Standard deviation	Average rank	Standard deviation
1	78	2.33	1.46	3.09	1.84
2	22	2.95	1.79	3.34	2.12
3	7	3.24	1.95	3.38	1.64
4	1	5.50	3.42	6.75	1.71
Sum Total	108	2.73	1.78	3.30	2.06

Table 1. Average and standard deviation of the rank of known TFs with p-values reported by different tools. The promoter regions are divided into four groups according to the number of TFs reported to bind in them.

all 130 genes from SCPD as our experimental data set. For each promoter region, we estimate the parameters of Markov model, and calculate the exact p-values for each of the 21 PWM motifs by using MotifRank on double strands and with specific factor.

We compare our results on SCPD data with a similar tool, MotifScanner [28], which estimates p-values. MotifScanner is a program that can be used to screen DNA sequences with pre-compiled motif models. The input of MotifScanner consists of four parts: DNA sequences, candidate motifs represented by Position Frequency Matrices, a prior probability, which allows the user to specify the degeneracy in the found instances of motifs, and a background model. We input the same SCPD data set and choose first-order Markov Model on double strands with prior probability 0.2, which is suggested by MotifScanner for yeast data. MotifScanner outputs the approximate p-values of PWM motifs which are regarded to be over-represented in the region. As a result, the outputs do not always contain the p-values for all the known motifs reported by biologists.

We rank the PWM motifs in each promoter region in increasing order of p-values, and compare the ranks of known motifs annotated in SCPD. As aforementioned, the value of multiplying factor in MotifRank determines the degeneracy of instances retrieved from PWM motifs, and so does the prior probability in MotifScanner. Both tools will fail to report the instances of some known motifs when a relatively strict constraint on degeneracy is imposed. To ensure that the comparison with MotifScanner is under similar degrees of degeneracy, we adopt a multiplying factor of 0.65 in MotifRank, with which the instances of known motifs reported by both tools are nearly the same. To better assess the results, we divide the promoter regions into different groups. Majority of promoter regions contain only one known motifs, but some contain more. For example, the CDC2 promoter region only has MCB motif reported but CDC9 promoter region has two motifs: MCB and REB1. We thus divide the regions into several groups according to the reported number of distinct motifs and compare ranks given by different tools within each group. We use the average rank of the known motifs within each region group to evaluate the rank results. As a criterion reflecting the overall performance, a lower

average rank of reported motifs means that the tool will assign a comparatively lower p-value to the transcription factors known to bind in promoter regions.

The average rank and its standard deviation using different tools are given in Table 1. Here we only consider the ranks of known motifs with p-values given by both tools, so there are totally 108 promoter regions in Table 1. From the average rank of known motifs, we can demonstrate the advantage of exact p-values over approximate ones (in the sense of giving a lower average rank for the motifs which are reported to bind in corresponding promoter region). Remarkably, our method has a higher performance on words that tend to have self-overlaps than MotifScanner does. For example, the motif of RAP1 binding sites "R(A/G)M(C/A)ACCCANACAY(C/T)" which has several self-overlaps could be ranked first by MotifRank when it is within most of its binding promoters, whereas MotifScanner would give a worse rank for it. More details are available in the supplement.

We also use the ranks given by MotifRank and MotifScanner to predict putative TFs that bind in the promoter region. Given an integer N , a PWM motif is said to be positive in a specific promoter region if its rank is among the top N , and the fraction of true-positives and false-positives can be calculated according to the known motifs reported by biologists. Thus we get a ROC curve over a range of possible N 's. The ROC curves of MotifScanner and MotifRank with multiplying factors ranging from 0.6 to 0.8, a range that allows moderate degeneracy of in the extraction of motif words, are illustrated in Figure 1. We can see that the area under the curves of MotifRank is significantly larger than that of MotifScanner, which shows the advantage of exact p-values over approximate ones in predicting the binding of TFs.

5.3 Performance on Simulated Data

We test our method on simulated data to further validate its ability to discriminate known TFBSs from backgrounds. We simulated a data set with 100 sequences according to the first-order Markov model and parameters extracted from SCPD. We randomly choose instances of motifs from the frequency matrix of 20 PWM motifs in SCPD and uniformly implanted them into foreground sequences. The number of motifs and their occurrences in each sequence is similar to that in real data




Fig. 1. ROC curves manifest the ability of MotifRank and MotifScanner to identify TFs that bind in promoters regions. Different points on a ROC curve correspond to cutoffs on the ranking list varying from 1 to 5. X-axis denotes false positives rate, i.e. $FP/(FP+TN)$, while y-axis denotes true positives rate, i.e. $TP/(TP+FN)$.

sets. Simultaneously, we generate 100 background sequences under the same model but without any implanted motif, and calculate the average p-value of each PWM motif. The average p-values of implanted PWM motifs on foreground and background sequences are given in Table 2.

