


This article was downloaded by:
Publisher: KKG Publications



Key Knowledge Generation

Publication details, including instructions for author and subscription information:

<http://kkgpublications.com/technology/>

EMS-GT2: An Improved Exact Solution for the (l, d) -Planted Motif Problem

MARK JOSEPH D. RONQUILLO¹, PROCESO L. FERNANDEZ²

^{1,2} Ateneo de Manila University, Quezon City, Philippines

Published online: 22 June 2017

To cite this article: M. J. D. Ronquillo and P. L. Fernandez, “EMS-GT2: An improved exact solution for the (l, d) -planted motif problem,” *International Journal of Technology and Engineering Studies*, vol. 3, no. 3, pp. 124-132, 2017.

DOI: <https://dx.doi.org/10.20469/ijtes.3.40005-3>

To link to this article: <http://kkgpublications.com/wp-content/uploads/2017/3/IJTES-40005-3.pdf>

PLEASE SCROLL DOWN FOR ARTICLE

KKG Publications makes every effort to ascertain the precision of all the information (the “Content”) contained in the publications on our platform. However, KKG Publications, our agents, and our licensors make no representations or warranties whatsoever as to the accuracy, completeness, or suitability for any purpose of the content. All opinions and views stated in this publication are not endorsed by KKG Publications. These are purely the opinions and views of authors. The accuracy of the content should not be relied upon and primary sources of information should be considered for any verification. KKG Publications shall not be liable for any costs, expenses, proceedings, loss, actions, demands, damages, expenses and other liabilities directly or indirectly caused in connection with given content.

This article may be utilized for research, edifying, and private study purposes. Any substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing, systematic supply, or distribution in any form to anyone is expressly verboten.

EMS-GT2: AN IMPROVED EXACT SOLUTION FOR THE (l, d) - PLANTED MOTIF PROBLEM

MARK JOSEPH D. RONQUILLO ^{1*}, PROCESO L. FERNANDEZ ²

^{1,2} Ateneo de Manila University, Quezon City, Philippines

Keywords:

Planted (l, d) -Motif Problem
 Bit-Based
 Exact Enumerative Algorithm

Received: 12 December 2016

Accepted: 18 March 2017

Published: 22 June 2017

Abstract. In this research, the EMS-GT2 algorithm, an extension of the Exact Motif Search-Generate and Test (EMS- GT), an exact enumerative algorithm for PMS, is being proposed. In EMS-GT2, a new speedup technique based on an important property is discovered and incorporated. This speedup has enabled a more efficient block-processing of candidate motifs. The C++ implementation of EMS-GT2 running on synthetic data for several PMS challenge instances demonstrates that it is competitive with both the EMS-GT and qPMS9, the two current best exact solutions for PMS. In particular, EMS-GT2 is able to reduce the run-times of EMS-GT by 20.3%, 15.8% and 22.6% for the (l, d) challenge instances (13, 4), (15, 5) and (17, 6) respectively. It also outperforms qPMS9, having runtime reductions of 91.6%, 79.3%, 82.0%, 59.4% and 9.7% for the (9, 2), (11, 3), (13, 4), (15, 5) and (17, 6) synthetic challenge instances respectively.

INTRODUCTION

DNA motif finding is a well-studied topic in computational biology. A motif is a short pattern of interest that occurs in a large amount of biological data. Detection of these motifs often leads to new biological discoveries. This may lead to finding transcription factor binding sites that help biologists understand gene functions, understand human diseases, identify potential therapeutic drug targets and conclude commonalities from different species. There are many variations of the motif search problem, such as the Simple Motif Search (SMS), the Edit-distance-based Motif Search (EMS) and the PMS which is also known as (l, d) -motif search. This study focuses on an exact enumerative algorithm for the PMS problem. PMS is formally defined in [1] as “Input are t sequences of length n each. Input also are two integers l and d . The problem is to find a motif (i.e., a sequence) M of length l . It is given that each input sequence contains a variant of M . The variants of interest are sequences that are at a hamming distance of d from M .”

In solving the PMS problem, traditional string matching is not efficient since these biological motifs are not typically exact but are subject to mutations. As a matter of fact, PMS has already been proven to be NP-hard, which means that it is very unlikely to have an algorithm that solves it in polynomial time [2].

Some of the terms used in this study are defined below:

- An l -mer is a string of length l in a DNA sequence of length m where $l < m$.
- The Hamming distance dH between two l -mers, of equal length, is equal to the number of positions where they have mismatches. Ex. $dH(\text{acttgca}, \text{actaaga}) = 3$.
- An l -mer x is considered a d -neighbor of another l -mer y if the Hamming distance between the two is at most d .
- The d -neighborhood of an l -mer x is the set $N(x, d)$ of all l -mers with at most d Hamming distance from x . i.e., $dH(x, x') \leq d$. Ex. ccgga , ccaaa , and gccta are all in $N(\text{cctta}, 2)$, where $l = 5$
- The d -neighborhood of a sequence S is the set $N(S, d)$ of all d -neighbors of all l -mers in sequence S . Ex. $N(\text{aattacg}, 2) = N(\text{aatta}, 2) \cup N(\text{attac}, 2) \cup N(\text{ttacg}, 2)$ where $l = 5$.

