

Identification of Transcription Factor Binding Sites in Promoter Regions by Modularity Analysis of the Motif Co-Occurrence Graph

Alexandre P. Francisco, Arlindo L. Oliveira, and Ana T. Freitas

INESC-ID/IST, Technical University of Lisbon, Portugal
{aplf,aml,atf}@inesc-id.pt

Abstract. Many algorithms have been proposed to date for the problem of finding biologically significant motifs in promoter regions. They can be classified into two large families: combinatorial methods and probabilistic methods. Probabilistic methods have been used more extensively, since their output is easier to interpret. Combinatorial methods have the potential to identify hard to detect motifs, but their output is much harder to interpret, since it may consist of hundreds or thousands of motifs. In this work, we propose a method that processes the output of combinatorial motif finders in order to find groups of motifs that represent variations of the same motif, thus reducing the output to a manageable size. This processing is done by building a graph that represents the co-occurrences of motifs, and finding communities in this graph. We show that this innovative approach leads to a method that is as easy to use as a probabilistic motif finder, and as sensitive to low quorum motifs as a combinatorial motif finder. The method was integrated with two combinatorial motif finders, and made available on the Web.

1 Introduction

An important open problem in computational biology is related with the accurate identification of biologically meaningful nucleotide sequences in promoter regions, that correspond to loci of attachment of transcription factors. These well conserved regions are usually referred to as consensus sequences or motifs. Motif finding is the problem of discovering these motifs without prior knowledge of their characteristics. Motif finding has been the subject of intense research and literally hundreds of papers have been published on this topic [1].

Currently available methods for motif finding can roughly be classified in two main classes: probabilistic and combinatorial.

Probabilistic methods have been extensively used, and they identify very well the strong signals present in the data, i.e., motifs that occur in a large fraction of the sequences. They have difficulties identifying weaker signals, that correspond to motifs that are present only in a subset of sequences, possibly superimposed with stronger signals.

Combinatorial methods, on the other hand, when executed with the right parameters, can identify both strong and weak signals. They suffer, however,

from a significant drawback. When executed with parameters that allow them to identify motifs that are present in only a small fraction of the sequences, they will deluge the user with a large, possibly huge, number of motifs, that correspond to many variations of the motifs of interest. In fact, since motifs are not perfectly conserved, many variations of the most common motifs will be reported by a combinatorial motif finder, since these variations will appear in a significant fraction of the sequences.

In this work, we propose a method for the identification of motifs that combines the advantages of probabilistic motif finders (easy to use, no parameters required) and combinatorial motif finders (ability to identify even the weaker signals) while avoiding the disadvantages of both.

We achieve this by post-processing the results of combinatorial motif finders, and identifying the motifs that are variations of the same signal. These motifs are then combined into a composed representation, and a position weight matrix (PWM) is generated for that set of motifs. The identification of the motifs that are, in reality, variations of the same motif, is done by computing the modules (or communities) in a graph. This graph has one node for each motif found, and one edge between two motifs if they have significant occurrence overlap.

2 Basic concepts and related work

2.1 Motif finders

The most used probabilistic algorithms for motif finding are based on the application of the Expectation-Maximization [2] method (EM) like PROJECTION [3] and MEME [4] or its stochastic counterpart, Gibbs sampling [5] used by ALIGNACE [6], BIOPROSPECTOR [7] and GIBBSDNA [5]. These methods use a two-phase iterative procedure where, in the first step the likeliest occurrences of the motif are identified, based on a model computed in the previous iteration. The second step adjusts the model for the motif (usually a position weight matrix) based on the occurrences determined in the previous step. In the first iteration the parameters of the initial model are usually set randomly. This iterative procedure converges usually in a rapid way to motifs that are present, possibly with mutations, in a large fraction of input sequences. They report their results in the form of PWMs, that represent directly the obtained model.

Combinatorial methods, which extract motifs consisting of plain nucleotide sequences work by enumerating the possible patterns, either explicitly or implicitly [8, 9], and counting their quorum. Consider a set of sequences under analysis, $\mathcal{S} = \{S_1, S_2, \dots, S_t\}$. The objective is to find motifs within a range of lengths $l_{\min}, \dots, l_{\max}$, which occur on $q \leq t$ of the sequences in \mathcal{S} with at most e mismatches, i.e., having at most e nucleotide substitutions. Algorithms that look for complex motifs have also been proposed [10, 11]. Complex motifs are built of two or more simple motifs, spaced by a number of bases that falls within a specific range.

For this work we selected the combinatorial motif finders MUSA [12] and RISO [10]. MUSA is an algorithm that does not require the user to specify

parameters (such as box lengths and distances between boxes) in order to extract motifs. This method relies on a biclustering algorithm that operates on a matrix of co-occurrences of small motifs. Requiring as input a list of gene promoter sequences, MUSA returns the list of structured or simple motifs found, ordered by their p-value, and their quorum. RISO is a complex motif extraction tool. It searches for complex motifs with certain characteristics specified by the user, through the assignment of a set of parameters such as the number and sizes of the boxes that form the structured motif, the distances between them and the minimum quorum expected. RISO returns the list of motifs found and their corresponding quorums.

2.2 Motif clustering

The idea of finding groups, or clusters, of motifs, in order to simplify the binding site studies and reduce the redundancy of the patterns found by motif finders is not new, and, indeed, has been proposed independently. Examples of tools that perform this clustering are MatAlign [13] and Stamp [14].

Although there are differences in the implementation, these and other existing methods work by defining a distance between two motifs and applying standard clustering methods to find motifs with similar patterns. The distance is typically obtained using dynamic programming to compute the best alignment between two motifs.

While this approach works well in some cases, it has some strong limitations. In particular, this approach is not able to identify that two motifs are part of the same pattern if they are poorly aligned, even if they represent different parts of the same, larger, motif. For instance, motifs ACCGTG and TGATTT may be frequent because the larger motif ACCGTGATTT is frequent, but no significant alignment will be found between motifs ACCGTG and TGATTT, specially if, for some reason, the larger motif is not identified.

The approach we propose avoids this difficulty by ignoring the actual pattern of the motifs, and considering only the sequences and positions where they occur. A significant amount of co-occurrence means that two motifs are linked, even though they may not be very similar.

For this, we need a method that finds communities in graphs of motifs, i.e., groups of motifs that are tightly linked by many co-occurrences.