There is a more than 30 fold difference relative to the background. Our method is indeed powerful to distinguish sites that contain true motifs from those that do not. Its high performance on the SCPD data is consistent in both this subsection and Subsection 5.2 in our paper.

5.4 Time Performance

We have proved that the worst case time complexity of the algorithm introduced in subsection 3.3 is exponential and there is no polynomial time algorithm for exact p-value calculation of the PWM motifs, unless $NP = P$. The algorithm is implemented using C++ and compiled under Microsoft Visual Studio 6.0. The test PC is equipped with a Pentium 4 running at 2.8GHz and 512MB of RAM. The average time to compute the p-value of one motif with a 1001 basepair promoter region is 0.7msec. This indicates that our DP algorithm is efficient enough for biological data in practice.

6 CONCLUSION AND FUTURE WORK

Accurate and efficient methods for calculating the p-value of DNA motifs are critical in aiding the discovery of TFBS in promoter regions. Although existing tools provide

TF Name	Length	Foreground	Background
ABF1	12	1.85E-08	0.94
CSRE	13	1.21E-13	1.00
SCB	7	2.38E-06	0.69
GCN4	6	4.42E-03	0.17
GCR1	5	1.86E-01	0.47
MCB	6	1.37E-04	0.59
MCM1	10	6.01E-09	0.74
MATalpha2	9	1.08E-04	0.84
MIG1	12	3.26E-09	0.93
PHO2	7	3.20E-02	0.12
PHO4	10	2.19E-10	0.91
RAP1	12	1.38E-05	0.85
REB1	7	5.65E-05	0.70
ROX1	12	8.34E-02	1.00
SWI5	12	3.82E-05	0.78
STE12	8	9.05E-06	0.69
TBP	7	4.10E-03	0.32
UASPHR	12	1.33E-01	0.79
XBPI	12	1.50E-05	0.96
repressor_of_CAR1	9	2.79E-03	0.65

Table 2. Average P-values on foreground and background simulated sequences. The first two columns show the name and length of the TFs whose PWM motifs are implanted into foreground sequences. The entries in the Foreground column are the average p-values over foreground sequences with corresponding PWM motifs implanted. In contrast, we have also run 100 sequences as background without implanting any PWM motifs and calculated the average p-values on them. The results are shown in the Background column.

approximate estimates of p-values, they are not accurate enough sometimes and there is a need for efficient algorithms to calculate exact p-values.

A rapid algorithm to compute exact p-values for DNA motifs is presented in this paper. We have validated the performance of the algorithm on both real and simulated biological data. Comparison with other tools confirms that the performance of exact p-values is better than approximate p-values in terms of evaluating the statistical significance of motifs.

Our algorithm for calculating exact p-values is useful and feasible wherever the evaluation of statistical significance for candidate motifs is desired. Furthermore, our method could be integrated into discovery of TFBSs in given promoter regions.

In future work, we will study the algorithms and experiments on evaluating statistical significance of synergistic motifs and TFBSs common to multiple sequences.

7 ACKNOWLEDGEMENT

ZJ' work is supported by the National Natural Science Foundation of China Grant 60553001 and the National

Basic Research Program of China Grant 2007CB807900, 2007CB807901. BJ and XZ are supported in part by NSFC grants 60234020, 30625012. ML's work was partially supported by the Chang Jiang Scholarship Program, NSERC, and Canada Research Chair program. MZ is partially supported by the Chang Jiang Scholarship Program and a NIH grant. JT's work has been partially funded by the Dutch BSIK/BRICKS project.

