



OSCAR: One-class SVM for Accurate Recognition of Cis-elements

Bo Jiang¹, Michael Q. Zhang^{2,1} and Xuegong Zhang^{1,*}

¹MOE Key Laboratory of Bioinformatics and Bioinformatics Div, TNLIST / Department of Automation, Tsinghua University, Beijing 100084, China

²Cold Spring Harbor Laboratory, Cold Spring Harbor, NY 11724, USA

ABSTRACT

Motivation: Traditional methods to identify potential binding sites of known transcription factors still suffer from large number of false predictions. They mostly use sequence information in a position-specific manner and neglect other types of information hidden in the proximal promoter regions. Recent biological and computational researches, however, suggest that there exist not only locational preferences of binding, but also correlations between transcription factors.

Results: In this paper, we propose a novel approach, OSCAR, which utilizes one-class SVM algorithms, and incorporates multiple factors to aid the recognition of transcription factor binding sites. Using both synthetic and real data, we find that our method outperforms existing algorithms, especially in the high sensitivity region. The performance of our method can be further improved by taking into account locational preference of binding events. By testing on experimentally-verified binding sites of GATA and HNF transcription factor families, we show that our algorithm can infer the true co-occurring motif pairs accurately, and by considering the co-occurrences of correlated motifs, we not only filter out false predictions, but also increase the sensitivity.

Availability: An online server based on OSCAR is available at <http://bioinfo.au.tsinghua.edu.cn/oscar>.

Contact: zhangxg@tsinghua.edu.cn

1 INTRODUCTION

Identification of transcription factor binding sites (TFBSs) and the corresponding motifs plays a pivotal role in the understanding of the transcriptional regulation mechanism. Recently, the advent of DNA microarray techniques and ChIP-chip experiments has significantly improved our ability to identify those important *cis*-regulatory elements. However, since these experimental procedures are expensive and time-consuming, computational methods are still needed for predicting TFBSs.

The task of computational identification of TFBSs can be divided into two main categories. One is motif discovery, *i.e.* to discover a motif as well as its putative sites in a collection

of genomic sequences that are expected to be bound by the same factor. Many state-of-the-art algorithms, such as MEME (Bailey and Elkan, 1994), Bioproscpector (Liu *et al.*, 2001), YMF (Sinha and Tompa, 2002) *etc.* (for a recent review, see Tompa *et al.*, 2005), were designed to this end. The other task is motif search, *i.e.* to scan potential binding sites of a known transcription factor, often, on a genomic scale. Several tools like MatInspector (Quandt *et al.*, 1995) and MATCHTM (Kel *et al.*, 2003) were developed to search for putative TFBSs in DNA sequences by using position weight matrices (PWMs; Stormo *et al.*, 1982) in TRANSFAC[®] (Wingender *et al.*, 2000) or JASPAR (Albin *et al.*, 2004) databases. In this paper, we focus on the latter task, which is to recognize the putative binding sites of known transcription factors.

Previous works to predict new putative TFBSs based on known binding sites still suffer from large number of false predictions when applied in a genomic scale. Most commonly used tools ignore the extra information in and beyond the TFBSs, such as the conservation of the binding site locations and the co-occurrences of other motifs in the promoter sequences. There is ample evidence that many *cis*-regulatory elements show preferred locations (Frith *et al.*, 2003; FitzGerald *et al.*, 2004; Xie *et al.*, 2005), and precise organizations (Boyer *et al.*, 2005; Odom *et al.*, 2006) within promoter sequences. A search of the binding sites on the entire genome without such information usually returns a large number of sites, many of which are not functional *in vivo*, although they would probably bind to the transcription factor if they were in proper genomic contexts.

Recently, many classification algorithms, such as support vector machines (SVMs), have been applied to discriminate false predictions from the true ones in TFBS recognition (Holloway, 2005; Sun *et al.*, 2006) and other applications. Jaakkola *et al.* (2000) proposed a discriminative framework with a variant of support vector machines to detect protein homology, which can potentially be applied to the detection of TFBSs. Special string kernels in SVMs were also designed to process the regulatory regions of genes, in order to recognize a given class of promoter region, and simultaneously identify a collection of relevant, discriminative se-

* To whom correspondence should be addressed.

quence motifs (Leslie *et al.*, 2004; Sharan *et al.*, 2005; Sonnenburg *et al.*, 2005a, 2005b; Ratsch *et al.*, 2005; Vert *et al.*, 2005). However, as with other classification techniques, a set of known positive and negative samples must be supplied to the SVMs. With limited instances of experimentally-verified binding sites, it is hard to determine a “negative” set of sequences, to which a transcription factor will certainly not bind (Hong *et al.*, 2005). In fact, there exists no site (separated from its context) that will never be bound by a transcription factor, but only sites that are unlikely to be bound. In this sense, recognition of TFBSs may not be well characterized as a standard two-class classification problem.

