

Learning interpretable SVMs for biological sequence classification

Sören Sonnenburg[†], Gunnar Rätsch^{†,‡}, Christin Schäfer[†]



Fraunhofer Institut
Rechnerarchitektur
und Softwaretechnik



MAX-PLANCK-GESELLSCHAFT

† Fraunhofer FIRST.IDA, Kekuléstr. 7, 12489 Berlin, Germany
‡ Friedrich Miescher Laboratory of the Max Planck Society,
Spemannstr. 39, 72076 Tübingen, Germany
{Soeren.Sonnenburg, Christin.Schaefer}@first.fraunhofer.de,
Gunnar.Raetsch@tuebingen.mpg.de




ROADMAP:

- The Motivating Application - Splice Site Prediction
- Multiple Kernel Learning (MKL)
- Derivation of the MKL Optimization Problem
- Algorithms
- Significance Analysis
- Results
- Outlook and Conclusion

THE MOTIVATING APPLICATION

Splice sites are locations on DNA at boundaries of

- **exons** (which code for proteins)
- **introns** (which do not)



DETECTION OF SPLICE SITES

Focus on Acceptor Splice Site Prediction

Intron Exon

```
AAACAAATAAGTAACTAATCTTTTAGGAAGAACGTTTCAACCATTTTGAG
AAGATTAAAAAAAAAACAAATTTTTCATTACAGATATAATAATCTAATT
CACTCCCAAATCAACGATATTTTAGTTCACTAACACATCCGTCTGTGCC
TTAATTTCACTTCCACATACTTCCAGATCATCAATCTCCAAAACCAACAC
TTGTTTTAATATTCAATTTTTTACAGTAAGTTGCCAATTCAATGTTCCAC
TACCTAATTATGAAATTAAAATTCAGTGTGCTGATGGAAACGGAGAAGTC
```

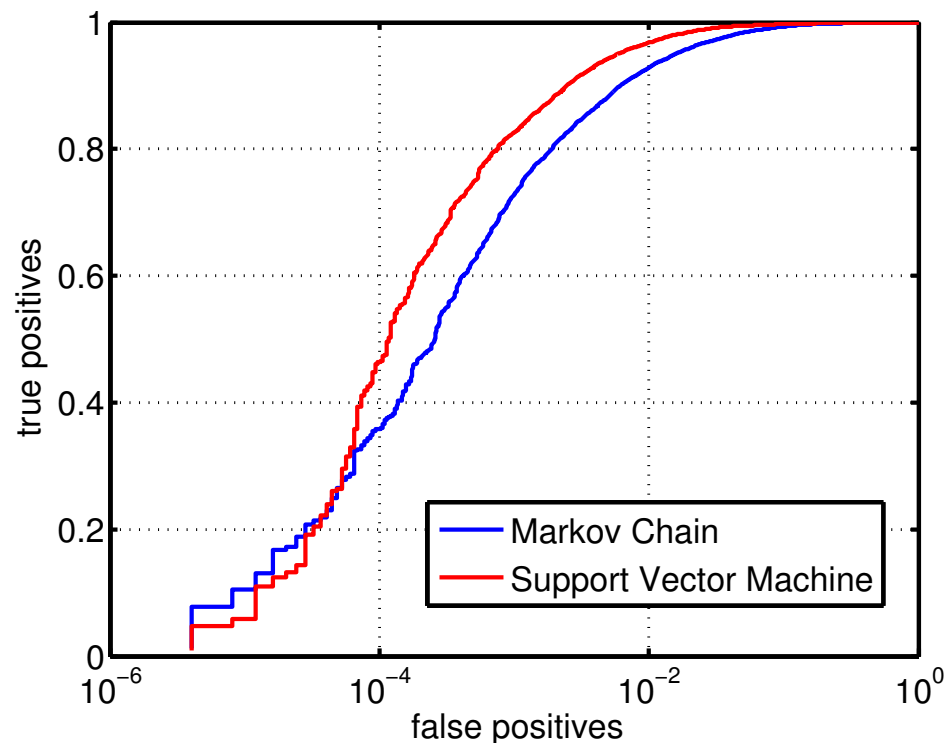
- aligned sequences of fixed length (AG always at same position)
- Task: distinguish *C. elegans* splice sites from fake splice sites

⇒ 2-class classification problem

STANDARD APPROACH: MARKOV CHAIN

Standard Approach: Use two higher order Markov Chains

- one model for fake splice sites and one for true splice sites trained on $\approx 600,000$ examples, of them only $\approx 36,000$ true splice sites




We can do better using SVMs!