2.3 Finding communities in graphs

Many algorithms have been developed to tackle the problem of finding communities in graphs [15–18]. Probably the best-known is the one proposed by Girvan and Newman [15] based on the betweenness centrality measure which runs in $O(|E|^2|V|)$ time, or $O(|V|^3)$ time for sparse graphs. However, it is no longer the most efficient and effective clustering algorithm. More recently Newman [17] has proposed a faster algorithm based on the greedy optimization of the modularity [16] which is substantially faster. It runs in $O((|E| + |V|)|V|)$, or $O(|V|^2)$ on sparse graphs. However, the running time of this algorithm can be improved by

exploiting some properties of the optimization problem and using more sophisticated data structures. Thus, Clauset *et al.* [18] proposed a greedy algorithm which runs in $O(|E|d \log |V|)$, where d is the depth of the “dendrogram” which describes the community structure. On sparse graphs with a hierarchical community structure this algorithm runs on average in $O(|V| \log^2 |V|)$ time. In what follows, we name this algorithm as *CNM (Clauset-Newman-Moore) algorithm*.

The concept of modularity is central to this problem [16]. Modularity is a property of the graph and of a specific division of the graph into communities. It measures the quality of the division by evaluating the number of edges within communities and the number of edges that connect vertices in different communities. Suppose the vertices are divided into k communities and let $1 \leq c_m \leq k$ denote the community where vertex $m \in V$ belongs. The *adjacency matrix* A of G and the *degree* d_m of a vertex $m \in V$ are respectively defined as

$$A_{mn} = \begin{cases} 1 & \text{if } (m, n) \in E, \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad d_m = \sum_{n \in V} A_{mn}. \quad (1)$$

We define the *modularity* Q of G with respect to the given division as

$$Q = \frac{1}{2|E|} \sum_{m, n \in V} \left[A_{mn} - \frac{d_m d_n}{2|E|} \right] \delta(c_m, c_n), \quad (2)$$

where the δ -function is such that $\delta(i, j) = 1$ if $i = j$ and $\delta(i, j) = 0$ otherwise. We note that the above sum runs over all possible pairs of vertices. Therefore, each edge is summed twice. If we split the sum in two terms,

$$\frac{1}{2|E|} \sum_{m, n \in V} A_{mn} \delta(c_m, c_n) \quad \text{and} \quad \frac{1}{2|E|} \sum_{m, n \in V} \frac{d_m d_n}{2|E|} \delta(c_m, c_n), \quad (3)$$

then the first term is the fraction of edges that fall within the communities, and the second term is the expected fraction of edges within the communities if the edges were randomly distributed respecting the vertices degrees. In particular, if the edges were randomly placed as mentioned, $d_m d_n / |E|$ is the probability of the existence of an edge between vertices m and n .

Thus, modularity measures the fraction of edges that connect vertices in the same component minus the expected value of the same quantity in a graph with the same components but random connections between the vertices [16]. Values near 1, the maximum value of Q , indicate strong community structure. Typically, values for graphs underlying common networks with known community structure are in the range from 0.3 to 0.7.

The CNM algorithm operates by finding the changes in Q which result from merging each pair of communities. It chooses the largest of such possible changes in a greedy way and performs the merging. Let ΔQ_{ij} be the change in Q that results from merging the communities i and j . Initially each vertex $m \in V$ is a community, by equation 2,

$$Q = -\frac{1}{2|E|} \sum_{m \in V} \frac{d_m d_m}{2|E|}. \quad (4)$$

And, for each community i and for each pair of communities i, j , we set

$$A_i = \frac{d_m}{2|E|}, \quad \text{and} \quad \Delta Q_{ij} = \begin{cases} \frac{1}{|E|} - 2A_i A_j & \text{if } i, j \text{ are connected,} \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

where m is the unique vertex in the community i . Thus, the main loop of the CNM algorithm consists in

1. select the largest ΔQ_{ij} and increment Q accordingly,
2. join the corresponding communities and, assuming that community i is merged into community j , update the ΔQ and the A values as follows:

$$\Delta Q_{jk} = \begin{cases} \Delta Q_{ik} + \Delta Q_{jk} & \text{if } k \text{ is connected to } i \text{ and } j, \\ \Delta Q_{ik} - 2A_j A_k & \text{if } k \text{ is connected to } i \text{ but not to } j, \\ \Delta Q_{jk} - 2A_i A_k & \text{if } k \text{ is connected to } j \text{ but not to } i; \end{cases} \quad (6)$$

$$A_j = A_j + A_i. \quad (7)$$

3 Finding co-occurring motifs

As described, our method builds a motif relation graph and finds communities of motifs, i.e., subgraphs such that the density of edges within it is greater than the density of edges between its vertices and those outside it. Each community is then processed in order to obtain the associated Position Weight Matrix (PWM).

3.1 Building the relation graph

Let \mathcal{S} be the set of sequences and \mathcal{M} be the set of motifs found in \mathcal{S} . For each $m \in \mathcal{M}$ and $s \in \mathcal{S}$, let $l(m, s)$ be the list of positions in s where m occurs. We say that $m, n \in \mathcal{M}$ *overlap* in a sequence $s \in \mathcal{S}$ if $x \in l(m, s)$ and $y \in l(n, s)$ exist such that one of the following two conditions verifies:

$$x < y < x + |m|; \quad y \leq x < y + |n|. \quad (8)$$

In such case we say that motifs m and n overlap in s with a *shift* σ equal to $y - x$.

In our method, we assume that a minimum overlap $0 < \alpha_o \leq 1$ and a minimum quorum $0 < \alpha_q \leq 1$ are given as parameters. The quorum represents the fraction of the number of common sequences in which a given pair of overlapping motifs must occur to be considered. Therefore, given two motifs $m, n \in \mathcal{M}$ we define the *minimum shift* and the *maximum shift*, for m and n , respectively, as

$$\sigma_{\min} = \alpha_o \min\{|m|, |n|\} - |n| \quad \text{and} \quad \sigma_{\max} = |m| - \alpha_o \min\{|m|, |n|\}. \quad (9)$$

To ensure that m and n overlap in a given sequence $s \in \mathcal{S}$ with at least $\alpha_o \min\{|m|, |n|\}$ common characters, we must check that such overlap occurs with a shift σ such that $\sigma_{\min} \leq \sigma \leq \sigma_{\max}$.

