

Base-pairing profile local alignment kernels for functional RNA analyses

KENGO SATO^{†1,†2,†3} YUTAKA SAITO^{†3}
and YASUBUMI SAKAKIBARA^{†3}

We have recently proposed novel kernel functions, called *base-pairing profile local alignment (BPLA) kernels* for discrimination and detection of functional RNA sequences using SVMs. We employ STRAL's scoring function which takes into account sequence similarities as well as upstream and downstream base-pairing probabilities, which enables us to model secondary structures of RNA sequences. In this paper, we develop a method for optimizing hyperparameters of BPLA kernels with respect to discrimination accuracy using a gradient-based optimization technique. Our experiments show that the proposed method can find a nearly optimal set of parameters much faster than the grid search on all parameter combinations.

1. Introduction

Recent research has discovered crucial roles of non-coding RNAs (ncRNAs) in cells, including post-transcriptional gene regulation, maturation of mRNAs, rRNAs and tRNAs. Most functional ncRNAs form secondary structures with hydrogen bonds such as Watson-Crick base-pairs (A-U and G-C) and Wobble base-pairs (G-U). It is well known that functions of ncRNAs are deeply related to their secondary structures rather than primary sequences. For example, transfer RNAs registered in the Rfam database⁽⁶⁾ have only 44 % average sequence identity, whereas they are folded into conserved clover leaf secondary structures.

Therefore, ncRNAs are modeled by means of their secondary structures for computational ncRNA analysis. One of such models is stochastic context-free grammars (SCFGs), which can represent RNA secondary structures without pseudoknots^(5),15). These grammatical methods have succeeded in modeling typical secondary structures of RNAs, and are commonly used for structural alignment of RNA sequences. However, such stochastic models are not capable of discriminating the member sequences of RNA families from nonmembers with sufficiently high accuracy to detect non-coding RNA regions in genome sequences.

On the other hand, a compelling attraction of support vector machines (SVMs) and other kernel methods has been increasing in many research fields, including bioinformatics⁽¹⁸⁾, because of their robustness. For example, Saigo *et al.* have proposed local alignment kernels for amino acid sequences^(13),14), and their kernels with SVMs have outperformed other state-of-the-art methods in benchmarks for remote homology detection of amino acid sequences. Therefore, we have been considering using kernel methods, including SVMs, for the analysis of functional ncRNAs^(12),16),17).

We have recently proposed novel kernel functions, called *base-pairing profile local alignment (BPLA) kernels* for discrimination and detection of functional RNA sequences using SVMs⁽¹²⁾. We extend the concept of local alignment kernels in such a way that it can handle RNA sequences using STRAL's scoring function⁽³⁾. The local alignment kernels measure the similarity between two sequences by summing the scores over all possible local alignments with gaps, whereas STRAL's scoring function takes into account sequence similarities as well as upstream and downstream base-pairing probabilities, which enables us to model secondary structures of RNA sequences. Our experiments have proved powerful discrimination ability of BPLA kernels for functional RNA families. Furthermore, we have demonstrated the applicability of our kernels to the problem of genome-wide search of snoRNA families in the *C. elegans* genome, and have confirmed that the expression is valid in 14 out of 48 of our predicted candidates by using qRT-PCR (p -value = 2.8×10^{32}).

Our previous experiments showed that parameters for the BPLA kernels are quite different depending on what kind of data (flanking regions or not) and what family of sequences should be trained. Our previous study employed a brute

†1 (社) バイオ産業情報化コンソーシアム
Japan Biological Informatics Consortium (JBIC)
†2 (独) 産業技術総合研究所生命情報工学研究センター
Computational Biology Research Center (CBRC), National Institute of Advanced Industrial
Sciences and Technology (AIST)
†3 慶應義塾大学理工学部生命情報学科
Department of Biosciences and Informatics, Keio University

force approach in which optimal parameters for each case were calibrated by the grid search on all parameter combinations from selected search space for each parameter. It is easily predicted that this approach should require exponential time in the number of parameters to be optimized.