Drawback: We loose interpretability of the result!

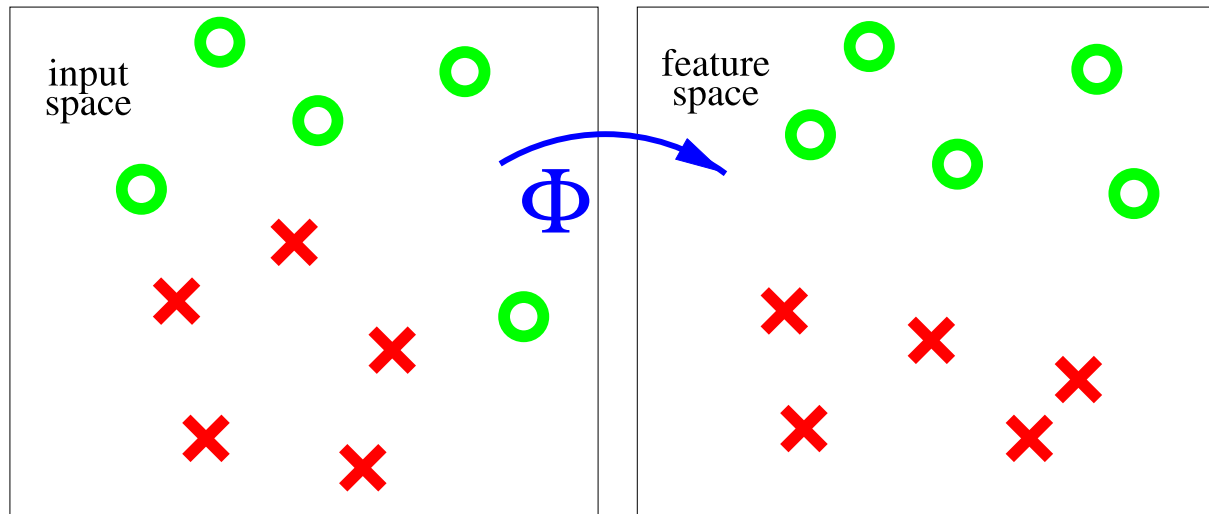
SUPPORT VECTOR MACHINE

- given: points (e.g. sequences) $s_i \in \mathcal{S}$ ($i = 1, \dots, N$) with respective labels $y_i \in \{-1, +1\}$
- in training hyperplane that maximizes **margin** is chosen