In this paper, we exploit the one-class SVM, rather than two-class approaches, to estimate the support of probability distribution of known TFBSs, *i.e.* the region where most known (positive) samples live in the feature space. Our method, named OSCAR (One-class Support vector machine for *Cis*-element Accurate Recognition), simultaneously considers the nucleotide composition of all the positions within a binding site, and further incorporates locational preferences and co-occurrences of DNA motifs. We first demonstrate that our novel approach outperforms existing algorithms on synthetic data. Applying our method to promoter regions in the Eukaryotic Promoter Database (EPD; Praz *et al.*, 2002) with annotated TFBSs from TRANSFAC[®], we find that the performance of our algorithm can be further improved by considering the locational preferences of binding events. We also show that the mutual interactions of transcription factors can be identified and exploited by using our method. Furthermore, the application of our method to recognize the binding sites of GATA and HNF transcription factor families indicates that the integration of co-occurrence information not only decrease false positives, but also increase the sensitivity of the prediction.

2 MATERIALS AND METHODS

2.1 Databases for identification of TFBSs

In this study we used two databases: (1) TRANSFAC[®] (release 9.4; Wingender *et al.*, 2000) is a database that provides data on transcription factors and their binding sites in promoters of eukaryotic genes as well as a library of PWMs. It also offers an additional profile, called “non-redundant profile for vertebrates”, by categorizing PWMs into groups on the basis of the linked transcription factors and selecting just one “best” matrix from each group. (2) The Eukaryotic Promoter Database (EPD; Praz *et al.*, 2002) is an annotated non-redundant collection of eukaryotic promoters, for which the transcription start site (TSS) has been determined experimentally. The annotation part of an entry includes description of the initiation site mapping data, cross-references to other databases including TRANSFAC[®].

2.2 One-class SVM and basic OSCAR algorithm

One-class Support Vector Machine (Scholkopf *et al.*, 2001) was proposed to estimate the support of a high-dimensional distribution, *i.e.* the region where most of the data live in the high-dimensional feature space. It returns a prediction function f that takes the value +1 in a “small” region capturing most of the data points, and -1 elsewhere.

Given known binding sites of a particular transcription factor, we first encode each site s with length L into a binary string \mathbf{x} with length $4L$. The nucleotide at each position of the binding site is mapped into a 4-dimension vector in $\{0, 1\}$ -space with following rule: A \mapsto 0001, C \mapsto 0010, G \mapsto 0100, and T \mapsto 1000. To determine whether a DNA sequence s_0 with length L is a putative binding site of the transcription factor, s_0 is transformed into a binary string \mathbf{x}_0 in the same way, and the prediction function is given as

$$f = \text{sign}((w \cdot \mathbf{x}_0) - \rho), \quad (1)$$

where the weight w and parameter ρ are determined by the one-class SVM training algorithm given below. Note that w can be denoted in a matrix form: $\mathbf{M}_{i,j} = w_{4xi+j}$, where $i = 1, \dots, L$, and $j = 1, \dots, 4$ correspond to four nucleotides. Regarding the parameter ρ in prediction function (1) as a threshold, we come to a scoring scheme similar to PWM-based methods, yet in our method nucleotide preferences at different positions within a binding site are considered simultaneously in SVM model.

To be more specific, consider the training data $\mathbf{x}_1, \dots, \mathbf{x}_l \in \Phi$, where \mathbf{x}_i is a binary string encoded by a known binding site, Φ is a $\{0, 1\}$ -vector space with $4L$ dimension, and l is the number of known binding sites. One-class SVM solves the following quadratic programming problem:

$$\min_{w \in \Phi, \xi_i \in \mathbf{R}^+, \rho \in \mathbf{R}} \frac{1}{2} \|w\|^2 + \frac{1}{\nu l} \sum_i \xi_i - \rho, \quad (2)$$

$$\text{subject to } (w \cdot \mathbf{x}_i) \geq \rho - \xi_i, \quad \xi_i \geq 0. \quad (3)$$

Here, $\nu \in (0, 1)$ is a user-defined parameter.