The *relation graph* G is a tuple $\langle V, E \rangle$ where V and E are defined as follows. The set of vertices V is the set of motifs found, i.e., $V = \mathcal{M}$. The set of edges E contains every pair $(m, n) \in \mathcal{M} \times \mathcal{M}$ for which exists $\sigma \in \mathbb{N}$ such that $\sigma_{\min} \leq \sigma \leq \sigma_{\max}$ and

$$\alpha_q |\mathcal{S}'| \leq |\{s \in \mathcal{S}' : m, n \text{ overlap in } s \text{ with shift } \sigma\}|, \quad (10)$$

where $\mathcal{S}' \subseteq \mathcal{S}$ is the set of sequences in which both m and n occur.

3.2 Implementation of the CNM algorithm

In its original version, the CNM algorithm iterates until a negative ΔQ_{ij} is selected, and stops when all vertices belong to the same community. However, in our case, the relation graph G may not be connected, and the selection of a negative value is an admissible stop condition because after that Q can only decrease.

The bound $O(|E|d \log |V|)$ in the running time can only be achieved if advanced data structures are used [18]. In our implementation, we store the ΔQ_{ij} values in a red-black tree for each community. Additionally, we maintain these values in binary heaps. Therefore the insertion, the selection and the maximum extraction can be done in $O(\log |V|)$ [19]. We also use the well known union-find data structure [19] to track the vertices in each community. Because the relation graph G is, in general, sparse, the CNM algorithm complexity is almost linear in the number of motifs.

3.3 Computing and ranking the PWM of a community of motifs

By applying the CNM algorithm to the motif relation graph and choosing the partition which grants maximum modularity, we obtain a set of communities of motifs. The third and final step of our method consists in processing each of these communities and computing the PWM for each one.

Thus, let $\mathcal{C} \subseteq \mathcal{M}$ be a community of motifs found in the graph G . First, we align the motifs in \mathcal{C} , which is simple because we already know the relative shift from the graph construction, and we compute the length of the PWM for this community. Second, for each edge $(n, m) \in E$ and using the best shift for equation 10, i.e., the shift which maximizes the right side of equation 10, we update the corresponding columns of the PWM by checking the symbols in the sequences where the pair of motifs occur, i.e, the sequences in the set \mathcal{S}' in equation 10.

Each community gets assigned a p-value that correspond to the lowest p-value of the motifs in that community. This p-value is used to rank the communities and corresponding PWMs. For each community a quorum is also computed. This quorum corresponds to the average number of sequences that support each edge in the community structure.

The method developed was implemented in **C**, including the CNM algorithm and all data structures. The resultant tool was integrated with two motif finders,

MUSA and RISO , and made available through the DISCOVERER platform in the YEASTRACT database [20]. Given that the complexity of the CNM algorithm is almost linear for sparse graphs, the computation of the relation graph is the most computational demanding step of our method, taking $\Omega(|M|^2)$ time. For all test examples, which have at most 3000 statistically significant motifs, we were able to compute motif communities and corresponding PWMs in less than one minute in a common workstation.

4 Results

In this section we only present the results obtained, with the motif finder MUSA, for the first two datasets described in Table 1. More detailed results for both motif finders and for all datasets are available, as supplementary material, at <http://kdbio.inesc-id.pt/mat/isbra08>.

To test the ability of the proposed method to find relevant motif communities, four real biological datasets were used. These datasets correspond to different sets of promoter sequences of *Sacharomyces cerevisiae* genes.

For all datasets the MUSA algorithm was executed with the default parameter values: $\lambda = 4$, $\epsilon = 1$ and the quorum equal to 30%. The motifs reported were ranked in accordance with their statistical significance. Motifs that have a p-value smaller than 10^{-3} were considered statistically significant and selected for further processing. To build the relation graph, for these motifs, the default values of α_o , the minimum overlap, and α_q , the minimum quorum, were 0.4 and 0.7, respectively.

Table 1 summarizes the results obtained. For each dataset it shows the number of sequences (N. seq), that were used to search for over-represented motifs, the total number of motifs found by the motif finder (T. mot), the number of motifs considered for processing (N. mot), the total number of edges in the relation graph (N. edg), the number of communities identified (N. com) and the modularity value achieved (Modul). It is clear from this table that the method

Table 1. Datasets content and results statistics.

<i>Datasets</i>	<i>N.seq</i>	<i>T.mot</i>	<i>N.mot</i>	<i>N.edg</i>	<i>N.com</i>	<i>Modul</i>
DeRisi [21]	25	1647	204	299	89	0.80
Aft2p [22]	193	2176	2026	3397	202	0.84
Yap1p [23]	225	2150	2065	3541	168	0.84
2,4D [24]	486	2088	2045	3143	271	0.86

effectively reduces the number of motifs that need to be analyzed, by up to one order of magnitude.

The first dataset, identified as the DeRisi set, corresponds to a list of 25 genes that were up-regulated in response to the expression of a point mutation in the PDR1 gene, that encodes a transcription factor (TF) involved in Pleiotropic Drug Resistance in yeast [21]. Due to the experimental procedure used, this set

Table 2. The top 15 motifs reported by MUSA for the DeRisi dataset

<i>ID</i>	<i>Motif</i>	<i>Quorum</i>	<i>P – value</i>
1	TCCGTGGA	12 of 25	2.79106e-17
2	TCCACGGA	12 of 25	2.79106e-17
3	AAGA (17,19) TTTC	18 of 25	3.57327e-16
4	GAAA (17,19) TCTT	18 of 25	3.57327e-16
5	CCGT (1,3) GAAA	13 of 25	7.22785e-15
6	TTTC (1,3) ACGG	13 of 25	7.22785e-15
7	CCACGGA	14 of 25	8.57847e-15
8	TCCGTGG	14 of 25	8.57847e-15
9	AAAA (4,6) AAAT	25 of 25	1.31737e-14
10	ATTT (4,6) TTTT	25 of 25	1.31737e-14
11	CCACGGAA	11 of 25	1.45924e-14
12	TTCCGTGG	11 of 25	1.45924e-14
13	AACA (43,45) CCTC	11 of 25	2.40999e-12
14	GAGG (43,45) TGTT	11 of 25	2.40999e-12
15	CAAAAG (3,5) AAAT	9 of 25	4.18278e-12

correspond to a small and very well characterized set of genes where the Pdr1p binding site can be easily identified.