Therefore, we propose a gradient-based parameter optimization for BPLA kernels. This work is based on Keerthi *et al.*⁽⁸⁾, in which gradients of accuracy measures (e.g. AUC) with respect to hyperparameters of kernel functions are calculated, then these hyperparameters are optimized with a numerical optimization algorithm such as the quasi-Newton's method. We develop a dynamic programming algorithm for calculating the gradients of BPLA kernels with respect to their parameters to be optimized by the Keerthi's method. Our experiments show that the proposed method can find a nearly optimal set of parameters much faster than the grid search on all parameter combinations.

This paper is organized as follows: In section 2, we first introduce BPLA kernels for functional RNA analysis and dynamic programming algorithms for calculating BPLA kernels and their gradient with respect to their hyperparameters. Then, we briefly describe the Keerthi's gradient-based optimization method. In section 3, computational experiments are shown to confirm the efficiency of our method, and we discuss our framework in section 4. Finally, we conclude this paper in section 5.

2. Methods

2.1 Base-pairing profile local alignment kernels

Local alignment kernels can calculate the similarity between a pair of sequences by taking into account only sequence homology⁽¹⁴⁾. However, since it is well known that ncRNAs form secondary structures which are related to their functions, it is necessary to consider the secondary structures of ncRNAs in order to compare two RNA sequences. To do this, we have proposed stem kernels⁽¹⁶⁾ and their approximation using directed acyclic graphs (DAGs) of ncRNAs⁽¹⁷⁾. The stem kernels can consider RNA secondary structures strictly, but require huge computational time for practical problems. Differing from them, BPLA kernels employ a kind of summary information of secondary structures, called *base-pairing profiles*, as well as STRAL⁽³⁾.

For each sequence, we first calculate a base-pairing probability matrix using the McCaskill algorithm⁽¹¹⁾ or the inside-outside algorithm⁽⁴⁾. The base-pairing probability matrix for a sequence \mathbf{x} consists of the base-pairing probabilities P_{ij} that the i -th and the j -th nucleotides of \mathbf{x} form a base pair, which is defined as:

$$P_{ij} = \mathbb{E}[I_{ij}|\mathbf{x}] = \sum_{y \in \mathcal{Y}(\mathbf{x})} p(y|\mathbf{x}) I_{ij}(y), \quad (1)$$

where $\mathcal{Y}(\mathbf{x})$ is an ensemble of all possible secondary structures of \mathbf{x} , $p(y|\mathbf{x})$ is the posterior probability of a secondary structure y given \mathbf{x} , and $I_{ij}(y)$ is an indicator function, which equals 1 if the i -th and the j -th nucleotides form a base-pair in y and 0 otherwise. In this study, we employ the Vienna RNA package⁽⁷⁾ for computing the expected counts (1) using the McCaskill algorithm.

Subsequently, for each position i of \mathbf{x} , the base-pairing probabilities are summarized into tree kinds of sums: the probability $P_i^{left} = \sum_{j>i} P_{ij}$ that a pair is formed with one of the downstream nucleotides, the probability $P_i^{right} = \sum_{j<i} P_{ji}$ that a pair is formed with one of the upstream nucleotides, and the probability $P_i^{unpair} = 1 - (P_i^{left} + P_i^{right})$ that the nucleotide is unpaired. A probability distribution consisting of these three probabilities is called a *base-pairing profile*⁽¹⁾.

In accordance with STRAL⁽³⁾, we define the match score between two nucleotides x_i and x_j using base-pairing profiles as follows:

$$\begin{aligned} S(x_i, y_j) &= \alpha(S_{struct}) + S_{seq} \\ &= \alpha \left(\sqrt{P_{x_i}^{left} P_{y_j}^{left}} + \sqrt{P_{x_i}^{right} P_{y_j}^{right}} \right) + \sqrt{P_{x_i}^{unpair} P_{y_j}^{unpair}} S'(x_i, y_j), \end{aligned} \quad (2)$$

where α is a weight parameter for structural information, and S' is a substitution scoring function between two nucleotides. Then, we employ equation (2) as the match scores for the local alignment kernels, which is defined as follows:

$$K(\mathbf{x}, \mathbf{y}) = \sum_{\pi \in \Pi(\mathbf{x}, \mathbf{y})} \exp \beta s(\mathbf{x}, \mathbf{y}, \pi), \quad (3)$$

where $\beta \geq 0$ is a concentration parameter, $\Pi(\mathbf{x}, \mathbf{y})$ is a set of all possible local alignments of \mathbf{x} and \mathbf{y} , and $s(\mathbf{x}, \mathbf{y}, \pi)$ is a score of alignment π defined by:

$$s(\mathbf{x}, \mathbf{y}, \pi) = \sum_{a,b} n_{a,b}(\mathbf{x}, \mathbf{y}, \pi) \cdot S(a, b) - n_g(\mathbf{x}, \mathbf{y}, \pi) \cdot g - n_e(\mathbf{x}, \mathbf{y}, \pi) \cdot e, \quad (4)$$

where $n_{a,b}(\mathbf{x}, \mathbf{y}, \pi)$ is the number of times that nucleotides a is aligned to nu-

Initialization:

for $i \in \{0, \dots, |\mathbf{x}|\}$ and $j \in \{0, \dots, |\mathbf{y}|\}$ **do**
 $M(i, 0) = I_X(i, 0) = I_Y(i, 0) = R_X(i, 0) = R_Y(i, 0) = 0$
 $M(0, j) = I_X(0, j) = I_Y(0, j) = R_X(0, j) = R_Y(0, j) = 0$
end for

Iteration:

for $i \in \{1, \dots, |\mathbf{x}|\}$ and $j \in \{1, \dots, |\mathbf{y}|\}$ **do**
 $M(i, j) = e^{\beta S(x_i, y_j)} (1 + I_X(i-1, j-1) + I_Y(i-1, j-1) + M(i-1, j-1))$
 $I_X(i, j) = e^{\beta g} M(i-1, j) + e^{\beta e} I_X(i-1, j)$
 $I_Y(i, j) = e^{\beta g} (M(i, j-1) + I_X(i, j-1)) + e^{\beta e} I_Y(i, j-1)$
 $R_X(i, j) = M(i-1, j) + R_X(i-1, j)$
 $R_Y(i, j) = M(i, j-1) + R_X(i, j-1) + R_Y(i, j-1)$
end for

Termination:

$K(\mathbf{x}, \mathbf{y}) = 1 + R_X(|\mathbf{x}|, |\mathbf{y}|) + R_Y(|\mathbf{x}|, |\mathbf{y}|) + M(|\mathbf{x}|, |\mathbf{y}|)$

Fig. 1 The dynamic programming of the BPLA kernel.

cleotides b in π , $n_g(\mathbf{x}, \mathbf{y}, \pi)$ and $n_e(\mathbf{x}, \mathbf{y}, \pi)$ are the numbers of gap opens and gap extensions, respectively, and g and e are penalties for gap opens and gap extensions, respectively. We call equation (3) with the match scoring function (2) the *base-pairing profile local alignment (BPLA) kernel*. Equation (3) can be computed by the dynamic programming technique shown as Figure 1.

The hyperparameter optimization algorithm, described in Section 2.2, employs the gradient of $K(\mathbf{x}, \mathbf{y})$ with respect to parameters α, β, g and e , which can also be computed by the dynamic programming. Figure 2 shows the dynamic programming of calculating the derivative with respect to α . The derivatives with respect to the other parameters can be computed as well as α .

In order to avoid length bias, we usually normalize kernel values. Therefore, we would like to optimize hyperparameters with respect to the normalized kernels:

$$K_n(x, y) = \frac{K(x, y)}{\sqrt{K(x, x)K(y, y)}}$$

$$K'_n(x, y) = \frac{K'(x, y)}{\sqrt{K(x, x)K(y, y)}} - \frac{K_n(x, y)}{2} \left(\frac{K'(x, x)}{K(x, x)} + \frac{K'(y, y)}{K(y, y)} \right).$$