For the \mathbf{x}_j 's that the equalities in (2) hold at the solutions of w and ρ , we have $\rho = (w \cdot \mathbf{x}_j)$ and \mathbf{x}_j is named a “support vector”. Since non-zero slack variables ξ_i are penalized in the objective function, we can expect that the prediction function (1) will be positive for most samples contained in the training set, while the regularization term $\|w\|$ will still be small. The actual trade-off between these two goals is controlled by ν . In fact, Scholkopf *et al.* (2001) have proved that ν is an upper bound on the fraction of outliers and a lower bound on the fraction of support vectors.

To solve w and ρ from problem (2) and (3), we use the LIBSVM package (Chang and Lin, 2001) for the implementation of a one-class SVM. In our applications, one-class SVM is used to construct prediction functions with training sets containing only experimentally-verified binding sites of known transcription factors. The algorithmic framework of a one-class SVM not only enables us to exploit the nucleotide composition of known binding sites, but

also allows us to consider their locational preferences and effects of TFBS co-occurrences.

2.3 Incorporating locational preference

Binding sites of many transcription factors show preferred locations within promoter regions, which are influenced by nucleosome occupancy (Segal et al., 2006), chromatin structure and other factors. The locational preferences of binding events can be conducive to the recognition of functional binding sites in genome. A one-class SVM can serve as a good algorithmic frame, which allows us to exploit this type of information when the relative locations to the TSS of known binding sites are available. To encode such information, we first divide the promoter regions that contain the binding sites in training data into K consecutive intervals relative to TSS, where K and the length of each interval are user-defined parameters. We recommend choosing the interval length from 150 bp to 300 bp, which is approximately the total length of a nucleosome unit and the linker region. The location of a known binding site, if available, is described by a binary string with length K , in which a non-zero value indicates the occurrence of binding sites in the corresponding interval (see examples in Figure 1). This additional string is concatenated to the original string that depicts the nucleotide composition of a binding site, as the training vector of the one-class SVM. Note that if the location of binding site is not known, the additional string will be all zero, and the corresponding method will become the basic OSCAR algorithm.

To search for putative binding sites of the transcription factor in a promoter region, the region is divided into intervals according to its TSS in the same way as described above. The sites in different intervals are encoded accordingly, and determined by the prediction function in the one-class SVM. To distinguish the algorithm that incorporates locational preferences from the basic OSCAR algorithm, we call this one OSCAR-L.




Figure 1. Illustration for encoding locational preferences in OSCAR-L. The arrow indicates transcription start site (TSS); colored intervals delineate the division of the promoter region; boxes denote the locations of binding sites, and the locational coding of each site is given below.

2.4 Exploiting co-occurrence information

Gene regulation is often achieved by a precise organization of multiple transcription factors. Some transcription factors are by nature moderately or poorly specific in their DNA binding and achieve higher specificity only in the context of other binding partners. In this subsection, we further extend

OSCAR to recognize putative binding sites with a proximal binding of other co-factors. We call this method OSCAR-C.

First, a set of M candidate motifs which may co-occur with a specific transcription factor (named “primary TF” here) is obtained from the database or predefined by the user. To search for the binding sites of the primary TF in a promoter region, M additional bits are added to the previous binary string, each of which indicates whether the corresponding candidate motif appears in this region. In the training process, the occurrences of a candidate motif in a promoter region can be determined from the annotations in the database. While in the predictions when we do not have the knowledge about the binding of candidate motifs, the prediction function obtained by the OSCAR-L algorithm described in the above subsection can be used to identify the occurrences of candidate motifs.

The candidate motif set may contain irrelevant motifs, and lead to over-fitting of the algorithm. Thus, the accuracy of prediction might be improved through the exclusion of candidate motifs with little relevance to the binding of the primary TF. Besides, knowing which candidate motifs are relevant can give insight to the interactions between transcription factors. Therefore, a criterion similar to SVM-RFE (Guyon et al., 2002) is proposed to perform the selection of co-occurring candidate motifs according to their contributions in predicting the binding of the primary TF. Assume that w_m is the weight corresponding to the occurrence of the m th candidate motif in prediction function, we use w_m^2 as the ranking score to select relevant candidate motifs. Let J denote the objective function in (2), which becomes $(1/2)\|w\|^2$ when data points are separable from the origin. Our ranking criterion is then explained by the OBD algorithm (LeCun et al., 1990), which approximates the change in objective function due to the removal of the i th feature by expanding objective function J in the Taylor series to the second order:

$$\Delta J(m) = \frac{\partial J}{\partial w_m} \Delta w_m + \frac{\partial^2 J}{\partial w_m^2} (\Delta w_m)^2, \quad (4)$$

where $m=1, \dots, M$. At the minimum of J , the first order term can be neglected. With $J = (1/2)\|w\|^2$, Equation (4) becomes $\Delta J(m) = (\Delta w_m)^2$, where $\Delta w_m = w_m$ corresponds to removing the m th feature.