For this set, using the default input parameters and considering both strands, the MUSA algorithm identified 1647 over-represented motifs. From these, only 204 motifs were classified as statistically significant and considered for further processing. Table 2 presents the first 15 motifs reported by MUSA. In this list, motifs 1 and 7 correspond to instances of the TF binding site of interest. Although these motifs are well positioned in the motifs list, it is possible to verify that they are only present in at most half of the input sequences. This low quorum hides the real importance of this binding site, something that is not expected in this particular dataset where all the genes were up-regulated by Pdr1p.

By inspecting Table 2 it is also possible to note that there is a large number of complex motifs in the top 15 motifs reported. For space reasons, we will not describe, in this article, the way complex motifs are handled, although a trivial extension of the method exists and has been implemented.

Figure 1 shows the 14 motifs which contribute to the most significant community obtained. The PWM description of this community is also presented. By inspecting this figure it is possible to see that the quorum of this community is 97%. This value is very important since it reflects the real importance of this binding site in this set. The community quorum is also an important feature in the evaluation of the community importance. Figure 2 shows the PWMs logos for the two most significant motif communities identified.

If a search for documented TFs binding sites is performed in YEASTRACT database using the PWM of the first community, the best match will be one of the documented Pdr1p-binding sites. If this search is also performed for the second most important community, shown in Figure 2, again one of the documented Pdr1p-binding sites will be found. This second motif is similar to the first one but not included in the same community because of the difference in the central

Community 1 (14 motifs) (quorum stat: min 0.71 avg 0.97 max 1.00)				
---TCCGTGGA-	0	14	8	2.79106e-17
---TCCGTGG--	1	17	9	8.57847e-15
--TTCCGTGG--	2	12	9	1.45924e-14
----CCGTGG--	3	26	10	2.03787e-09
---TCCGTG---	4	18	9	2.16867e-07
--TTCCGT----	5	28	13	4.31347e-07
-----CGTGG--	6	45	16	1.13802e-05
-----CGTGGA-	7	23	10	5.08829e-05
----CCGTG---	8	33	13	0.000103709
-GTTCCG-----	9	15	8	0.000258935
---TCCGTC---	10	16	5	0.00026753
CCCTC-----	11	36	14	0.000360318
-----GTGGA-	12	48	17	0.000506259
-----GTGAA	13	21	8	0.000797191

A:	0.26	0.11	0.12	0.05	0.01	0.00	0.00	0.00	0.00	0.00	0.90	0.63
C:	0.12	0.28	0.10	0.04	0.97	1.00	0.00	0.00	0.01	0.00	0.03	0.03
G:	0.18	0.36	0.07	0.03	0.01	0.00	1.00	0.00	0.99	0.99	0.06	0.17
T:	0.44	0.25	0.71	0.88	0.01	0.00	0.00	1.00	0.00	0.01	0.01	0.17

Fig. 1. Depiction of the most significant community found for the DeRisi dataset: motifs alignment and PWM description.




Fig. 2. Depiction of the PWM logos for the first and second most important communities identified.

nucleotides: an GT switched for a AC. In fact, since MUSA searched for motifs in both strands, this second community roughly contains the reverse complement of the motifs contained in the first community. This community could be trivially merged with the first one, however when reporting only one of the communities we have found out that some users do not recognize the motif they were looking for. Thus, we decided to report both.

The second dataset, identified as Aft2p, includes 193 promoter sequences of Aft2p-target genes. This TF is involved in the regulation of iron homeostasis and associated oxidative stress [25]. There is evidence supporting the direct binding of Aft2p to the promoter region of many of these 193 target-genes [26]. However, for some of them, evidence of a regulatory association is indirect, coming from the comparison of gene transcript levels in the wild type and in a mutant devoid of AFT2 [22]. In this test case the list of genes considered includes direct and indirect targets of Aft2p and, due to this fact, this TF binding site is not expected to be a strong signal in the sequences.

From a total of 2176 motifs extracted by the motif finder, only 2026 were classified as statistically significant and considered for further processing. For this set, 202 communities were identified. Table 3 presents the top 15 motifs extracted by MUSA.

Table 3. The top 15 motifs reported by MUSA for the Aft2p dataset

<i>ID</i>	<i>Motif</i>	<i>Quorum</i>	<i>P – value</i>
1	TTTT (13,15)CACC	64 of 193	6.54141e-34
2	ACATAT	139 of 193	1.63126e-31
3	ATATGT	139 of 193	1.63126e-31
4	CACCC	165 of 193	6.43219e-31
5	GGGTG	165 of 193	6.43219e-31
6	GAAGAA	149 of 193	6.3155e-30
7	TTCTTC	149 of 193	6.3155e-30
8	GTATAT	124 of 193	1.36158e-29
9	CATATA	130 of 193	3.35912e-29
10	AAGAAG	145 of 193	1.04215e-28
11	TATTCT	145 of 193	1.31282e-28
12	CAAGAA	148 of 193	5.27289e-28
13	TTCTTG	148 of 193	5.27289e-28
14	GCACC	155 of 193	4.69138e-27
15	GGTGC	155 of 193	4.69138e-27

**Fig. 3.** Depiction of the PWM logo for the Aft2p-target genes binding site.

Figure 3 shows the PWM logo of the fourth most significant community found, that corresponds precisely to the documented Aft2p-binding site [22]. The binding site previously described for this TF allows some variability in the peripheral nucleotides. It is interesting to note that the identified community also exhibits a central conserved core region, the motif CACCC, flanked by less conserved peripheral nucleotides. This core motif, that correspond to motif 4 in Table 3, is statistically significant and is present in 85% of the input sequences. The correspondent community still presents a better quorum, 91%.

The first three most significant communities found for this dataset were compared with the documented TF binding sites described in the YEASTRACT database. The third most significant community was also associated with the documented Aft2p-binding site. The second most significant community matched the *TATA-box*, a well characterized core promoter element also expected to be a strong signal in this dataset.

For the first community found, the alignments obtained were very poor, suggesting that there is no documented TF binding site with such characteristics. To further investigate the existence of a similar TF binding site, the original PWM was trimmed. In this case, the best match was with the binding site of the Rap1p transcription factor. This TF is described as a DNA-binding protein involved in either activation or repression of transcription, depending on binding site context. However, the trimmed PWM aligned only with a short part of the Rap1p TF binding site. Although presumably not related with the documented

Aft2p binding site or other documented TF binding site, this motif can have an important biological meaning.

5 Discussion

In this paper we proposed a methodology that assembles a list of individual simple motifs into communities of motifs, leading to a simplified analysis of the motif finders results.