In this study, we introduce EMS-GT2, an improvement of EMS-GT that is an exact enumerative algorithm for the planted motif problem consisting of two phases - - Generate and Test. The EMS-GT2 algorithm leverages on the block processing of candidate motifs during its Test phase. The original EMS-GT algorithm was previously evaluated using (l, d) -challenge instances (9, 2), (11, 3), (13, 4), (15, 5) and (17, 6). It was already shown in [3] that the current implementation of EMS-GT is faster than the then state-of-the-art algorithms

*Corresponding author: Mark Joseph D. Ronquillo

†Email: markronquillo23@gmail.com

PMS8 and qPMS9 for all (l, d) -challenge instances mentioned except in (17, 6). Even though that EMS-GT2 algorithm can only solve (l, d) -instances when $l \leq 17$, studies have shown that DNA motif lengths are usually around 10 bp (i.e., base pairs, or the number of characters in the string representation) in eukaryotes and 16 bp in prokaryotes [4]. Thus, the EMS-GT2 algorithm is still significant in practical biological applications.

LITERATURE REVIEW

Motif finding has been studied extensively in the previous years. Numerous algorithms have been made for motif finding and for PMS. Each of these algorithms may be categorized as either an approximate algorithm or an exact algorithm. Approximate algorithms, although they are fast, do not guarantee the exact solution all the time. Heuristic algorithms for PMS that perform local search such as Gibbs Sampling [5], Expectation Maximization (EM) [6], and Projections [7], [8] have been previously explored in the literature.

Most of these algorithms initially work on a tuple of alignment positions that corresponds to l -mers across different string sequences in the dataset. They then iteratively refine the alignment until a certain criterion is met. MEME [6] is a tool for motif finding that implements Expected Maximization.

WINNOWER [9] reduces the PMS problem to finding a large clique in a multipartite graph. Instead of looking for the motif directly, the algorithm applies a winnowing technique to remove spurious edges that trims the graph representation, making it easier to find the motif. Other approximate algorithms are MULTIPROFILER [10], Pattern Branching, Profile Branching [11] and CONSENSUS [12].

Exact algorithms are not as fast as approximate algorithms but are guaranteed to return the correct answer for any given problem instance. Furthermore, these exact algorithms can be categorized based on their approach in solving the problem. One approach is to generate a common neighborhood out of all $(ml + 1)^n$ possible positions or l -mers from all string sequences. This approach is called sample-driven. Another approach is called pattern-driven that checks from $|\sum|^l$ possible l -mers which are the motifs over a base alphabet \sum .

[1] proposed a series of exact algorithms for the (l, d) -motif search problem. The algorithm PMS1 [1] is one of these algorithms. PMS1 solves the problem by enumerating the d -neighborhood of all the sequences in the dataset and intersects them; the result is a set of motifs. PMS_i and PMSP [13] are algorithms based on PMS1. PMS_i improves the memory space requirement of PMS1 by processing only two sequences at a time. PMSP works by generating all the d -neighborhood of each l -mer in the first sequence and testing each d -neighbor if it exists in the remaining sequences. PMSPrune [13] works

very similar to PMP with some improvements. It generates the neighborhood of an l -mer using a branch and bound approach and implements a pruning strategy to speed up the testing of l -mers.

Succeeding algorithms like PMS5 [14] and PMS6 [15] extend the ideas of PMS1 and PMSPrune. PMS5 generates the common neighborhood of three l -mers from different sequences at a time and uses Integer Linear Programming for the pruning process. PMS6 only differs from PMS5 in the way it determines the three l -mers. Quorum PMS (qPMS) is a generalized version of the (l, d) -motif search problem.

Instead of finding an l -mer that exists in all n sequences, it only considers up to q where $q \leq n$. We can see that a qPMS problem is equal to PMS when $q = n$. The qPMS7 [16] is one algorithm that solves the qPMS problem. Algorithm qPMS7 is a generalized version of qPMSPrune (quorum version of PMSPrune) combined with the pruning strategy of PMS5 algorithm. PMS8 [17] is an algorithm that combines the sample-driven approach and the pattern-driven approach.