Initialization:

for $i \in \{0, \dots, |\mathbf{x}|\}$ and $j \in \{0, \dots, |\mathbf{y}|\}$ **do**
 $\frac{\partial}{\partial \alpha} M(i, 0) = \frac{\partial}{\partial \alpha} I_X(i, 0) = \frac{\partial}{\partial \alpha} I_Y(i, 0) = \frac{\partial}{\partial \alpha} R_X(i, 0) = \frac{\partial}{\partial \alpha} R_Y(i, 0) = 0$
 $\frac{\partial}{\partial \alpha} M(0, j) = \frac{\partial}{\partial \alpha} I_X(0, j) = \frac{\partial}{\partial \alpha} I_Y(0, j) = \frac{\partial}{\partial \alpha} R_X(0, j) = \frac{\partial}{\partial \alpha} R_Y(0, j) = 0$
end for

Iteration:

for $i \in \{1, \dots, |\mathbf{x}|\}$ and $j \in \{1, \dots, |\mathbf{y}|\}$ **do**
 $\frac{\partial}{\partial \alpha} M(i, j) = \beta M(i, j) \left(\sqrt{P_{x_i}^{left} P_{y_j}^{left}} + \sqrt{P_{x_i}^{right} P_{y_j}^{right}} \right)$
 $+ e^{\beta S(x_i, y_j)} \left(\frac{\partial}{\partial \alpha} I_X(i-1, j-1) + \frac{\partial}{\partial \alpha} I_Y(i-1, j-1) + \frac{\partial}{\partial \alpha} M(i-1, j-1) \right)$
 $\frac{\partial}{\partial \alpha} I_X(i, j) = e^{\beta g} \frac{\partial}{\partial \alpha} M(i-1, j) + e^{\beta e} \frac{\partial}{\partial \alpha} I_X(i-1, j)$
 $\frac{\partial}{\partial \alpha} I_Y(i, j) = e^{\beta g} \left(\frac{\partial}{\partial \alpha} M(i, j-1) + \frac{\partial}{\partial \alpha} I_X(i, j-1) \right) + e^{\beta e} \frac{\partial}{\partial \alpha} I_Y(i, j-1)$
 $\frac{\partial}{\partial \alpha} R_X(i, j) = \frac{\partial}{\partial \alpha} M(i-1, j) + \frac{\partial}{\partial \alpha} R_X(i-1, j)$
 $\frac{\partial}{\partial \alpha} R_Y(i, j) = \frac{\partial}{\partial \alpha} M(i, j-1) + \frac{\partial}{\partial \alpha} R_X(i, j-1) + \frac{\partial}{\partial \alpha} R_Y(i, j-1)$
end for

Termination:

$\frac{\partial}{\partial \alpha} K(\mathbf{x}, \mathbf{y}) = \frac{\partial}{\partial \alpha} R_X(|\mathbf{x}|, |\mathbf{y}|) + \frac{\partial}{\partial \alpha} R_Y(|\mathbf{x}|, |\mathbf{y}|) + \frac{\partial}{\partial \alpha} M(|\mathbf{x}|, |\mathbf{y}|)$

Fig. 2 The dynamic programming of the derivative of the BPLA kernel with respect to α .

2.2 Gradient-based optimization for SVMs

Tuning hyperparameters of kernel functions is known as an important task to achieve enough accuracy in practical problems. Although brute force approaches like the grid search are usually employed for such tasks, huge computational time will obviously be required if the search space of hyperparameters is large. Keerthi *et al.* have proposed a more efficient method which optimizes hyperparameters via the gradient-based optimization⁸⁾.

We assume that n training data $\{(x_i, y_i)\}_{i=0}^n$ ($y_i \in \{-1, +1\}$) are given. In SVM training, w and b are optimized to satisfy the following condition:

$$\min_{w, b} \frac{1}{2} \|w\|^2 + C \sum_i l(o_i, y_i),$$

where l is a loss function such as the hinge-loss $l(o_i, y_i) = \max\{0, 1 - y_i o_i\}$, C is a weight for the loss function, and o_i is the output of the SVM for the training data:

$$o_i = w \cdot \phi(x_i) - b = \sum_j \alpha_j y_j K(x_i, x_j) - b.$$

We can obtain a linear system from the Karush-Kuhn-Tucker (KKT) condition:

$$P \begin{pmatrix} \alpha \\ b \end{pmatrix} = q. \quad (5)$$

The training data and their coefficients α are divided into three parts: a part I_0 for $\alpha_i = 0$, a part I_u for $0 < \alpha_i < C$, and a part I_c for $\alpha_i = C$. The linear system (5) can be rewritten as follows:

$$\alpha_0 = 0, \quad \alpha_c = Ce_c, \quad \begin{pmatrix} \Omega_{uu} & -y_u \\ -y_u^T & 0 \end{pmatrix} \begin{pmatrix} \alpha_u \\ b \end{pmatrix} = \begin{pmatrix} e_u - \Omega_{uc}\alpha_c \\ y_c^T \alpha_c \end{pmatrix},$$

where $\Omega_{ij} = y_i y_j K_{ij}$ and e_u is a unit vector. These α and b can be computed by solving quadratic programming.

For given validation data $\{(\bar{x}_l, \bar{y}_l)\}$, we can calculate:

$$\bar{o}_l = \sum_i \alpha_i y_i K(\bar{x}_l, x_i) - b,$$

and can rewrite it with a matrix form:

$$\bar{o}_l = \psi_l^T \beta,$$

where β is a vector containing α and b , and ψ_l is a vector containing $y_i K_{li}, i = 1, \dots, n$ and -1 as the last element (corresponding to b).

Our objective is to maximize f (which is a kind of accuracy measures such as specificity, sensitivity, F-value, or AUC) with respect to hyperparameters θ of a given kernel function K . Here, \dot{f} denotes a derivative of f with respect to θ , that is, $\dot{f} = \frac{\partial f}{\partial \theta}$, which can be computed as:

$$\begin{aligned} \dot{f} &= \sum_l \delta_l \dot{\bar{o}}_l = \sum_l \delta_l \left(\psi_l^T P^{-1} (\dot{q} - \dot{P}\beta) + \dot{\psi}_l^T \beta \right) \\ &= d^T (\dot{q} - \dot{P}\beta) + \left(\sum_l \delta_l \dot{\psi}_l \right)^T \beta, \end{aligned} \quad (6)$$

where d is the solution of

$$P^T d = \left(\sum_l \delta_l \psi_l \right), \quad (7)$$

and

$$\delta_l = \frac{\partial f}{\partial \bar{o}_l}.$$

Eq (7) can efficiently be solved by the conjugate gradient method. After f and its gradients \dot{f} are computed, we can apply the gradient-based optimization by the L-BFGS¹⁰⁾ (or the L-BFGS-B²⁾ with box-constraints).

We optimize the hyperparameters of BPLA kernels by means of AUC, that is, $f = auc$:

$$auc = \frac{1}{n^+ n^-} \sum_{i: \bar{y}_i = +1} \sum_{j: \bar{y}_j = -1} \text{sign}(\bar{o}_i - \bar{o}_j), \quad (8)$$

where n^+ and n^- are the numbers of the positive and negative examples in the validation dataset, respectively. The step function $\text{sign}(x)$ is approximated by using the sigmoidal function $s(x) = \frac{1}{1 + \exp(-\sigma x)}$ to make (8) differentiable.

3. Results

In order to illustrate the efficiency of our gradient-based optimization, we compared the proposed method with the brute-force grid search. In each experiment, we first calculated the AUC for all the parameter combinations in the predefined search space, and then performed our optimization procedure with the corresponding box constraints to evaluate whether it could approach the AUC value at the optimal grid point. We used the same data set as in our previous paper¹⁶⁾, which includes five ncRNA families taken from the Rfam database⁶⁾. For each families, 100 sequences were chosen as positive samples in such a way that the pairwise identity is not above 80% for any pair of sequences, and as negative samples, 100 randomly shuffled sequences with the same dinucleotide composition as the positives were generated. AUC was evaluated using four-fold cross-validation, and its gradient was computed by summing over the four partitions.