Given a cut-off on the ranking list, we can obtain candidate motifs relevant to the binding of the primary TF. The indicators for the binding of selected motifs are used in the one-class SVM algorithm together with other features, and an integrated prediction function is finally obtained.

3 EXPERIMENT RESULTS

3.1 Results on synthetic data

We begin by evaluating our methods on synthetic data. To this end, we retrieve the binding sites of transcription factors

from TRANSFAC[®] (release 9.4) non-redundant profiles, and 63 vertebrate transcription factors with more than 10 annotated binding sites are selected. To avoid over-fitting effects in evaluating algorithms, Leave-One-Out Cross Validation (LOOCV) is performed together with the synthetic procedure: for a set with N known binding sites of a transcription factor, we first randomly implant a site into an artificial sequence that is generated according to uniform nucleotide compositions, and then we apply the basic OSCAR algorithm, which uses the remaining $N-1$ sites as training data, to predict putative binding sites in the synthetic sequence. This step is repeated N times with each site implanted into a random sequence. To make comparisons, we test the performance of three other methods: simple PWM based on log-odds ratio (LOPWM), MATCH[™] (Kel *et al.*, 2003) and P-Match (Chekmenev *et al.*, 2005), which are closely interconnected with TRANSFAC[®] database. MATCH[™] is a representative PWM-based tool for searching putative transcription factor binding sites in DNA sequences, while P-Match is a newly developed tool that combines pattern matching and weight matrix approaches.

In synthetic data, the locations of binding events are known with certainty, which enable us to evaluate the predicting ability of different tools according to known binding sites in simulated sequences. A predicted position is a true-positive (TP) prediction if it coincides with a known binding site, and it is a false-positive (FP) prediction otherwise. To assess the performance of algorithms, we calculate the true positive rate (TPR, also known as “sensitivity”), which is the ratio of the number of true positives to the total number of known binding sites, and false positive rate (FPR), which is the ratio of the number of false positives to the total number of non-site positions. A curve of FPR versus TPR can be plotted while a threshold parameter is varied. This curve, which is called ROC (Receiver Operating Characteristics) curve, is a comprehensive and objective way to compare the performance of methods as a tradeoff between specificity and sensitivity. ROC curves are obtained by setting ν of one-class SVM in a range from 0.1 to 0.9. To evaluate MATCH and P-Match, we vary the two cut-offs of each algorithm in a complete range of values, and report the maximum TPR achieved at a given FPR.

In Figure 2, the area under the ROC curve indicates that OSCAR is able to pick out more true-positive sites given fixed number of false predictions in the region of high FPR (with a relatively small ν). In the low FPR region, the performance of OSCAR is comparable with LOPWM and MATCH algorithms. This can be explained by the fact that the number of support vectors in the training samples increases as ν becomes larger, and as a result, the weight vectors of the basic OSCAR algorithm become similar to the entries of traditional PWMs.




Figure 2. ROC curves show the performances of different methods on synthetic data. Note that a logarithmic scale transformation on the X-axis is used to show the low FPR performances of different methods.

3.2 Results on EPD promoter regions

To verify our methods on real biological data, we determine the locations of TFBSs in TRANSFAC[®] database according to the annotations in EPD (release 85). Among all the transcription factors in TRANSFAC[®] non-redundant profiles, binding sites of 45 vertebrate transcription factors can be identified, corresponding to totally 124 promoters in EPD85 (data available from our website). Promoter regions from upstream 1000 bp to downstream 300 bp are retrieved. To test the sensitivity of algorithms, we use the LOOCV procedure as follows: for a set with N known binding sites of a transcription factor, we obtain a prediction function on $N-1$ sites and test the remaining site based on the score. This step is repeated N times with each site serving as testing sample. TPR is equal to the number of predicted binding sites divided by total number of known sites. For the purpose of test, we construct a background set, which presumably contains few true TFBSs, by randomly permuting the sequences of promoter regions in the training set with conserved dinucleotides (Coward, 1999). Predictions were made on the background set by moving a scanning window of the motif length at the step of 1 base-pair. The FPR is calculated as the ratio of the number of predictions to the total number of base-pairs on background sequences. The evaluations of LOPWM, MATCH, P-Match and the basic OSCAR algorithms are performed by using the same procedure. Besides, we also apply the OSCAR-L algorithm to incorporate the locational preference, by dividing each promoter region into 5 consecutive intervals with equal length of 260 bp. ROC curves in Figure 3 show the overall performance of different methods on EPD promoters for all 45 transcription factors.