First, it chooses k -tuple T of l -mers from k different sequences and it makes sure that all l -mers in T have a common neighbor. Each l -mer that belongs to the common neighborhood of the tuple T is checked if it appears in the remaining $n - k$ sequences. One of the current state-of-the-art algorithms, qPMS9 [18], improves the sample-driven approach of PMS8 by prioritizing l -mers that are highly distant from those already in the tuple, resulting in a smaller size of common d -neighborhood to test and enables the algorithm to process the quorum version of the PMS problem.

Several exact algorithms solve the PMS problem using suffix trees and other related data structures. RISO [19], RISOTTO [20], SPELLER [21] and SMILE are all exact algorithms that use suffix trees. MITRA [22] improves the excessive memory requirement of sample-driven approach by using a mismatch tree. Two other algorithms that have some similarity with our algorithm are the Voting algorithm and Bit-based algorithm.

Voting algorithm [23] maintains a hash table that tracks the number of occurrences of every possible l -mer and makes sure that every l -mer is only counted once in each sequence. An l -mer is considered a motif if its total occurrences are equal to the total number of sequences in the dataset.

Bit-based algorithm [24] generates the neighborhood of each sequence and intersects it to get the set of motifs. Unlike PMS1, the Bit-based algorithm maps every l -mer to its corresponding integer value. It uses an array of size $|\sum|^l$ to represent the neighborhood of a sequence and uses the integer representation of an l -mer to flag if it is a member of the array.

It generates the neighborhood of all sequences and

merges it using the logical operator AND. The resulting array represents the set of motifs.

The EMS-GT Algorithm

The Exact Motif Search - Generate and Test Algorithm (EMS-GT) for the planted motif search problem is composed of two phases, the Generate phase, and the Test phase. The Generate phase takes the first n' number of string sequences in the dataset and generates the set d -neighborhood one sequence at a time then intersects it.

This accumulates and outputs the set of candidate motifs C and is composed of l -mers that have at least one d -neighbor in each of the first n' sequences. The Test phase evaluates each candidate motif $c \in C$ by checking if c has at least one d -neighbor in each of the remaining $n - n'$ string sequences. The phases are formally defined below:

Generate Candidates

This step operates on the first n' sequences. The intersection of the d -neighborhood of each sequence results to the set of candidate motifs C .

$$C = N(S_1, d) \cap N(S_2, d) \cap \dots \cap N(S_{n'}, d). \quad (1)$$

Test Candidates

Each candidate motif in C is evaluated if it has a d -neighbor in all of the remaining $n - n'$ string sequences. If a candidate motif passes the test, it is then included in the set of motifs M .

Speedup Strategies. To further improve the algorithm performance, we introduced a number of speedup strategies, and these techniques are defined in this section:

1. Integer mapping of l -mers.

The l -mers are represented using binary representation of integer values. Each character in the l -mer is translated using 2 bits (a=00, c=01, g=10, t=11). Ex. actg maps to 00011110 and has an integer value of 30.

2. Bit-based set representation and l -mer enumeration.

The EMS-GT implementation maintains a 4^l array for enumerating all the possible l -mer values. The l -mer's integer value is used as the index value for the array. It uses the value of 1 if the l -mer is a member of the set else it sets the value to 0.

3. Bit-array compression.

To efficiently store these l -mers and save memory space, EMS-GT implements an approach that compresses the search space array using integer value bit flags. Instead of one l -mer per index value, the implementation can flag up to 32 l -mers (since we are using 32-bit integers) per index value. An illustration on how the algorithm accesses the bit flag is provided below:

Ex. gacgt maps to 1000011011 = 539 in decimal.

bit position = $539 \bmod 32 = 27$; array index = $539 / 32 = 16$; The bit flag for gacgt is in the 27th least significant bit of the integer at array index 16.

4. XOR-based Hamming distance computation.

The mapping of an l -mer to its integer value has an additional advantage in computing for mismatch positions. Applying the Boolean operator exclusive-or (XOR) between two integer values will return another integer value that contains a nonzero value for mismatch positions. Counting this nonzero bit pair positions result to the hamming distance value. An example of this computation is shown below:

Ex. aacgt maps to 0000011011

tacgc maps to 1100011001

XOR produces 1100000010 = 2 mismatches (Note, the mismatches are counted per pair)

5. Recursive neighborhood generation.

The Generate step of the algorithm produces the d -neighborhood of a string sequence by generating the d -neighborhood of all l -mers in that sequence. Our implementation of EMS-GT uses a recursive approach for generating the d -neighborhood of an l -mer. The recursive generation can be visualized by a tree $T(x)$ of height d that is generated in a depth-first manner. Each node is a tuple of (w, p) where w is an l -mer and p corresponds to a position in the l -mer $0 \leq p \leq l$. At a given node (w, p) and $p \neq l$, three children nodes are generated where each node is a variant of w that has a different character starting at the $p + 1$ position. The root node is $(x, 0)$, and any l -mer in nodes at depth t has a hamming distance of t from the l -mer x (see Figure 1). Given this, the expected size of $N(x, d)$ can be computed using the equation:

$$|N(x, d)| = \sum_{i=0}^d \binom{l}{i} 3^i \quad (2)$$

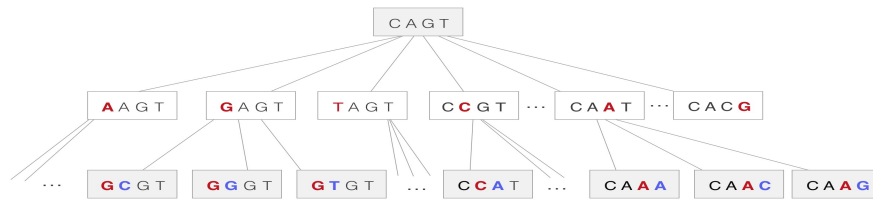


Fig. 1. Illustration of the recursive generation of neighborhood

6. Block-based optimization for neighborhood generation.

The way our implementation of the algorithm represents the neighborhood of 4^l bit flags array opens up a new way to improve the generation of neighborhood $N(x, d)$. The algorithm maintains an array of 32-bit integers where each bit represents an l -mer. Setting the bit value to 1 means that the l -mer is in $N(x, d)$, otherwise 0 means it is not. The improvement generates the neighborhood by blocks of size k where $0 < k < l$ instead of per bit. It is observed that dividing the bit array N into 4^k consecutive blocks results to blocks conforming to one of the $(k + 2)$ possible block patterns. For each possible k -mer, block patterns are pre-generated according to the remaining number of allowed mismatch d' where $0 \leq d' \leq d$. There is no need to pre-generate the block pattern for 0 value since it only corresponds to a block pattern where exactly one l -mer is set and the block pattern for d value since it corresponds to a block pattern where all bits are set. Given an l -mer x , the generation of $N(x, d)$ now divides the l -mer to its $(l - k)$ -length prefix y and its k -length suffix z . The algorithm then generates all prefixes in $N(x, d)$ by recursively generating $N(y, d)$. For each prefix $y \in N(y, d)$, a block pattern is applied to the neighborhood array based on the remaining number of allowed mismatch d' computed using $d - dH(y, y')$ and the suffix z . This is based on the observation that the distance between two l -mers is equal to the sum of the distance between their prefixes and the distance between their suffixes. The run-time complexity for generating the d -neighborhood of l -mer x is now reduced to $O(4^k \sum_{i=0}^{l-k} \binom{l-k}{i} 3^i)$

METHOD AND MATERIALS

This section states how the EMS-GT2 and the proposed speedup techniques were implemented and evaluated using synthetic datasets for some challenge (l, d) instances. A study [18] proposed that (l, d) instances where d is the largest integer value for which the expected number of motifs of length l would occur in the input by random chance and does not exceed a constant value (500) are categorized as challenging instances. The following (l, d) instances (9, 2), (11, 3), (13, 4),

(15, 5) and (17, 6) are the only challenge instances where l is between 9 and 17, and these challenge instances were used for the evaluation of the algorithms.

Datasets

Algorithms that solve PMS [9], [17] [18] use a dataset containing 20 string sequences where each nucleotide is in $\Sigma = \{a, c, g, t\}$. Each string sequence is 600 Base Pairs (bp) long and each nucleotide is randomly generated with an equal chance of being selected. A motif is then generated and for each string sequence in the dataset, a d -neighbor is planted at a random position. In this study, a generator was run to produce a dataset of PMS problem instances with this configuration, and this dataset was then used to evaluate the algorithms. Furthermore, a converter program was used to translate the dataset into FASTA format in order to execute qPMS9.

Implementation

The EMS-GT2 maintains a 4^l motif search space that is represented by an array of bits for space efficiency. The previous speedup technique (used in EMS-GT) exploits this manner of representing the search space by generating the neighborhood of an l -mer by blocks instead of per bit. In this study, we took advantage of this block-processing approach in the testing of candidate motifs. The Test phase checks if a candidate motif c is in the remaining $n - n'$ sequences by comparing if there is at least one l -mer in each sequence that is within d -distance from c . If a candidate motif x is eliminated for failing to have a d -neighbor in some input sequence S_i , then it is possible to reduce the testing for another candidate motif y on the same sequence S_i if y is within the same k -block as x .

Parameter Fine Tuning

The EMS-GT2 defines an integer value n' ($1 < n' < n$) that divides the dataset into two set of sequences. The first n' sequences are used in the Generate phase while the remaining are assigned to the Test phase. Previous experimentations [3] showed that it is efficient for the algorithm to set the value of n' to 10. Technically, n' dictates how big is the size of the set of

candidate motifs C to be evaluated if they are in the remaining $n - n'$ sequences. In line with this, we run an experimentation that records the average runtime of the algorithm with the speedup techniques over 5 tests and having different values for n' . The

values for n' range from 5 to 10 in this experiment, since we only want to make the candidate motif set large enough for our speedup technique to take effect.