In our first experiment, we attempted to simultaneously optimize all the four parameters α , β , g , and e in BPLA kernels, as well as C in the loss functions of SVMs. To do this, we defined the search space with five-dimensional box constraints: $0 < \alpha \leq 20$, $0 < \beta \leq 0.2$, $-30 \leq g < 0$, $-1 \leq e < 0$, and $0 < C \leq 40$. Subsequently, we carried out the grid search in which we tried six values for each parameter at even intervals, *i.e.* $\alpha \in \{1 \times 10^{-9}, 4, 8, 12, 16, 20\}$, $\beta \in \{1 \times 10^{-9}, 0.04, 0.08, 0.12, 0.16, 0.20\}$, and similarly for the others. Then, we performed the gradient-based optimization. Since the proposed method might fall into one of local maxima, differences in initial values of the parameters could change resulting AUC. Therefore, we decided to execute the optimization procedure from various start points. These start points were chosen by combining the first or third quantile of each parameter, *i.e.* $\alpha \in \{5, 15\}$, $\beta \in \{0.05, 0.15\}$,

and similarly for the others, and we evaluated the mean AUC value of all the 2^5 trials. Table 1 presents the experimental results. Although the grid search attained the best discrimination for all the ncRNA families tested, it took the huge computational time because of brute-force calculation for the 7776 ($= 6^5$) combinations of the parameters. Note that even the relatively coarse grid with only six values for one dimension required such massive computation. This clearly shows that the grid search is not a practical way to adapt the parameters of BPLA kernels against various kinds of data set. On the other hand, our gradient-based optimization successfully found the nearly optimal set of the parameters much faster than the grid search. The AUC was substantially improved from the start points on average, and convergence was obtained within a few dozen of iterations or so. In particular, we achieved the comparable discrimination ability to the grid search for 5S rRNAs, tRNAs and miRNAs showing the effectiveness of the proposed method.

The second experiment was organized to visually demonstrate how our method works in the parameter space. The idea was that we would show the trajectory of the optimization procedure on the heat map of AUC. For this purpose, we focused only on α and β , and fixed the other parameters. Thus, we defined the same box constraints for α and β as in the first experiment, whereas we consistently used $g = 27$, $e = 1$, and $C = 1$ for the others. We calculated the AUC values over the finer grid with 100 values for each of α and β , and initiated the gradient-based optimization from the 2^2 different points (α, β) chosen by the same manner as in the previous experiment. Figure 3 shows the corresponding heat maps and the trajectory plots with the different colors according to their start points. We only present the typical results for tRNAs and C/D snoRNAs. In the results for tRNAs, β moved sharply through the optimization procedure, and converged into the well calibrated values, allowing all the trajectories to terminate in the region of the highest AUC. To the contrary, the moves of α were much more limited except the blue trajectory which once seemed to terminate near the convergence point of red one, but suddenly slipped to the boundary of the box constraints. This observation suggests that our method might not work well when the gradient values of some parameter are much smaller than that of the others. The same problem was more obvious in the results of C/D snoRNA,

Table 1 Comparison between the gradient-based optimization and the grid search. Family: name of the target ncRNA family. N: number of hyperparameter sets tried. AUC: area under the ROC curve, the value in parentheses is the AUC calculated at the start point of the optimization. time: computational time on a 2.8 GHz AMD Opteron processor. In the columns of the gradient-based optimization, mean values of the 2^5 trials from the different start points are shown.

Family	gradient-based optimization			grid search		
	N	AUC	time (h)	N	AUC	time (h)
5S rRNA	20	0.996 (0.956)	2.1	7776	0.997	149.0
tRNA	16	0.992 (0.957)	0.8	7776	0.995	67.5
C/D snoRNA	24	0.837 (0.773)	2.0	7776	0.887	124.8
H/ACA snoRNA	17	0.873 (0.832)	3.0	7776	0.906	248.0
miRNA	20	0.994 (0.971)	1.2	7776	0.995	89.0

where all the trajectories were caught in local optima because of insufficient adaptation of α , while β was successfully adjusted under the given α .