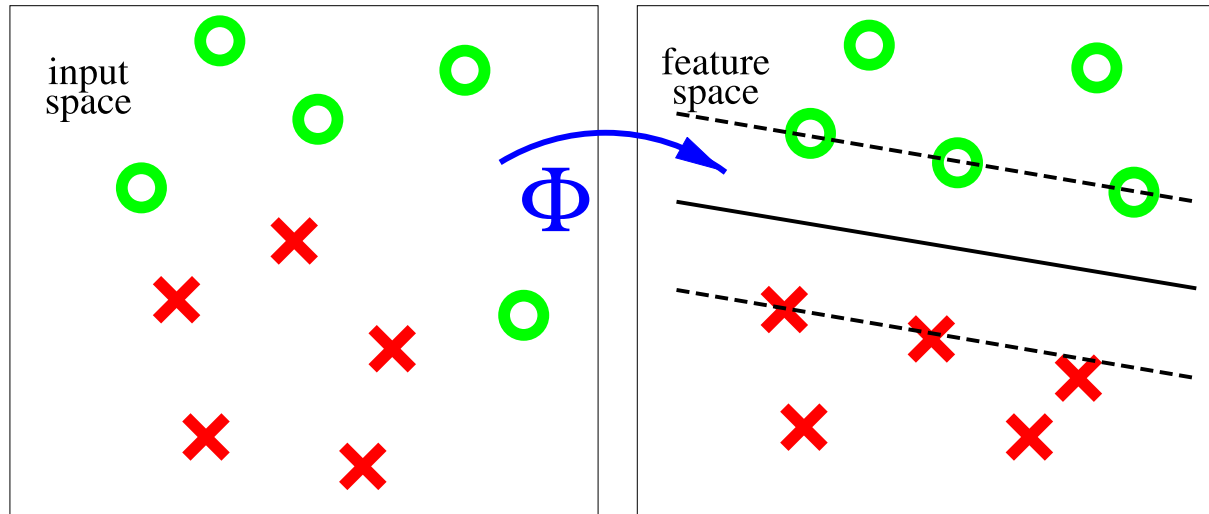


$$\text{Decision function } f(s) = w \cdot s + b$$

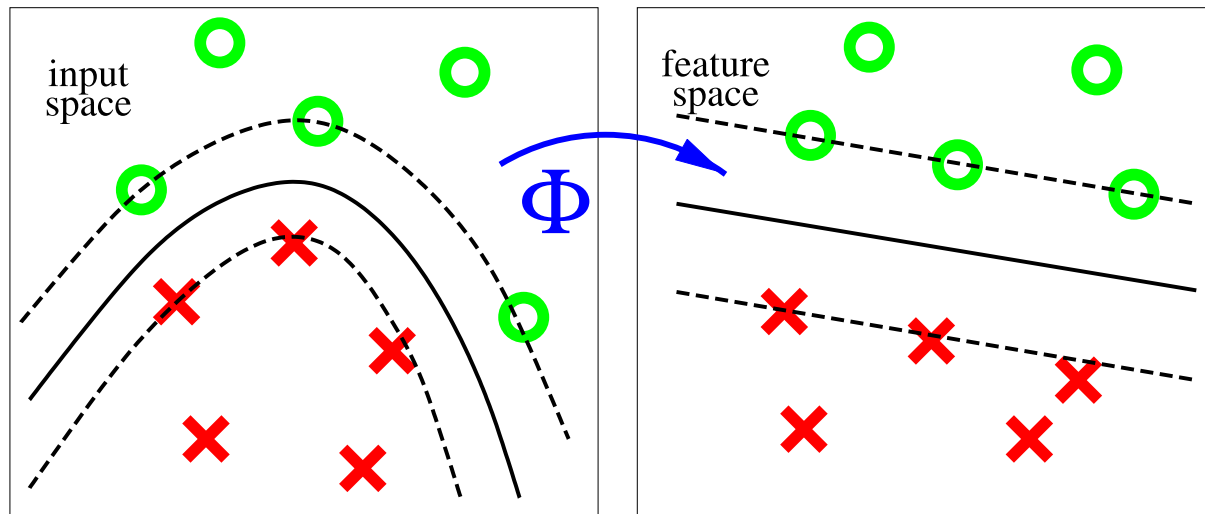
SVM WITH KERNELS



SVM WITH KERNELS



SVM WITH KERNELS



- SVM decision function in kernel feature space:

$$f(\mathbf{s}) = \sum_{i=1}^N y_i \alpha_i \underbrace{\Phi(\mathbf{s}) \cdot \Phi(\mathbf{s}_i)}_{=k(\mathbf{s}, \mathbf{s}_i)} + b \quad (1)$$


- find parameters α by solving quadratic optimization problem

Problem: Decision function (1) is hard to interpret


UNDERSTANDING THE SVM DECISION

Splice Sites

1. Which positions in the sequence are important for discrimination?



2. What characterizes those positions?



3. Which motifs at which position are important?

APPROACH: OPTIMIZE COMBINATION OF KERNELS

- Define Kernel as Convex Combination of Subkernels:

$$\mathbf{k}(\mathbf{s}, \mathbf{s}') = \sum_{l=1}^L \beta_l \mathbf{k}_l(\mathbf{s}, \mathbf{s}')$$