TABLE 1
 RUNTIME EVALUATION OF THE ALGORITHM WITH SPEEDUP OVER DIFFERENT n' VALUES

Sequence	(9, 2)	(11, 3)	(13, 4)	(15, 5)	(17, 6)
5	0.07 s	0.42 s	2.37 s	17.86 s	148.62 s
6	0.06 s	0.29 s	1.54 s	12.64 s	116.01 s
7	0.05 s	0.21 s	1.01 s	10.70 s	119.94 s
8	0.05 s	0.18 s	0.84 s	10.38 s	119.82 s
9	0.03 s	0.14 s	0.75 s	11.01 s	132.10 s
10	0.03 s	0.13 s	0.76 s	11.90 s	146.10 s

Table 1 shows the ideal n' value for (l, d) -challenge instances. For (l, d) instances where $l \leq 11$, the ideal value for n' is still 10. This is true due to the efficiency of the Generate phase in small instances of the problem. For (13, 4), (15, 5) and (17, 5), different n' were used in the evaluation which are 9, 8 and 7 respectively.

Evaluation

For evaluation of the algorithms, we compared the EMS-GT2 to EMS-GT and the algorithm qPMS9. We used the challenging (l, d) instances defined in [13] [18]. The instances used in the evaluation are the following: (9, 2), (11, 3), (13, 4), (15, 5) and (17, 6).

RESULTS

This section describes EMS-GT2 with the two proposed speedup techniques and briefly discusses the observation where

the idea for the improvement originated. Runtime evaluation of EMS-GT2 is also discussed in this section.

Faster Candidate Motif Elimination through Block Processing

The EMS-GT algorithm tests candidate motifs in a brute-force approach. A candidate motif is tested by checking if it has at least one d -neighbor in each of the remaining $n - n'$ sequences. In testing a candidate motif c , if there is a sequence S_i in the remaining $n - n'$ sequences where c does not have any d -neighbor, then candidate motif c is automatically eliminated. In our implementation, the search space is represented by a compressed bit array and the l -mers are enumerated alphabetically. L -mers that are near each other do not differ that much. We used this observation in improving the way the algorithm tests the candidate motifs (see Figure 2 for the illustration of the search space).




Fig. 2. First 8 rows of the 4^5 search space with random flag values

In EMS-GT, each testing of candidate motif is independent of each other. We propose here a new speedup technique that processes these candidate motifs by blocks. We first partition the search space by blocks containing 4^k l -mers. This results into l -mers that share the same $(l - k)$ -prefix characters

where $2 < k < l$. Since every row in the search space represents exactly 32 l -mers (32-bit integers), the height of every block is computed using $(4^k/32)$ (see Figure 3 for an example of a partitioned search space). Additionally, any two l -mers within a block had at most k hamming distance value between them.