4. Discussion

Our current implementation of the hyperparameter optimizer for BPLA kernels does not optimize the substitution scoring function for RNAs (S' in (2)) although derivatives of substitution scores can easily be calculated. In fact, Saigo *et al.* have proposed a method for optimizing substitution scores for local alignment kernels using the gradient decent with respect to a confidence score for remote homology search¹³⁾ (not with respect to the SVM classification). However, the optimization of the substitution scoring matrix using the L-BFGS-B, which our method employed, cannot hold the positive semidefiniteness of S' and BPLA kernels, which is the essential basis of kernel methods. This kind of optimization problems can be solved by a method for semidefinite programming (SDP)¹⁹⁾. Because of small effect of changing substitution scoring matrices in our method (data not shown), we fixed the substitution scoring matrix to a modified RI-BOSUM 85-60⁹⁾, in which its smallest eigenvalue is subtracted from each of its diagonal elements in order to make it positive semidefinite.

5. Conclusion

We have developed a method for optimizing hyperparameters of BPLA kernels with respect to discrimination accuracy using a gradient-based optimization

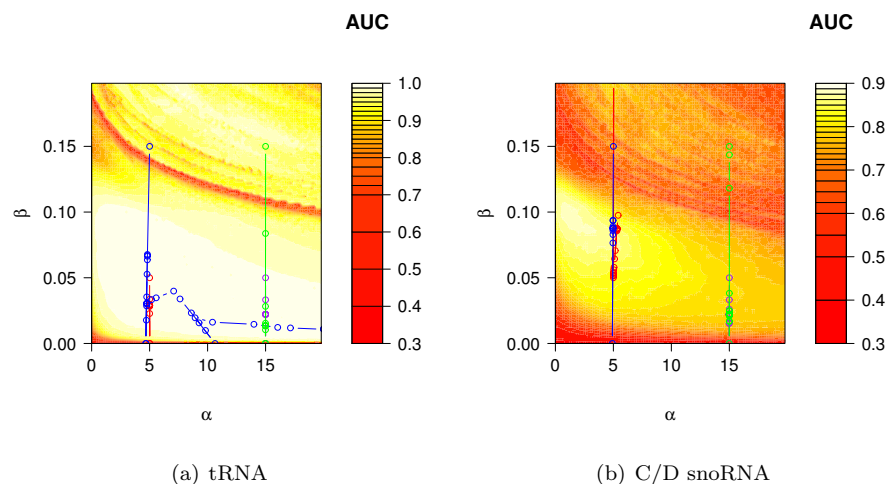


Fig. 3 Trajectories of the gradient-based optimization on the heat map of AUC. The trajectory plots are colored according to their start points (α, β) : red for (5, 0.05), blue for (5, 0.15), purple for (15, 0.05), and green for (15, 0.15). (a) tRNA. The trajectory of each color terminates, at (5.0, 0.029) for red; (20.0, 0.011) for blue; (15.0, 0.022) for purple; (15.0, 0.013) for green. (b) C/D snoRNA. The trajectory of each color terminates, at (5.3, 0.087) for red; (5.0, 0.088) for blue; (15.0, 0.022) for purple; (15.0, 0.022) for green.

technique. Our experiments showed that the proposed method can find a nearly optimal set of parameters much faster than the grid search on all parameter combinations.

We are planning to apply BPLA kernels to comparative analysis of functional RNAs, which can consider the sequence conservation between species as well as the structural conservation. For this purpose, the profile-profile comparison should be implemented.

Acknowledgments This work was supported in part by a grant from “Functional RNA Project” funded by the New Energy and Industrial Technology Development Organization (NEDO) of Japan, and was also supported in part by Grant-in-Aid for Scientific Research on Priority Area “Comparative Genomics” No. 17018029 from the Ministry of Education, Culture, Sports, Science and Tech-

nology of Japan. We thank our colleagues from the RNA Informatics Team at the Computational Biology Research Center (CBRC) for fruitful discussions.