Figure 3. ROC curves show the performances of different methods on EPD promoter regions, and confirm the advantage of incorporating locational preference of binding events. Note that a logarithmic scale transformation on the X-axis is used to show the low FPR performances of different methods. True positive rate is obtained by LOOCV, and false positive rate is estimated on randomly shuffled EPD promoter regions.

As we can see, OSCAR and P-Match are generally more accurate than LOPWM and MATCH in the area of high FPR, while OSCAR, LOPWM and MATCH are comparable in the low FPR region, which are consistent with observations in literature (Chekmenev *et al.*, 2005) and the results on our synthetic data. As expected, the overall performance can be improved by incorporating locational preferences into the basic OSCAR algorithm (OSCAR-L). At a given sensitivity level, OSCAR-L will result in a smaller false positive rate than the basic OSCAR algorithm when averaged over all 45 transcription factors (see Supplementary Table 1).

3.3 Results on GATA and HNF family

In order to check the ability of our method to reveal known TFBSs through identifying and utilizing co-occurring motifs (OSCAR-C), we test OSCAR-C and other methods on two additional data sets involving GATA factor and HNF factor family, which have relatively large numbers of binding sites that can be identified on EPD promoters. To apply the OSCAR-C algorithm, 141 motifs in TRANSFAC[®] database for vertebrates constitute the candidate motif set. Their co-occurrences with the binding sites of the primary TF are input as additional features to obtain the prediction function. To further evaluate the statistical significance of co-occurring motifs selected by the OSCAR-C algorithm, we randomly permute the bases in the promoter sequences, while preserving the confirmed sites of the primary factor. The permutation is repeated for 100 times, and the OSCAR-

C algorithm is applied to each permuted set. The distribution of the largest and second largest absolute values of weights are used to assess the significance of the top two motifs on the ranking list, respectively.

GATA factor binding sites Transcription factors in the GATA family are so-called because they bind to the consensus DNA sequence “(A/T)GATA(A/G)”. They are shown to play critical roles in development, including in cell-fate specification, regulation of differentiation and control of cell proliferation and movement (Patient *et al.* 2002). We extract 13 EPD promoter regions which contain 25 TRANSFAC annotated binding sites of a GATA transcription factor (V\$GATA_Q6). The two most relevant motifs that co-occur with the binding of GATA factor is identified by the OSCAR-C algorithm and given in the first row of Table 1, together with the p-values evaluated by the procedure described above. The co-occurrences of GATA with NF-Y and Sp1 factors are consistent with biological evidence in the literature (*e.g.* Huang *et al.*, 2004; Furusawa, *et al.*, 2003).

We incorporate co-occurrences of two motifs to predict binding sites of GATA factor. Again, we use LOOCV procedure to test the TPRs of MATCH, P-Match, the basic OSCAR, OSCAR-L, and OSCAR-C algorithms. FPRs are also evaluated by randomly shuffling the promoter regions in training set. Table 2 shows the true positive rates of different methods given fixed number of false predictions. A complete plot of ROC curves is given in Supplementary Figure 1. We can see that there is a clear advantage from utilizing co-occurrence information, indicating that the integration of co-occurring motifs not only decreases false predictions, but also increases sensitivity by enhancing the signal strength.

HNF family binding sites Hepatocyte nuclear factors 1, 3 and 4 (HNF-1, -3 and -4) are liver-enriched transcription factors that function in the regulation of several liver-specific genes (Ktistaki *et al.*, 1997). Totally 22 promoter regions are extracted from EPD85 with annotated binding sites for transcription factor HNF-1 (on 7 promoters), HNF-3 (on 7 promoters) and HNF-4 (on 13 promoters) in TRANSFAC[®]. Motifs co-occurring with different factors in HNF family are identified and evaluated separately, and their TRANSFAC IDs and corresponding p-values are shown in Table 1. These results are supported by evidence from the biological experimental literature (*e.g.* Hiesberger *et al.*, 2004; Antes *et al.*, 2001; Elholm *et al.*, 1996; Kahn, 1997; Group *et al.*, 1994).

The co-occurrence information is exploited by OSCAR-C. Predictions are made separately for different factors in HNF family, and the results are summed up in Table 3 and Supplementary Figure 2. From the results in Table 3 and ROC curves in Supplementary Figure 2, we can find that they are consistent with what we have observed in the other studies,

and that the performance of OSCAR-C is superior over other methods.