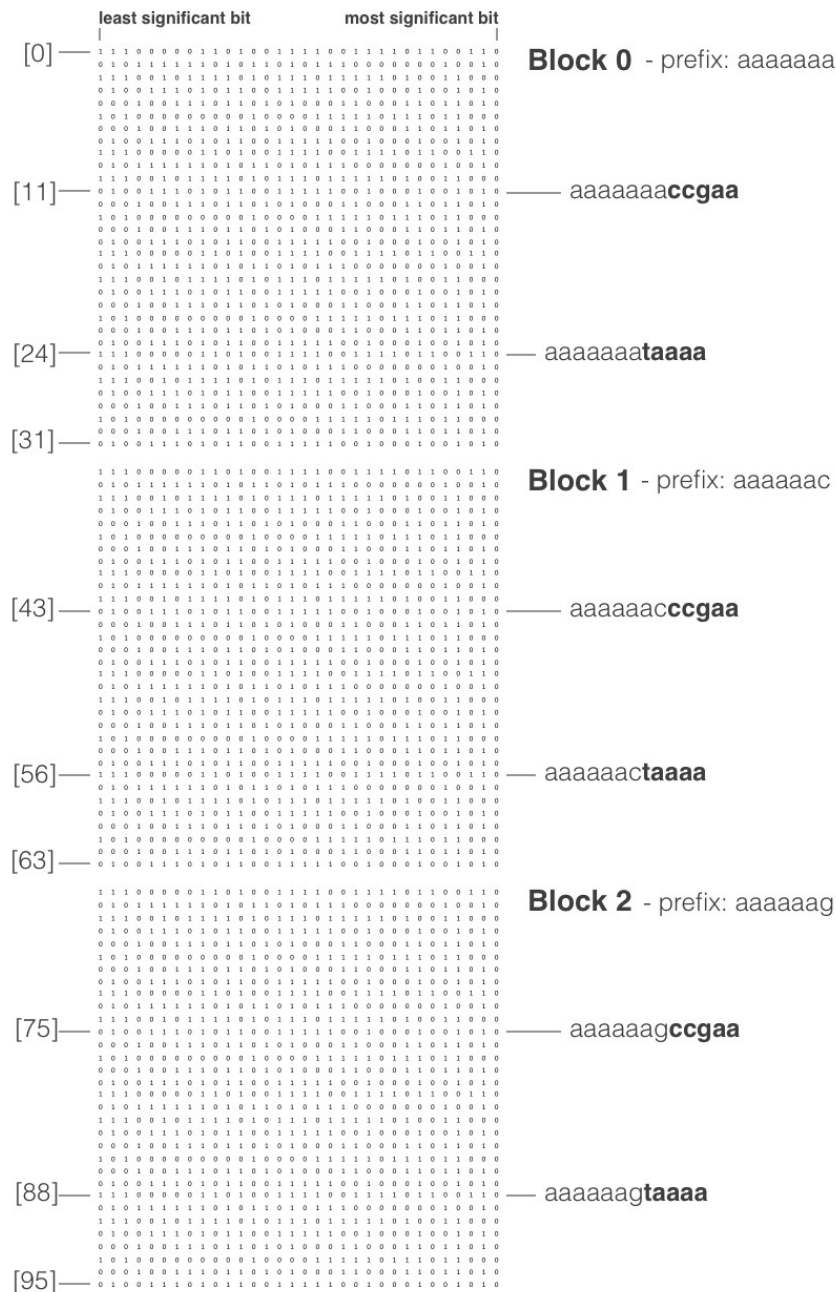


Fig. 3. Illustration of the block partitioning of a 4^{12} search space. There are 4^5 l -mers in each block and has a height of 32. The illustration shows the first 3 blocks only

We process the testing of candidate motifs now by blocks. If candidate motifs x and y are within a block and x has been eliminated as a candidate motif in sequence S_i ($n' \leq i \leq n$), we can filter out l -mers $z \in S_i$ where $dH(x, z) > d + k$. We collect the remaining l -mers in S_i and use them for testing the remaining candidate motifs in the block along with the other l -mers in the remaining sequences in $\{S'_n, S'_n + 1, \dots, S_n\} \{S_i\}$. The theorem below formalizes the main property used in this speedup technique.

Theorem 1. Let x and y be l -mers in a block in the search space containing 4^k l -mers. Let d be the number of allowed mutations in the problem instance. Let z be another l -mer. If $dH(x, z) > (k + d)$ then $dH(y, z) > d$, and therefore z is not in $N(y, d)$.

Proof. Using proof by contradiction, we first suppose that $dH(y, z) \geq d$. Since the l -mers x and y belong to the same block, then $dH(x, y) \leq k$. We use these bounds in the triangle of inequality $dH(x, z) \leq dH(x, y) + dH(y, z)$ to derive the result $dH(x, z) \leq k + d$. This result contradicts the given condition that $dH(x, z) > (k + d)$. Hence, we are sure that $dH(y, z) > d$.

The k -value affects the number of l -mers that is filtered in a sequence. The lower its value, the larger the number of filtered l -mers and faster candidate motif testing will be. But since k also affects the number of l -mers in a block, the lower its value, the fewer the candidate motifs that might benefit from

the speedup technique. In our implementation, we use $k = 5$ where every block has 32 rows of 32 bit flags representing a total of 4^5 l -mers.

Pre-Computation of Mismatch Values

The EMS-GT2 uses the hamming distance computation heavily during the Test phase. As discussed earlier for the original EMS-GT, the hamming distance of two binary represented l -mers, can be efficiently computed using the Boolean operator XOR. In EMS-GT2, instead of repeatedly counting these nonzero pairs of bits every time we compute the hamming distance, we use a pre-computed lookup table to help reduce computational time. A naive pre-computation of these nonzero pair counts for all possible l -mers (each represented by $2l$ bit values) will introduce an unacceptable overhead computation time when l is sufficiently large, i.e., when $l \geq 10$ (based on actual runs on our current machine configurations). A more efficient approach is to pre-compute only up to l -mers of length $l' < l$, which requires $b = 2l'$ number of bits. Then we determine the hamming distance by looking up the nonzero counts in the XOR results, b number of bits at a time, as described in Algorithm 1. In our experimentation the maximum required bits to represent l -mers are 34 bits (for (17, 6)-instance). Given this, we pre-compute up to 18 bits ($l'=9$) values only and use the lookup table twice for the computation of the actual hamming distance between any given pair of l -mers.