REFERENCES

- [1] T.L. Bailey and M. Gribskov. (1998) Combining evidence using p-values: application to sequence homology searches, *Bioinformatics*, 14, 48–54.
- [2] T.L. Bailey and C. Elkan.(1994) Fitting a mixture model by expectation maximization to discover motifs in biopolymers. *Proceedings of the Second International Conference on Intelligent Systems for Molecular Biology*, 28–36.
- [3] M. Beckstette et al. (2006) Fast index based algorithms and software for matching position specific scoring matrices in BMC Bioinformatics, 7
- [4] G. Bejerano, N. Friedman, N. Tishby. (2004) Efficient exact p-value computation for small sample, sparse, and surprising categorical data. *J. of Comp. Bio.*, 11, 867–886.
- [5] V. Brendel, J. Beckmann, E. Trifonov. (1986) Linguistics of nucleotide sequences: Morphology and comparison of vocabularies, *J. Biomol. Struct. Dynamics*, 4.
- [6] O. Chryssaphinou, S. Papastavridis. (1988) A limit theorem for the number of non-overlapping occurrences of a pattern in a sequence of independent trials, *J. Appl. Prob.*, 25, 428–431.
- [7] A. Denise, M. Régnier, M. Vandenberghe. (2001) Assessing statistical significance of overrepresented oligonucleotides. In *WABI'01*, 85–97.
- [8] M.C. Frith, Y. Fu, L. Yu, J. F. Chen, U. Hansen, Z. Weng. (2004) Detection of functional DNA motifs via statistical over-representation, *N.A.R.*, 32(4), 1372–81.
- [9] J. Gentleman, R. Mullin. (1989) The distribution of the frequency of occurrence of nucleotide subsequence, based on their overlap capability. *Biometrics*, 45, 35–52.
- [10] A.P. Godbole. (1991) Poisson approximations for runs and patterns of rare events, *Adv. Appl. Prob.*, 23, 851–865.
- [11] I.P. Goulden and D.M. Jackson (1983) Combinatorial Enumeration.
- [12] L.J. Guibas, A.M. Odlyzko. (1981) *Journal of Combinatorial Theory*, Series A 30.
- [13] G.Z. Hertz and G.D. Stormo. (1999) Identifying DNA and protein patterns with statistically significant alignment of multiple sequences. *Bioinformatics*, 15.
- [14] L. Hertzberg, O. Zuk, G. Getz and E. Domany. (2005) Finding Motifs in Promoter Regions, *J. Comp. Biol.*, 12(3), 314–330.
- [15] U. Keich, M. Li, B. Ma, and J. Tromp. (2004) On spaced seeds for similarity search. *Discrete Applied Mathematics* 138(3), 253–263.
- [16] J. Kleffe, U. Langbecker. (1990) Exact computation of pattern probabilities in random sequences generated by Markov chains. *Comp. Applic. Biosci.*, 6, 347–353.
- [17] M.Y. Leung, G.M. Marsh, and T.P. Speed. (1996) Over and underrepresentation of short DNA words in Herpesvirus genomes, *J. Comp. Biol.*, 3, 345–360.
- [18] M. Li, B. Ma, L. Zhang. Superiority and complexity of the spaced seeds. *SODA'06*
- [19] B. Ma, J. Tromp, and M. Li. (2002) PatternHunter: faster and more sensitive homology search. *Bioinformatics* 18, 440–445.
- [20] N. Nagarajan, N. Jones, and U. Keich. (2005) Computing the p-value of the information content from an alignment of multiple sequences. *Bioinformatics*, Vol. 21, Suppl 1, *ISMB '05*, i311–i318.
- [21] N. Rajewsky, M. Vergassola, U. Gaul and E. D. Siggia. (2002) Computational detection of genomic cis regulatory modules, applied to body patterning in the early *Drosophila* embryo, *BMC Bioinformatics*.
- [22] M. Régnier. (2001) A unified approach to word occurrence probabilities. *Discrete Applied Mathematics*, 104, 259–280.
- [23] R. Staden. (1998) Methods for calculating the probabilities of finding patterns in sequences. *Comput. Appl. Biosci.*, 5, 89–96.
- [24] S. Schbath. (1995) Compound Poisson approximation of word counts in DNA sequences, *ESAIM: Probability and Statistics*, 1, 1–16.
- [25] R.J. Simes. (1986) An improved Bonferroni procedure for multiple tests of significance. *Biometrika*, 73, 751–754.
- [26] G.D. Stormo. (2000) DNA binding sites: representation and discovery. *Bioinformatics*, 16, 16–23.
- [27] F. Tronche, F. Ringeisen, M. Blumenfeld, M. Yaniv, M. Pontoglio (1997) Analysis of the Distribution of Binding Sites for a Tissue-specific Transcription Factor in the Vertebrate Genome. *J. Mol. Biol* 266, 234–245
- [28] G. Thijs, K. Marchal, M. Lescot, S. Rombauts, B. De Moor, P. Rouzé and Y. Moreau. (2001) A Gibbs Sampling method to detect over-represented motifs in upstream regions of coexpressed genes, *Proceedings Recomb'2001*, 305–312.
- [29] M.S. Waterman. (1995) *Introduction to Computational Biology*. Chapman & Hall.
- [30] J. Zhu, M.Q. Zhang. (1999) SCPD: A promoter database of yeast *Saccharomyces cerevisiae*. *Bioinformatics*, 15, 607–611.