Table 1. Motifs co-occurring with the binding of GATA factor and HNF family identified by the OSCAR-C algorithm

TF family	Primary Motif	Co-occurring Motifs	P-value
GATA	V\$GATA_Q6	V\$NFY_Q6_01	< 0.01
		V\$SP1_Q2_01	0.02
HNF	V\$HNF1_Q6_01	V\$PIT1_Q6	0.01
		V\$AP4_Q6_01	0.02
	V\$HNF3_Q6_01	V\$CEBP_Q3	0.01
		V\$NF1_Q6_01	0.01
	V\$HNF4_Q6_01	V\$SRF_Q5_02	< 0.01
		V\$USF_Q6_01	0.03

Table 2. True positive rates (%) of different method to identify GATA factor binding sites given fixed levels of false positive rate

FPR (%)	OSCAR-C	OSCAR-L	OSCAR	P-Match	MATCH
1.0	100.0	96.0	92.0	68.0	48.0
0.5	96.0	84.0	84.0	48.0	40.0
0.1	84.0	64.0	48.0	NA [§]	20.0
0.05	68.0	48.0	44.0	NA [§]	20.0

[§]The lowest false positive rate that P-Match could achieve on this data set is 0.23%, corresponding to a true positive rate of 36.0%.

Table 3. True positive rates (%) of different method to identify TFBSs of transcription factors in HNF family given fixed levels of false positive rate

FPR (%)	OSCAR-C	OSCAR-L	OSCAR	P-Match	MATCH
1.0	96.7	93.3	93.3	43.3	53.3
0.5	93.3	90.0	90.0	40.0	26.7
0.1	86.7	80.0	76.7	33.3	10.0
0.05	76.7	73.3	66.7	30.0	3.3

4 DISCUSSION

In this paper, we use the one-class SVM algorithm to identify *cis*-regulatory elements in promoter regions when only limited, “positive” samples are available. To improve the accuracy of prediction, we further consider the preferences of binding site locations, and the co-occurrences of other motifs in promoter regions. Applying our method to synthetic as well as real data, we have demonstrated the advantage of the proposed strategy, and confirmed the conclusion that incorporation of locational preference can improve the performance of prediction. We also illustrate how the method can be used to identify and utilize co-occurring motif pairs in predicting binding sites of two transcription factor families.

Note that, in the OSCAR algorithm, we construct a linear prediction function, which assumes an additive contribution from each position towards the score. Different from PWM-based methods, our method allows for encoding each bind-

ing site as a unitary vector. Some higher-order models, such as Bayesian tree structure (Barash *et al.*, 2003) or dinucleotide interactions (Zhou and Liu., 2004), were exploited to improve the prediction accuracy. However, even in cases where intra-site interactions exist, the additive model has been suggested to be a good approximation (Benos *et al.*, 2002). In our applications, we limit to construct linear prediction rules. The “kernel trick” in SVM, which allows for non-linearly mapping into a high-dimensional feature space, can be used as well. We have tried both polynomial and RBF kernels in our experiments, but little improvement has been observed. Since high-dimensional feature spaces may suffer from low generalization ability, we recommend using the linear kernel in this context.

An extensive investigation of the correlated motifs in prediction can help us to better understand gene regulations mechanisms. In this paper, we focus on how to recognize the binding sites of a particular transcription factor more accurately by exploring the related information more efficiently, but the method presented in this paper can also be used to infer the co-occurrences of motifs in the future. In addition, we use discrete interval indicators for the division of promoter regions in order to utilize the locational information. Given larger amount of available data, the algorithm may be further improved by using the exact positions of binding events directly. Moreover, other types of information, such as the occupancy of nucleosome (an algorithm that can use nucleosome occupancy information in TFBS prediction is in preparation), evolutionary conservation between related species, and the tissue specificity of the downstream genes, may be further exploited by incorporating them into our method.

ACKNOWLEDGEMENTS

We thank Dr. Andrew Smith for his help with the Shufflet program, and Dr. Xiaoyue Zhao for helpful discussions.

This work is partially supported by NSFC grants 60540420569 (XZ and MQZ), 30625012 (XZ), the National Basic Research Program 2004CB518605 (XZ), the Changjiang Professorship Award (MQZ) of China and NIH grant HG001696 (MQZ) of USA.

REFERENCES

- Antes, T.J., Levy-Wilson, B. (2001) HNF-3 beta, C/EBP beta, and HNF-4 act in synergy to enhance transcription of the human apolipoprotein B gene in intestinal cells. *DNA Cell Biol.*, 20, 67-74.
- Bailey, T.L. and Elkan, C. (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, AAAI Press, Menlo Park, California, 1994.
- Barash, Y., Elidan, G., Friedman, N. and Kaplan, T. (2003) Modeling dependence in protein-DBA binding sites. *RECOMB '03*. Berline, Germany, April 2003.
- Benos, P.V., Bulyk, M.L. and Stormo, G.D. (2002) Additive in protein-DNA interactions: how good an approximation is it? *Nucleic Acids Research*, 30, 4442-4451.