Algorithm 1 HAMMING DISTANCE COMPUTATION USING PRE-COMPUTED MISMATCH VALUES

Input: l -mer mappings u and v and
 MC \triangleright array of pre-computed count of mismatch positions

Output: Hamming distance $d_H(u, v)$

```

1:  $d_H(u, v) \leftarrow 0$ 
2:  $z \leftarrow u \oplus v$ 
3: while  $z > 0$  do
4:    $l \leftarrow z \& ((1 \ll 18) - 1) \triangleright$  process 18 bits at a time
5:    $d_H(u, v) \leftarrow d_H(u, v) + MC[l]$ 
6:    $z \leftarrow z \gg \gg 18 \triangleright$  shift 18 bits to the right
7: end while
8: return  $d_H(u, v)$ 

```

Algorithm 1: Shows the improved way of computing the mismatch value for the Hamming distance computation

Performance of EMS-GT with Speedup Techniques

The EMS-GT2, EMS-GT, and qPMS9 were evaluated in terms of actual runtime on an Intel Xeon, 2.10 Ghz machine. The performance of each algorithm was averaged over

20 synthetic datasets for each (l, d) -challenge instance where $l \leq 17$. Table 2 shows the runtime results between EMS-GT2 vs. EMS-GT while Table 3 shows the runtime results between the EMS-GT2 vs. the qPMS9.

TABLE 2
EMS-GT AND EMS-GT2 RUNTIME EVALUATION

(l, d)	EMS-GT	EMS-GT2	% Speedup
(9, 2)	0.04 s	0.05 s	-
(11, 3)	0.17 s	0.26 s	-
(13, 4)	1.03 s	0.82 s	20.3%
(15, 5)	12.39 s	.43 s	15.8%
(17, 6)	143.87 s	111.22 s	22.6%

The additional speedup techniques become more effective as the l value in the (l, d) -instance grows. For every (l, d) -challenge instance mentioned where $l \geq 13$, EMS-GT2 has improved the runtime over the EMS-GT by at least 15%. Unfortunately, the speedup techniques in EMS-GT2 failed to compensate for their additional overhead computations in both (9, 2) and (11, 3) challenge instances and failed to improve the overall runtime of the implementation. Previous implementations of the original EMS-GT failed to beat qPMS9 in

(17, 6)-challenge instance. The proposed EMS-GT2 not only produced improved runtimes but also succeeded in beating the qPMS9 in this challenge instance. The improved EMS-GT is now faster than the state-of-the-art qPMS9 in all of the (l, d) -challenge instances where $l \leq 17$. Even though the implementation of EMS-GT can only run in (l, d) -challenge instances where $l \leq 17$ because of computer memory constraint, studies have shown that the typical length of motifs is around 10 Base Pairs (BP) [4] anyway.

TABLE 3
EMS-GT2 AND QPMS9 RUNTIME EVALUATION

(l, d)	qPMS9	EMS-GT2	% Speedup
(9, 2)	0.60 s	0.05 s	91.6%
(11, 3)	1.26 s	0.26 s	79.3%
(13, 4)	4.58 s	0.82 s	82.0%
(15, 5)	25.73 s	10.43 s	59.4%
(17, 6)	123.17 s	111.22 s	9.7%

CONCLUSION AND RECOMMENDATIONS

We have presented EMS-GT2, an improved exact solution for the planted motif search problem. EMS-GT2 was able to efficiently test candidate motifs within the same block by filtering out l -mers using a property of the search space array that we have discovered and proven in this paper. The previous implementation of EMS-GT already outperforms the state-of-the-art algorithm qPMS9 in (l, d) -challenge instances

(9, 2), (11, 3), (13, 4) and (15, 5) but failed to beat qPMS9 in (17, 6)-challenge instance. EMS-GT2 improved the original algorithm's performance on (13, 4), (15, 5) and (17, 6) and was able to beat qPMS9 in all (l, d) -challenge instances where $l \leq 17$.

Declaration of Conflicting Interests

No conflicts of interest are present in the current study.

REFERENCES

- [1] S. Rajasekaran, S. Balla and C. H. Huang, "Exact algorithms for planted motif problems," *Journal of Computational Biology*, vol. 12, no. 8, pp. 1117-1128, 2005.
- [2] P. A. Evans, A. D. Smith and H. T. Wareham, "On the complexity of finding common approximate substrings," *Theoretical Computer Science*, vol. 306, no. 1, pp. 407-430, 2003.
- [3] M. C. I. D. Sia, J. Q. Nabos and P. L. Fernandez, "An efficient exact solution for the (l, d) -planted motif problem," in *8th AUN/SEED-Net Regional Conference on Electrical and Electronics Engineering*, Manila, Philippines, Nov. 16-17, 2015.
- [4] A. J. Stewart, S. Hannenhalli and J. B. Plotkin, "Why transcription factor binding sites are ten nucleotides long." *Genetics*.

vol. 192, no. 3, pp. 973-985, 2012.