References

- 1) Bonhoeffer, S., McCaskill, J. S., Stadler, P. F. and Schuster, P.: RNA multi-structure landscapes. A study based on temperature dependent partition functions., *Eur Biophys J*, Vol.22, No.1, pp.13–24 (1993).
- 2) Byrd, R.H., Lu, P., Nocedal, J. and Zhu, C.: A limited memory algorithm for bound constrained optimization, *SIAM Journal on Scientific and Statistical Computing*, Vol.16, No.5, pp.1190–1208 (1995).
- 3) Dalli, D., Wilm, A., Mainz, I. and Steger, G.: STRAL: progressive alignment of non-coding RNA using base pairing probability vectors in quadratic time., *Bioinformatics*, Vol.22, No.13, pp.1593–1599 (2006).
- 4) Durbin, R., Eddy, S., Krogh, A. and Mitchison, G.: *Biological Sequence Analysis*, Cambridge University Press (1998).
- 5) Eddy, S.R. and Durbin, R.: RNA sequence analysis using covariance models., *Nucleic Acids Res*, Vol.22, No.11, pp.2079–2088 (1994).
- 6) Gardner, P.P., Daub, J., Tate, J.G., Nawrocki, E.P., Kolbe, D.L., Lindgreen, S., Wilkinson, A.C., Finn, R.D., Griffiths-Jones, S., Eddy, S.R. and Bateman, A.: Rfam: updates to the RNA families database., *Nucleic Acids Res*, Vol.37, No.Database issue, pp.D136–D140 (2009).
- 7) Hofacker, I.L.: Vienna RNA secondary structure server., *Nucleic Acids Res*, Vol.31, No.13, pp.3429–3431 (2003).
- 8) Keerthi, S.S., Sindhiani, V. and Chapelle, O.: An efficient method for gradient-based adaptation of hyperparameters in SVM Models, *Advances in Neural Information Processing Systems 19* (Schölkopf, B., Platt, J. and Hoffman, T., eds.), Cambridge, MA, MIT Press, pp.673–680 (2007).
- 9) Klein, R.J. and Eddy, S.R.: RSEARCH: finding homologs of single structured RNA sequences., *BMC Bioinformatics*, Vol.4, p.44 (2003).
- 10) Liu, D.C. and Nocedal, J.: On the limited memory BFGS method for large-scale optimization, *Mathematical Programming*, Vol.45, pp.503–528 (1989).
- 11) McCaskill, J.S.: The equilibrium partition function and base pair binding probabilities for RNA secondary structure., *Biopolymers*, Vol.29, No.6-7, pp.1105–1119 (1990).
- 12) Morita, K., Saito, Y., Sato, K., Oka, K., Hotta, K. and Sakakibara, Y.: Genome-wide searching with base-pairing kernel functions for noncoding RNAs: computational and expression analysis of snoRNA families in *Caenorhabditis elegans*., *Nucleic Acids Res*, Vol.37, No.3, pp.999–1009 (2009).
- 13) Saigo, H., Vert, J.-P. and Akutsu, T.: Optimizing amino acid substitution matrices with a local alignment kernel., *BMC Bioinformatics*, Vol.7, p.246 (2006).

- 14) Saigo, H., Vert, J.-P., Ueda, N. and Akutsu, T.: Protein homology detection using string alignment kernels., *Bioinformatics*, Vol.20, No.11, pp.1682–1689 (2004).
- 15) Sakakibara, Y., Brown, M., Hughey, R., Mian, I.S., Sjölander, K., Underwood, R.C. and Haussler, D.: Stochastic context-free grammars for tRNA modeling., *Nucleic Acids Res*, Vol.22, No.23, pp.5112–5120 (1994).
- 16) Sakakibara, Y., Pependorf, K., Ogawa, N., Asai, K. and Sato, K.: Stem kernels for RNA sequence analyses., *J Bioinform Comput Biol*, Vol.5, No.5, pp.1103–1122 (2007).
- 17) Sato, K., Mituyama, T., Asai, K. and Sakakibara, Y.: Directed acyclic graph kernels for structural RNA analysis., *BMC Bioinformatics*, Vol.9, No.1, p.318 (2008).
- 18) Schölkopf, B., Tsuda, K. and Vert, J.P.: *Kernel Methods in Computational Biology*, MIT Press, Cambridge, MA (2004).
- 19) Vandenberghe, L. and Boyd, S.: Semidefinite programming., *SIAM Review*, Vol.38, pp.49–95 (1996).