- Boyer, L.A., Lee, T.I., Cole, M.F., Johnstone, S.E., Levine, S.S., Zucker, J.P., Guenther, M.G., Kumar, R.M., Murray, H.L., Jenner, R.G., Melton, D.A., Gifford, D.K., Jaenisch, R. and Young, R.A. (2005) Core Transcriptional Regulatory Circuitry in Human Embryonic Stem Cells. *Cell*, 122: 947-956.
- Chang, C.C., and Lin, C. (2001) LIBSVM : a library for support vector machines. <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- Chekmenov, D.S., Haid, C. and A.E. Kel. (2005) P-Match: transcription factor binding site search by combining patterns and weight matrices. *Nucleic Acids Research*, 33, W432-W437.
- Coward, E. (1999) Shufflet: shuffling sequences while conserving the k-let counts. *Bioinformatics*, 15: 1058-1059.
- Elholm, M., Bjerking, G., Knudsen, J., Kristiansen, K. and Mandrup, S. (1996) Regulatory elements in the promoter region of the rat gene encoding the acyl-CoA-binding protein. *Gene*, 173, 233-8.
- FitzGerald, P.C., Shlyakhtenko, A., Mir, A.A. and Vinson, C. (2004) Clustering of DNA sequences in human promoters. *Genome Research*, 14, 1562-1574.
- Frith, M.C., Li, M.C. and Weng, Z. (2003) Cluster-Buster: Finding dense clusters of motifs in DNA sequences. *Nucleic Acids Research*, 31, 3666-3668.
- Furusawa, M., Taira, T., Iguchi-Arigo, S.M., Ariga, H. (2003) Molecular cloning of the mouse AMY-1 gene and identification of the synergistic activation of the AMY-1 promoter by GATA-1 and Sp1. *Genomics*, 81, 221-33.
- Group, E.R., Crawford, N., and Locker J. (1994) Characterization of the distal alpha-fetoprotein enhancer, a strong, long distance, liver-specific activator. *J. Biol. Chem.*, 269, 22178-2218.
- Guyon, I., Weston, J., Barnhill, S. and Vapnik, V. (2002) Gene selection for cancer classification using support vector machines. *Machine Learning*, 46, 389-422.
- Hiesberger, T., Bai, Y., Shao, X., McNally, B.T., Sinclair, A.M., Tian, X., Somlo, S. and Igarashi, P. (2004) Mutation of hepatocyte nuclear factor-1beta inhibits Pkhd1 gene expression and produces renal cysts in mice. *J. Clin. Invest.*, 113, 814-825.
- Holloway, D.T., Kon, M. and DeLisi, C. (2005) Integrating genomic data to predict transcription factor binding. *Genome Informatics*, 16: 83-94.
- Hong, P., Liu, X.S., Zhou, Q., Lu, X., Liu, J.S. and Wong, W.H. (2005) A boosting approach for motif modeling using ChIP-chip data. *Bioinformatics*, 21(11): 2636-2643.
- Huang, D.Y., Kuo, Y.Y., Lai, J.S., Suzuki, Y., Sugano, S., and Chang, Z.F. (2004) GATA-1 and NF-Y cooperate to mediate erythroid-specific transcription of Gfi-1B gene. *Nucleic Acids Research*, 32, 3935-3946.
- Jaakkola, T., Diekhans, M. and Haussler, D. (2000) A discriminative framework for detecting remote protein homologies. *Journal of Computational Biology*, 7, 95-114.
- Kahn, A. (1997) Transcriptional regulation by glucose in the liver. *Biochimie*, 79, 113-118.
- Kel, A.E., Gößling, E., Reuter, I., Cheremushkin, E., Kel-Margoulis, O.V. and Wingerder, E. (2003) MATCHTM: A tool for searching transcription factor binding sites in DNA sequences. *Nucleic Acids Research*, 31, 3576-3579.
- Ktistaki, E. and Talianidis, I. (1997) Modulation of hepatic gene expression by hepatocyte nuclear factor 1. *Science*, 277, 109-112.
- LeCun, Y., Denker, J. S. and Solla, S. A. (1990) Optimum brain damage. *Advances in neural information processing systems 2*, 598-605. San Mateo, CA: Morgan Kaufmann.
- Leslie, C.S., Eskin, E., Cohen, A., Weston, J. and Noble, W.S. (2004) Mismatch string kernels for discriminative protein classification. *Bioinformatics*, 20:467-76.