For the test-cases presented, the results show that this method is able to identify the most important motif communities. In fact this approach is very useful in reducing the number of motifs to be inspected, leading to a more tractable output, easier to interpret by humans. The PWM representation of the community highlights the motifs degeneracy, being more informative than the consensus representation usually reported by combinatorial motif finders. The quorum of the community reveals the real importance of the motifs in the dataset.

Compared to the first test-case, the results obtained for the second test-case seem less precise. However the results achieved are still remarkably important. The post-processing of the motif finder results allowed the identification of the Aft2p binding site and suggested new putative binding sites. The third and fourth test-cases show that even for more noisy datasets this approach can provide interesting clues on how transcription factors interact with their target genes.

References

1. Sandve, G., Drablos, F.: A survey of motif discovery methods in an integrated framework. *Biology Direct* **1**(1) (2006) 11
2. Segal, E., Sharan, R.: A discriminative model for identifying spatial cis-regulatory modules. *Journal of Computational Biology* **12**(6) (2005) 822–834
3. Buhler, J., Tompa, M.: Finding motifs using random projections. *Journal of Computational Biology* **9**(2) (2002) 225–242
4. Bailey, T., Elkan, C.: Fitting a mixture model by expectation maximization to discover motifs in biopolymers. In: *Proceedings of the Second International Conference on Intelligent Systems for Molecular Biology*. (1994) 28–36
5. Lawrence, C.E., Altschul, S.F., Boguski, M.S., Liu, J.S., Neuwald, A.F., Wootton, J.C.: Detecting subtle sequence signals: A Gibbs sampling strategy for multiple alignment. *Science* **262**(5131) (1993) 208–214
6. Roth, F.P., Hughes, J.D., Estep, P.W., Church, G.M.: Finding DNA regulatory motifs within unaligned noncoding sequences clustered by whole-genome mRNA quantitation. *Nature Biotechnology* **16** (1998) 939–945
7. Liu, X., Brutlag, D.L., Liu, J.S.: BioProspector: discovering conserved DNA motifs in upstream regulatory regions of co-expressed genes. *Pacific Symposium on Biocomputing* **6** (2001) 127–138
8. Sagot, M.: Spelling approximate repeated or common motifs using a suffix tree. *Latin* **98** (1998) 111–127

9. Pevzner, P.A., Sze, S.H.: Combinatorial approaches to finding subtle signals in DNA sequences. *Proceedings of the International Conference on Intelligent Systems for Molecular Biology* **8** (2000) 269–278
10. Carvalho, A., Freitas, A., Oliveira, A., Sagot, M.F.: An efficient algorithm for the identification of structured motifs in DNA promoter sequences. *IEEE Transactions on Computational Biology and Bioinformatics* **3**(2) (2006) 126–140
11. Marsan, L., Sagot, M.: Algorithms for extracting structured motifs using a suffix tree with an application to promoter and regulatory site consensus identification. *Journal of Computational Biology* **7**(3-4) (2000) 345–362
12. Mendes, N., Casimiro, A., Santos, P., S-Correia, I., Oliveira, A., Freitas, A.: MUSA: A parameter free algorithm for the identification of biologically significant motifs. *Bioinformatics* **22** (2006) 2996–3002
13. Kankainen, M., Loytynoja, A.: MATLIGN: a motif clustering, comparison and matching tool. *BMC Bioinformatics* **8**(1) (2007) 189
14. Mahony, S., Benos, P.: STAMP: a web tool for exploring DNA-binding motif similarities. *Nucleic Acids Research* (2007)
15. Girvan, M., Newman, M.E.J.: Community structure in social and biological networks. *Proceedings of the National Academy of Sciences* **99** (2002) 7821
16. Newman, M.E.J., Girvan, M.: Finding and evaluating community structure in networks. *Physical Review E* **69** (2004) 026113
17. Newman, M.E.J.: Fast algorithm for detecting community structure in networks. *Physical Review E* **69** (2004) 066133
18. Clauset, A., Newman, M.E.J., Moore, C.: Finding community structure in very large networks. *Physical Review E* **70** (2004) 066111
19. Cormen, T.H., Leiserson, C.E., Rivest, R.L., Stein, C.: *Introduction to Algorithms*. MIT Press (2001)
20. Teixeira, M.C., Monteiro, P., Jain, P., Tenreiro, S., Fernandes, A.R., Mira, N.P., Alenquer, M., Freitas, A.T., Oliveira, A.L., Sá-Correia, I.: The YEASTRACT database: a tool for the analysis of transcription regulatory associations in *saccharomyces cerevisiae*. *Nucleic Acids Research* **34** (2006) D446–D451
21. DeRisi, J., van den Hazel, B., Marc, P., Balzi, E., Brown, P., Jack, C., Goffeau, A.: Genome microarray analysis of transcriptional activation in multidrug resistance yeast mutants. *FEBS Letters* **470** (2000) 156–160
22. Courel, M., Lallet, S., Camadro, J., Blaiseau, P.: Direct activation of genes involved in intracellular iron use by the yeast iron-responsive transcription factor Aft2 without its paralog Aft1. *Molecular Cell Biology* **25**(15) (2005) 6760–6771
23. Cohen, B., Pilpel, Y., Mitra, R., Church, G.: Discrimination between paralogs using microarray analysis: application to the Yap1p and Yap2p transcriptional networks. *Molecular Biology of the Cell* **13**(7) (2002) 1608–1614
24. Teixeira, M., Fernandes, A., Mira, N., Becker, J., Sá-Correia, I.: Early transcriptional response of *Saccharomyces cerevisiae* to stress imposed by the herbicide 2, 4-dichlorophenoxyacetic acid. *FEMS Yeast Research* **6**(2) (2006) 230–248
25. Blaiseau, P., Lesuisse, E., Camadro, J.: Aft2p, a novel iron-regulated transcription activator that modulates, with Aft1p, intracellular iron use and resistance to oxidative stress in yeast. *Journal of Biological Chemistry* **276**(36) (2001) 34221–34226
26. Harbison, C., Gordon, D., Lee, T., Rinaldi, N., Macisaac, K., Danford, T., Hannett1, N., Tagne, J., Reynolds, D., Yoo, J., Jennings, E., Zeitlinger, J., Pokholok, D., Kellis, M., Rolfe, P., Takusagawa, K., Lander, E., Gifford, D., Fraenkel, E., Young, R.: Transcriptional regulatory code of a eukaryotic genome. *Nature* **431**(7004) (2004) 99–104

Identification of Transcription Factor Binding Sites in Promoter Regions by Modularity Analysis of the Motif Co-Occurrence Graph

Alexandre P. Francisco, Arlindo L. Oliveira, and Ana T. Freitas

INESC-ID/IST, Technical University of Lisbon, Portugal
apl@ist.utl.pt, {aml,atf}@inesc-id.pt

1 Supplementary material

In this document we present the results obtained, with the motif finders MUSA and RISO, for all datasets described in Table 1.