- [5] C. E. Lawrence, S. F. Altschul, M. S. Boguski, J. S. Liu, A. F. Neuwald and J. C. Wootton, "Detecting subtle sequence signals: A Gibbs sampling strategy for multiple alignment," *Science*, vol. 262, no. 5131, pp. 208-214, 1993.
- [6] T. L. Bailey, N. Williams, C. Mischel and W. W. Li, "MEME: Discovering and analyzing DNA and protein sequence motifs," *Nucleic Acids Research*, vol. 34, pp. W369-W373, 2006.
- [7] H. Buhler and M. Tompa, "Finding motifs using random projections," in *proceedings of the Fifth Annual International Conference on Computational Biology*, New York, NY, 2001, pp. 69-76.
- [8] H. Huo, Z. Zhao, V. Stojkovic and L. Liu, "Combining genetic algorithm and random projection strategy for (l, d) -motif discovery," in *Fourth International Conference on Bio-Inspired Computing*, Beijing, China, Oct. 16-19, 2009. pp. 1-6.
- [9] P. A. Pevzner and S. H. Sze, "Combinatorial approaches to finding subtle signals in DNA sequences," in *International Conference on Intelligent Systems for Molecular Biology*, San Diego, CA, Aug. 19-23, 2000.
- [10] U. Keich and P. Pevzner, "Finding motifs in the twilight zone," *Bioinformatics*, vol. 18, no. 10, pp. 1374-1381, 2002.
- [11] A. Price, S. Ramabhadran and P. A. Pevzner, "Finding subtle motifs by branching from sample strings," *Bioinformatics*, vol. 19, no. 2, pp. 149-155, 2003.
- [12] G. Z. Hertz and G. D. Stormo, "Identifying DNA and protein patterns with statistically significant alignments of multiple sequences," *Bioinformatics*, vol. 15, no. 7, pp. 563-577, 1999.
- [13] J. Davila, S. Balla and S. Rajasekaran, "Fast and practical algorithms for planted (l, d) motif search," *IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB)*, vol. 4, no. 4, pp. 544-552, 2007.
- [14] H. Dinh, S. Rajasekaran and V. K. Kundeti, "PMS5: An efficient exact algorithm for the (l, d) -motif finding problem," *BMC Bioinformatics*, vol. 12, no. 1, pp. 01-10, 2011.
- [15] S. Bandyopadhyay, S. Sahni and S. Rajasekaran, "PMS6: A fast algorithm for motif discovery," *International Journal of Bioinformatics Research and Applications 2*, vol. 10, no. 4-5, pp. 369-383, 2014.
- [16] H. Dinh, S. Rajasekaran and J. Davila, "qPMS7: A fast algorithm for finding (l, d) -motifs in DNA and protein sequences," *PLoS One*, vol. 7, no. 7, pp. 01-08, 2012.
- [17] M. Nicolae and S. Rajasekaran, "Efficient sequential and parallel algorithms for planted motif search," *BMC Bioinformatics*, vol. 15, no. 1, pp. 01-10, 2014.
- [18] M. Nicolae and S. Rajasekaran, "qPMS9: An efficient algorithm for quorum planted motif search," *Scientific Reports*, vol. 5, pp. 01-017, 2015.
- [19] A. M. Carvalho, A. T. Freitas, A. L. Oliveira and M. F. Sagot, "A highly scalable algorithm for the extraction of cis-regulatory regions," in *3rd Asia-Pacific Bioinformatics Conference*, Singapore, Jan. 17-21, 2005.
- [20] N. Pisanti, A. M. Carvalho, L. Marsan and M. F. Sagot, "RISOTTO: Fast extraction of motifs with mismatches," in *Latin American Symposium on Theoretical Informatics*, Valdivia, Chile, March 20-24, 2006, pp. 757-768.
- [21] M. Sagot, "Spelling approximate repeated or common motifs using a suffix tree," in *Lecture Notes in Computer Science*, C. L. Lucchesi and A. V. Moura, Eds. Berlin, Germany: Springer.
- [22] E. Eskin and P. A. Pevzner, "Finding composite regulatory patterns in DNA sequences," *Bioinformatics*, vol. 18, no. 1, pp. S354-S363, 2002.
- [23] F. Y. Chin and H. C. Leung, "Voting algorithms for discovering long motifs," in *3rd Asia-Pacific Bioinformatics Conference*, Singapor, Jan. 17-21, 2005, pp. 261-271.
- [24] N. S. Dasari, R. Desh and M. Zubair, "An efficient multicore implementation of planted motif problem," in *International Conference on High Performance Computing and Simulation (HPCS)*, Caen, France, June 28- July 2, 2010, pp. 9-15.

— This article does not have any appendix. —