- Liu, X., Brutlag, D.L. and Liu J.S. (2001) BioProspector: discovering conserved DNA motifs in upstream regulatory regions of co-expressed genes. *Proceedings of the Pacific Symposium on Biocomputing*, 2001:127-138.
- Odom, D.T., Dowell, R.D., Jacobsen, E.S., Nekludova, L., Rolfe, P.A., Danford, T.W., Gifford, D.K., Fraenkel, E., Bell, G.I. and Young, R.A. (2006) Core transcriptional regulatory circuitry in human hepatocytes. *Molecular Systems Biology*, 2:2006.0017
- Patient, R.K., Mcghee, J.D. The GATA family (vertebrates and invertebrates). (2002) *Curr Opin Genet Dev*, 12, 416-422
- Praz, V., Perier, R., Bonnard, C. and Bucher, P. (2002) The Eukaryotic Promoter Database, EPD: new entry types and links to gene expression data. *Nucleic Acids Research*, 30: 322-324.
- Quandt, K., Frech, K., Karas, H., Wingender, E. and Werner, T. (1995) MatInd and MatInspector: new fast and versatile tools for detection of consensus matches in nucleotide sequence data. *Nucleic Acids Research*, 23, 4878-4884.
- Rätsch, G., Sonnenburg, S. and Schölkopf, B. (2005) RASE: recognition of alternatively spliced exons in *C.elegans*. *Bioinformatics*, 21: i369-i377.
- Segal, E., Fondufe-Mittendorf, Y., Chen, L., Thåström, A., Field, Y., Moore, I.K., Wang, J.Z. and Widorn J. (2006) A genomic code for nucleosome positioning. *Nature*, 442, 772-778.
- Sandelin, A., Alkema, W., Engström, P., Wasserman, W. and Lenhard, B. (2004) JASPAR: an open-access database for eukaryotic transcription factor binding profiles, *Nucleic Acids Research*, 32, 91-94.
- Schölkopf, B., Platt, J., Shawe-Taylor, J., Smola, A.J. and Williamson, R.C. (2001) Estimating the support of a high-dimensional distribution. *Neural Computation*, 13, 1443-1471
- Sharan, R. Myers, E.W. (2005) A motif-based framework for recognizing sequence families. *Bioinformatics*, 21: i387-i393.
- Sinha, S. and Tompa, M. (2002) Discovery of novel transcription factor binding sites by statistical overrepresentation. *Nucleic Acids Research*, 30, 5549-5560.
- Sonnenburg, S., Rätsch G. and Schölkopf B. (2005a) Large Scale Genomic Sequence SVM Classifiers. *Proceedings of the 22nd International Conference on Machine Learning*, 849-856.
- Sonnenburg, S., Rätsch, G. and Schäfer, C. (2005b) Learning interpretable SVMs for biological sequence classification. *RECOMB 2005, LNBI 3500*, 389-407.
- Stormo, G.D., Schneider, T.D., Gold, L. and Ehrenfeucht, A. (1982) Use of the 'Perceptron' algorithm to distinguish translational initiation sites in *E.coli*. *Nucleic Acids Research*, 10, 2997-3011.
- Sun, Y., Robinson, M., Adams, R., Boekhorst, R., Rust, A.G. and Davey, D. (2006) Using feature selection filtering methods for binding site prediction. *5th IEEE International Conference on Cognitive Informatics*, 2006.
- Tompa, M., et al. (2005) Assessing computational tools for the discovery of transcription factor binding sites. *Nature Biotechnology*, 23,137-144.
- Vert, J.P., Thurman, R. and Noble, W.S. (2005) Kernels for gene regulatory regions. *Advances in Neural Information Processing Systems 18*.
- Wingender, E., Chen, X., Hehl, R., Karas, H., Liebich, I., Matys, V., Meinhardt, T., Pruss, M., Reuter, I. and Schachere, F. (2000) TRANSFAC[®]: an integral system for gene expression regulation. *Nucleic Acids Research*, 28, 316-319.
- Xie, X., Lu, J., Kulbokas, E.J., Golub, T.R., Mootha, V., Lindblad-Toh, K., Lander, E.S. and Kellis, M. (2005) Systematic discovery of regulatory motifs in human promoters and 30 UTRs by comparison of several mammals. *Nature*, 434, 338-345
- Zhou, Q. and Liu, J.S. (2004) Modeling within-motif dependence for transcription factor binding site predictions. *Bioinformatics*, 20, 909-916.