The motifs extracted by both algorithms were ranked in accordance with their statistical significance. Motifs that have a p-value smaller than 10^{-3} were considered statistically significant and selected for further processing.

To build the relation graph, for the statistically significant motifs, we used the following default values for α_o , the minimum overlap, and for α_q , the minimum quorum:

- MUSA algorithm: 0.4 and 0.7, respectively,
- RISO algorithm: 0.6 and 0.9, respectively.

1.1 Datasets

To test the proposed method, four real biological datasets were used. These datasets correspond to different sets of promoter sequences of *Sacharomyces cerevisiae* genes. Table 1 presents the datasets, showing the number of sequences (N. seq) in each one.

Table 1. Datasets content.

<i>Datasets</i>	<i>N.seq</i>
DeRisi [1]	25
Aft2p [2]	193
Yap1p [3]	225
2,4D [4]	486

The first dataset, identified as the DeRisi set, corresponds to a list of 25 genes that were up-regulated in response to the expression of a point mutation in the PDR1 gene, that encodes a transcription factor (TF) involved in Pleiotropic Drug Resistance in yeast [1]. Due to the experimental procedure used, this set of Pdr1p-target genes correspond to a small and very well characterized set of genes where the Pdr1p binding site can be easily identified. This site will correspond to a set of motifs that are over-represented in the promoter sequences.

The second dataset, Aft2p, includes 193 promoter sequences of Aft2p-target genes. This TF is involved in the regulation of iron homeostasis and associated oxidative stress [5]. There

is evidence supporting the direct binding of Aft2p to the promoter region of many of these 193 target-genes [6]. However, for some of them, evidence of a regulatory association is indirect, coming from the comparison of gene transcript levels in the wild type and in a mutant devoid of AFT2 [2]. In this test case the list of genes considered includes direct and indirect targets of Aft2p and, due to this fact, this TF binding site is not expected to be a strong signal in the sequences.

The third dataset, Yap1p, correspond to a set of 225 promoter sequences of the genes involved in stress response and showing Yap1p-dependent expression [3]. Yap1p is a bZIP transcription factor involved in multiple stress responses and controlling a set of genes involved in detoxifying the effects of reactive oxygen species. Since Yap1p are implicated in the response to cellular stresses, we selected the expression measurements for cells treated with the oxidizing agent hydrogen peroxide (H_2O_2). The dataset was created including the genes whose transcript level increased at least 2 fold in response to the stress induced and that decreased its expression at least 25% in the *yap1* Δ mutant. This gene list includes direct and indirect targets of Yap1p.

The fourth and last set, identified as the 2,4D set, includes 446 promoter sequences of the genes that have shown significant response to the herbicide 2,4-dichlorophenoxyacetic acid (2,4-D) [4]. This dataset was created from the treated microarray data, including the yeast genes whose transcript level increased at least 2 fold in response to the stress induced by 2,4-D. This particular set mimics the typical data that can be obtained from a microarray experiment. It includes direct and indirect targets for a set of different transcription factors. It is also expected that even for the main biological processes involved in this stress answer, the TFs binding sites will not correspond to strong signals in the sequences.

1.2 Results with MUSA algorithm

MUSA (Motif finding using an UnSupervised Approach) is a motif finder algorithm that does not require the user to specify parameters (such as box lengths and distances between boxes) in order to extract motifs [7]. The algorithm can therefore either be used autonomously to search for motifs, or to estimate the search parameters to be used in other motif discovery tools. The method relies on a biclustering algorithm that operates on a matrix of co-occurrences of small motifs. The performance of this method is independent of the composite structure of the motifs being sought.

Requiring as input a list of genes, MUSA returns the list of structured or simple motifs found, ordered by their p-value, and the proportion of sequences containing each motif (the quorum).

For all datasets, MUSA algorithm was executed with the default parameter values: $\lambda = 4$, $\epsilon = 1$ and the quorum equal to 30%.

The complete set of results obtained for the four tested datasets can be found through the following links:

- DeRisi dataset: <http://kdbio.inesc-id.pt/mat/isbra08/derisi/musa>
- Aft2p dataset: <http://kdbio.inesc-id.pt/mat/isbra08/aft2p/musa>
- Yap1p dataset: <http://kdbio.inesc-id.pt/mat/isbra08/yap1p/musa>
- 2,4D dataset: <http://kdbio.inesc-id.pt/mat/isbra08/24d/musa>

Each link correspond to a web page that contains a link to the MUSA's output and a table listing all the motif communities identified. Each entry contains a logo depicting the PWM (Position Weight Matrix) of the community, the p-value of the most statistically significant

motif within the community and a link to a file containing the list of motifs in the community and the PWM itself.

Table 2 summarizes the results obtained. For each dataset it shows the total number of motifs found by the motif finder (T. mot), the number of motifs considered for processing (N. mot), the total number of edges in the relation graph (N. edg), the number of communities identified (N. com) and the modularity value achieved (Modul).

Table 2. Results statistics.

<i>Datasets</i>	<i>T.mot</i>	<i>N.mot</i>	<i>N.edg</i>	<i>N.com</i>	<i>Modul</i>
DeRisi	1647	204	299	89	0.80
Aft2p	2176	2026	3397	202	0.84
Yap1p	2150	2065	3541	168	0.84
2,4D	2088	2045	3143	271	0.86

Since MUSA and RISO algorithms are integrated in the YEASTRACT-DISCOVERER system, it is possible to use other facilities available to help the results analysis. Each of the PWMs, in the previous described links with output results, can be compared to the transcription factors binding sites (TFBS) contained in the YEASTRACT database, by following these steps:

- Select a family’s PWM in the column labeled ”Select”
- Select the parameters to be used in the alignment of the PWMs
- Press the button labeled ”Match!”, to compare the selected PWM with the database’s TFBS

The default metric is *Sum of the Squared Distances*, and the input PWM can also be trimmed. Trimming removes the columns at the edges of the PWM that have an information content below the selected threshold.

The comparison of the input PWM with the TFBS of the YEASTRACT database is done using a procedure described in [8]. First of all, the TFBS of the YEASTRACT database are converted to PWMs, using the IUPAC rules and assuming equiprobability between the nucleotides. The input PWM is then locally aligned (using the Smith-Waterman local alignment algorithm) with each of the TFBS PWMs, with the selected column distance metric to perform the alignment. Four distance metrics were implemented but the default metric was the one used in this context.

A detailed analysis of the results obtained for the first two datasets, in table 1, was performed. This analysis is presented here and in the paper. For the results obtained with the two other sets, further analysis is needed.

DeRisi dataset For this set, using the default input parameters and considering both strands, the MUSA algorithm identified 1647 over-represented motifs. From these, only 204 motifs were classified as statistically significant and considered for further processing.

Table 3 presents the first 20 motifs reported by MUSA. In this list, motifs 1 and 7 correspond to instances of the TF binding site of interest. Although these motifs are well positioned in the motifs list, it is possible to verify that they are only present in at most half of the input

Table 3. The top 20 motifs extracted for the DeRisi dataset

<i>ID</i>	<i>Motif</i>	<i>Quorum</i>	<i>P – value</i>
1	TCCGTGGA	12 of 25	2.79106e-17
2	TCCACGGA	12 of 25	2.79106e-17
3	AAGA (17,19) TTTC	18 of 25	3.57327e-16
4	GAAA (17,19) TCTT	18 of 25	3.57327e-16
5	CCGT (1,3) GAAA	13 of 25	7.22785e-15
6	TTTC (1,3) ACGG	13 of 25	7.22785e-15
7	CCACGGA	14 of 25	8.57847e-15
8	TCCGTGG	14 of 25	8.57847e-15
9	AAAA (4,6) AAAT	25 of 25	1.31737e-14
10	ATTT (4,6) TTTT	25 of 25	1.31737e-14
11	CCACGGAA	11 of 25	1.45924e-14
12	TTCCGTGG	11 of 25	1.45924e-14
13	AACA (43,45) CCTC	11 of 25	2.40999e-12
14	GAGG (43,45) TGTT	11 of 25	2.40999e-12
15	CAAAAG (3,5) AAAT	9 of 25	4.18278e-12
16	GACG (35,37) CATA	9 of 25	2.10875e-11
17	TATG (35,37) CGTC	9 of 25	2.10875e-11
18	CAACA (5,7) ATAA	9 of 25	6.84029e-11
19	AAAGAAAA	18 of 25	1.39523e-10
20	TTTTCTTT	18 of 25	1.39523e-10

sequences. This quorum uncovers the real importance of this binding site, something that is not expected in this particular dataset where all the genes were up-regulated by the Pdr1p transcription factor.

Figure 1 shows the 14 motifs which contribute to the most significant community obtained.

By inspecting this figure it is possible to see that the quorum of this community is 97%. This value is very important since it reflects the real importance of this binding site in this set. The community quorum is also an important feature in the evaluation of the community importance.

If a search for documented TFs binding sites is performed in YEASTRACT database using the PWM of the first community, the best match will be precisely one of the documented Pdr1p-binding sites. If this search is also performed for the second most important community, shown in Figure 1, again one of the documented Pdr1p-binding sites will be found. This second motif is similar to the first one but not included in the same community because of the difference in the central nucleotides: an GT switched for a AC. In fact, since MUSA searched for motifs in both strands, this second community roughly contains the reverse complement of the motifs contained in the first community.

Aft2p dataset For this set and considering the default parameters MUSA extracted 2176 motifs. From these, only 2026 were classified as statistically significant and considered for further processing. For this set, 202 communities were identified. Table 4 presents the top 20 motifs extracted by MUSA.

Figure 2 shows the PWM logo of the fourth most significant community found, that corresponds precisely to the documented Aft2p-binding site [2, 9]. The list of motifs in the alignment is not presented in this document since this community includes 124 motifs.

Community 1 (14 motifs) (quorum stat: min 0.71 avg 0.97 max 1.00)

---TCCGTGGA-	0	14	8	2.79106e-17
---TCCGTGG--	1	17	9	8.57847e-15
--TTCCGTGG--	2	12	9	1.45924e-14
----CCGTGG--	3	26	10	2.03787e-09
---TCCGTG---	4	18	9	2.16867e-07
--TTCCGT----	5	28	13	4.31347e-07
-----CGTGG--	6	45	16	1.13802e-05
-----CGTGGA-	7	23	10	5.08829e-05
----CCGTG---	8	33	13	0.000103709
-GTCCG-----	9	15	8	0.000258935
---TCCGTC---	10	16	5	0.00026753
CCCTC-----	11	36	14	0.000360318
-----GTGGA-	12	48	17	0.000506259
-----GTGGAA	13	21	8	0.000797191

A: 0.26 0.11 0.12 0.05 0.01 0.00 0.00 0.00 0.00 0.00 0.90 0.63

C: 0.12 0.28 0.10 0.04 0.97 1.00 0.00 0.00 0.01 0.00 0.03 0.03

G: 0.18 0.36 0.07 0.03 0.01 0.00 1.00 0.00 0.99 0.99 0.06 0.17

T: 0.44 0.25 0.71 0.88 0.01 0.00 0.00 1.00 0.00 0.01 0.01 0.17



Fig. 1. Depiction of the most significant community found for the DeRisi dataset: motifs alignment, PWM description and PWM logo



Fig. 2. Depiction of the trimmed PWM logo of community that matches the Aft2p-binding site

Table 4. The top 20 motifs reported by MUSA for the Aft2p dataset

<i>ID Motif</i>	<i>Quorum</i>	<i>P – value</i>
1 TTTT (13,15)CACC	64 of 193	6.54141e-34
2 ACATAT	139 of 193	1.63126e-31
3 ATATGT	139 of 193	1.63126e-31
4 CACCC	165 of 193	6.43219e-31
5 GGGTG	165 of 193	6.43219e-31
6 GAAGAA	149 of 193	6.3155e-30
7 TTCTTC	149 of 193	6.3155e-30
8 GTATAT	124 of 193	1.36158e-29
9 CATATA	130 of 193	3.35912e-29
10 AAGAAG	145 of 193	1.04215e-28
11 TATTCT	145 of 193	1.31282e-28
12 CAAGAA	148 of 193	5.27289e-28
13 TTCTTG	148 of 193	5.27289e-28
14 GCACC	155 of 193	4.69138e-27
15 GGTGC	155 of 193	4.69138e-27
16 CTTTTC	149 of 193	5.89157e-27
17 GAAAAG	149 of 193	5.89157e-27
18 AGGAAA	143 of 193	1.54727e-26
19 TTTCCCT	143 of 193	1.54727e-26
20 AAGGAA	142 of 193	3.50713e-26

The binding site previously described for this TF allows some variability in the peripheral nucleotides. It is interesting to note that the identified community also exhibits a central conserved core region, the motif CACCC, flanked by less conserved peripheral nucleotides. This core motif, that correspond to motif 4 in Table 4, is statistically significant and is present in 85% of the input sequences. The correspondent community still presents a better quorum, 91%.

The first three most significant communities found for this dataset were compared with the documented TF binding sites described in the YEASTRACT database. The third most significant community was also associated with the documented Aft2p-binding site. This community roughly contains the reverse complement of the motifs present in the fourth community. The second most significant community matched the *TATA-box*, a well characterized core promoter element also expected to be a strong signal in this dataset.

For the first community found, the alignments obtained were very poor, suggesting that there is no documented TF binding site with such characteristics. To further investigate the existence of a similar TF binding site, the original PWM was trimmed. Trimming removes the edges of the PWM with low information content, considering only the most informative columns for the alignments. In this case, the best match was with the binding site of the Rap1p transcription factor. This TF is described as a DNA-binding protein involved in either activation or repression of transcription, depending on binding site context. However, the trimmed PWM aligned only with a short part of the Rap1p TF binding site, which motivated us to search further for a community originating a better match with a documented TF binding site. Although presumably not related with the documented Aft2p binding site or other documented TF binding site, this motif can have an important biological meaning.

1.3 Results with RISO algorithm

RISO [10] is an algorithm that searches for simple and structured motifs in a given set of DNA sequences. The motifs that are found by RISO comply with certain characteristics which are specified by the user, such as the number and size of the boxes that form the structured motif, the distances between them and the minimum quorum expected. It is also possible to specify a number of substitutions for each box.

RISO searches for motifs in the input sequences by using all the input sequences to construct what is known as a factor tree, a kind of suffix tree pruned at a given depth (defined by the length of the boxes of the complex motif). RISO uses a new data structure, the box-link, in order to efficiently extract structured motifs from the set of promoters of the input genes.

For all datasets, RISO algorithm was instructed to identify simple motifs, sized between 4 and 7, with 0 substitutions and a quorum equal to 60%.

The complete set of results obtained for the four tested datasets can be found through the following links:

- DeRisi dataset: <http://kdbio.inesc-id.pt/mat/isbra08/derisi/riso>
- Aft2p dataset: <http://kdbio.inesc-id.pt/mat/isbra08/aft2p/riso>
- Yap1p dataset: <http://kdbio.inesc-id.pt/mat/isbra08/yap1p/riso>
- 2,4D dataset: <http://kdbio.inesc-id.pt/mat/isbra08/24d/riso>

Table 5 summarizes the results obtained. For each dataset it shows the total number of motifs found by the motif finder (T. mot), the number of motifs considered for processing (N. mot), the total number of edges in the relation graph (N. edg), the number of communities identified (N. com) and the modularity value achieved (Modul).

Table 5. Results statistics.

<i>Datasets</i>	<i>T.mot</i>	<i>N.mot</i>	<i>N.edg</i>	<i>N.com</i>	<i>Modul</i>
DeRisi	1353	264	249	88	0.89
Aft2p	1198	1094	527	682	0.97
Yap1p	1233	1133	698	612	0.96
2,4D	1204	1158	675	706	0.95

References

1. DeRisi, J., van den Hazel, B., Marc, P., Balzi, E., Brown, P., Jack, C., Goffeau, A.: Genome microarray analysis of transcriptional activation in multidrug resistance yeast mutants. *FEBS Letters* **470** (2000) 156–160
2. Courel, M., Lallet, S., Camadro, J., Blaiseau, P.: Direct activation of genes involved in intracellular iron use by the yeast iron-responsive transcription factor Aft2 without its paralog Aft1. *Molecular Cell Biology* **25**(15) (2005) 6760–6771
3. Cohen, B., Pilpel, Y., Mitra, R., Church, G.: Discrimination between paralogs using microarray analysis: application to the Yap1p and Yap2p transcriptional networks. *Molecular Biology of the Cell* **13**(7) (2002) 1608–1614
4. Teixeira, M., Fernandes, A., Mira, N., Becker, J., Sá-Correia, I.: Early transcriptional response of *Saccharomyces cerevisiae* to stress imposed by the herbicide 2, 4-dichlorophenoxyacetic acid. *FEMS Yeast Research* **6**(2) (2006) 230–248

5. Blaiseau, P., Lesuisse, E., Camadro, J.: Aft2p, a novel iron-regulated transcription activator that modulates, with Aft1p, intracellular iron use and resistance to oxidative stress in yeast. *Journal of Biological Chemistry* **276**(36) (2001) 34221–34226
6. Harbison, C., Gordon, D., Lee, T., Rinaldi, N., Macisaac, K., Danford, T., Hannett1, N., Tagne, J., Reynolds, D., Yoo, J., Jennings, E., Zeitlinger, J., Pokholok, D., Kellis, M., Rolfe, P., Takusagawa, K., Lander, E., Gifford, D., Fraenkel, E., Young, R.: Transcriptional regulatory code of a eukaryotic genome. *Nature* **431**(7004) (2004) 99–104
7. Mendes, N., Casimiro, A., Santos, P., S-Correia, I., Oliveira, A., Freitas, A.: MUSA: A parameter free algorithm for the identification of biologically significant motifs. *Bioinformatics* **22** (2006) 2996–3002
8. Mahony, S., Auron, P., Benos, P.: Dna familial binding profiles made easy: Comparison of various motif alignment and clustering strategies. *PLoS Computational Biology* **3**(3) e61. doi:10.1371/journal.pcbi.0030061
9. Rutherford, J., Jaron, S., Winge, D.: Aft1p and Aft2p mediate iron-responsive gene expression in yeast through related promoter elements. *Journal of Biological Chemistry* **278**(30) (2003) 27636–27643
10. Carvalho, A., Freitas, A., Oliveira, A., Sagot, M.F.: An efficient algorithm for the identification of structured motifs in DNA promoter sequences. *IEEE Transactions on Computational Biology and Bioinformatics* **3**(2) (2006) 126–140