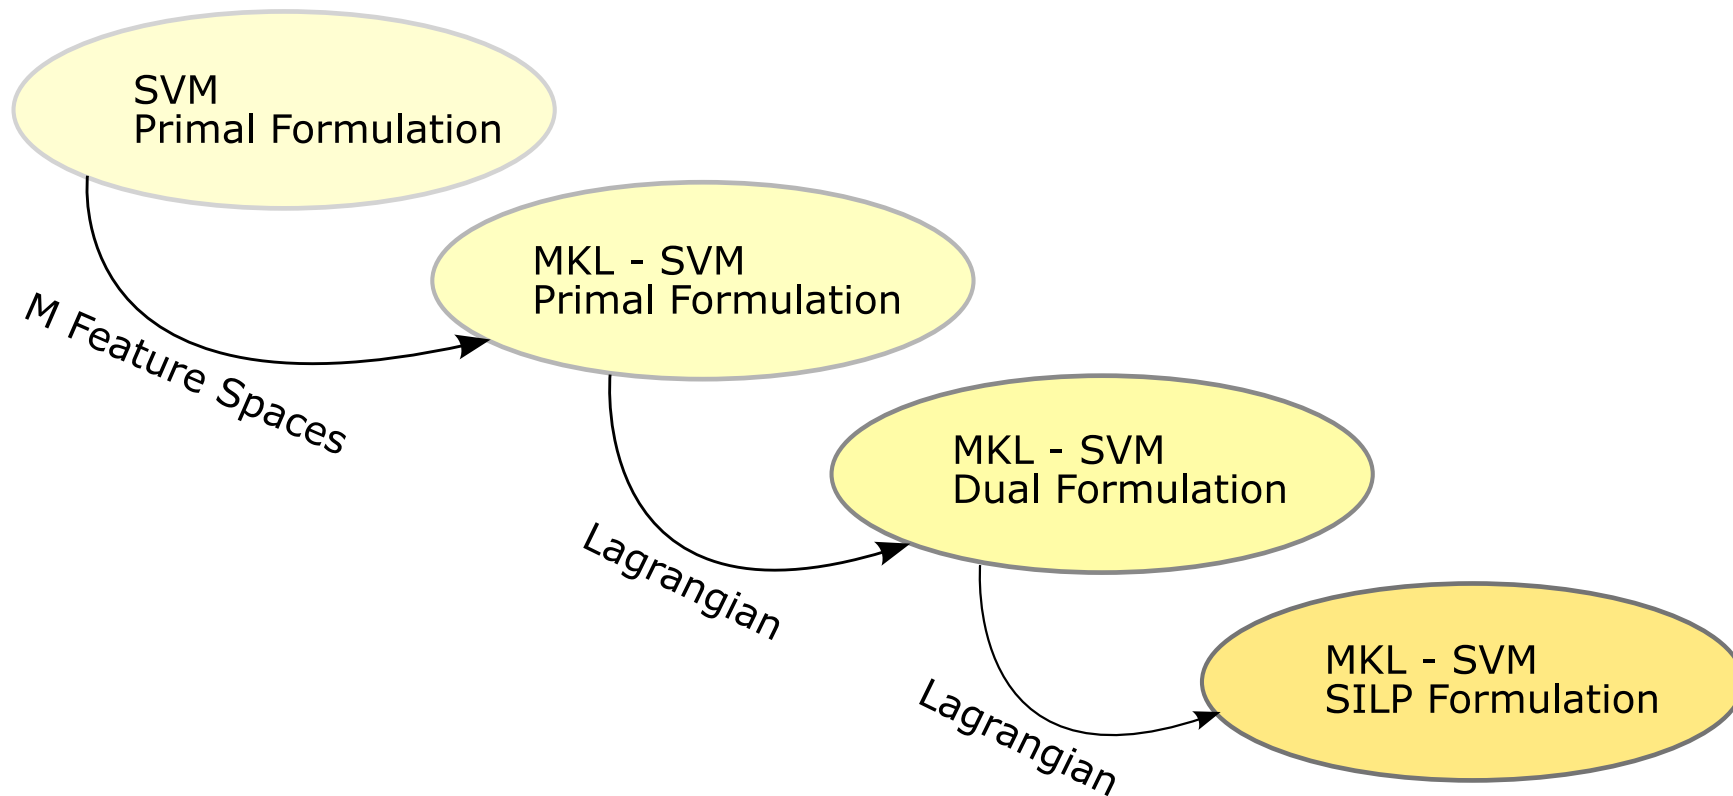
e.g. Weighted Degree Kernel

$$\mathbf{k}(\mathbf{s}, \mathbf{s}') = \sum_{l=1}^L \beta_l \sum_{k=1}^d \mathbb{I}(\mathbf{u}_{k,l}(\mathbf{s}) = \mathbf{u}_{k,l}(\mathbf{s}'))$$

- optimize weights β such that margin is maximized
 - ⇒ determine (β, α, b) simultaneously
 - ⇒ **Multiple Kernel Learning** (Bach, Lanckriet and Jordan 2004)
 - ⇒ leads to Second Order Cone Problem - *very slow!*
 - ⇒ we propose new *efficient* algorithm

REWRITING THE MKL FORMULATION

From the Standard SVM Primal to the Semi-Infinite Linear Program:



MKL OPTIMIZATION PROBLEM

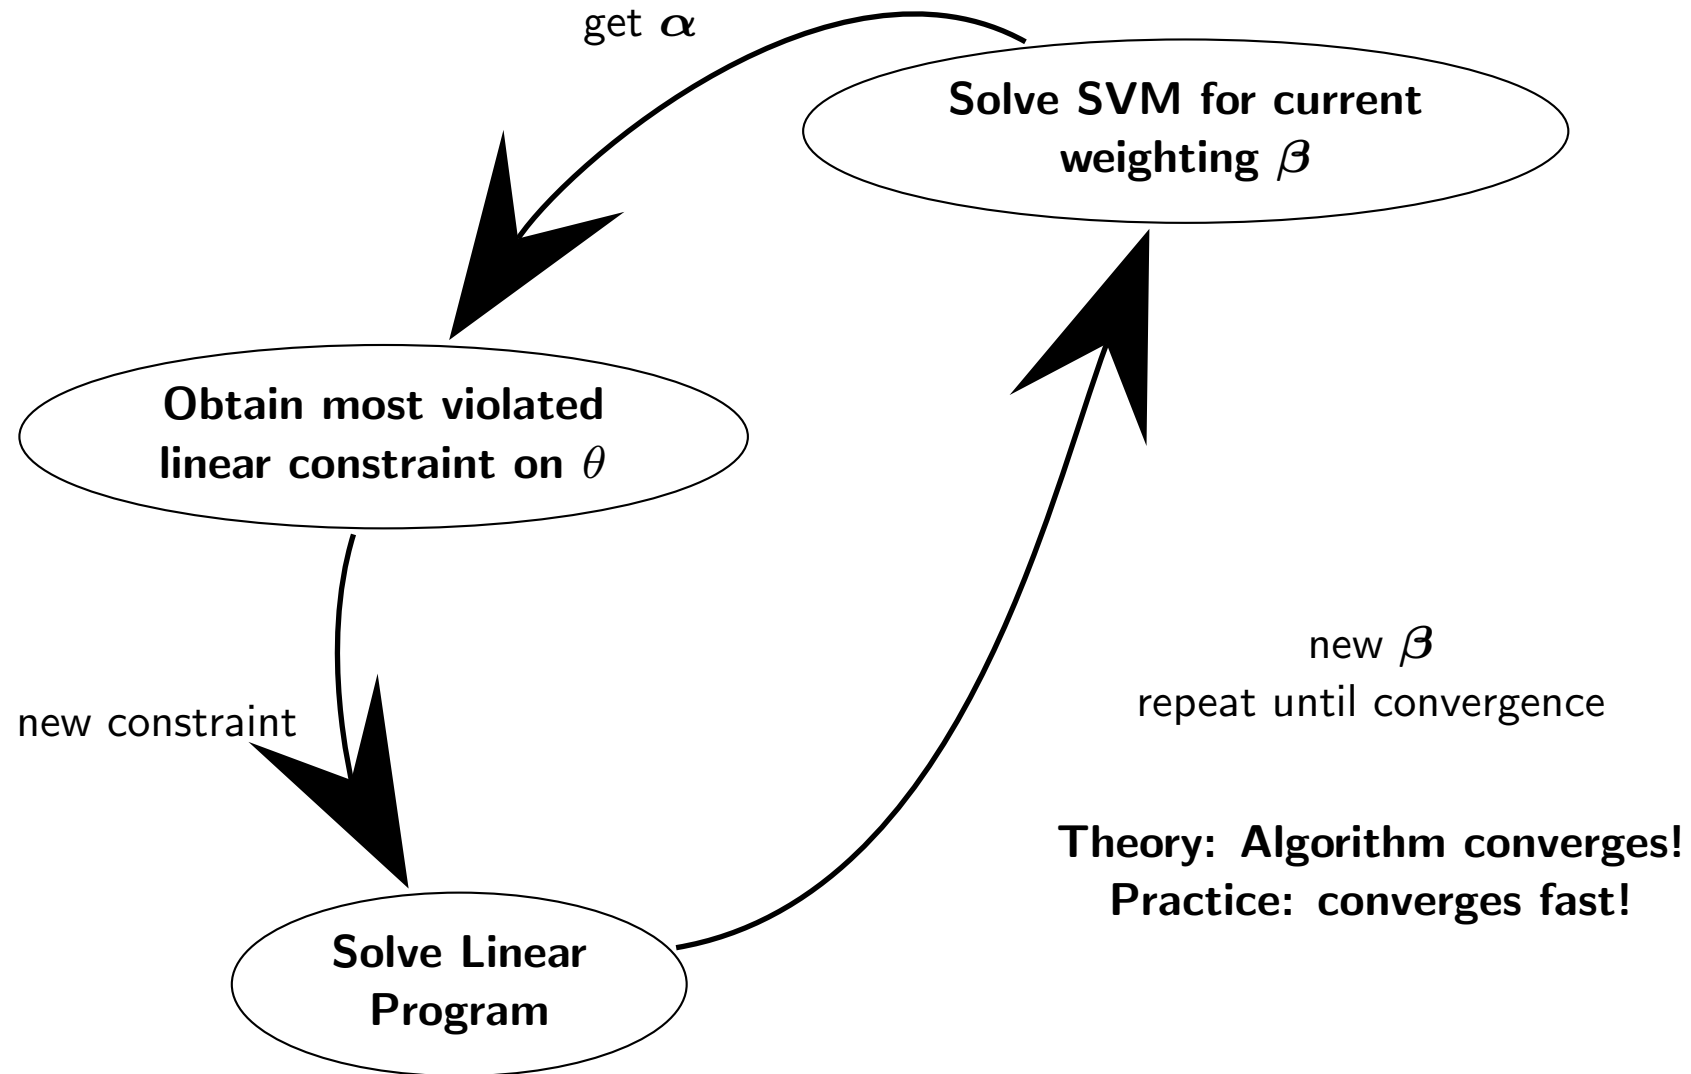
Semi-Infinite Linear Program Formulation:

$$\begin{array}{ll}
 \max & \theta \\
 \text{w.r.t.} & \theta \in \mathbb{R}, \beta \in \mathbb{R}_+^M \text{ with } \sum_{j=1}^M \beta_j = 1 \\
 \text{s.t.} & \sum_{j=1}^M \beta_j \left(\frac{1}{2} \sum_{r=1}^N \sum_{s=1}^N \alpha_r \alpha_s y_r y_s K_j(\mathbf{s}_r, \mathbf{s}_s) - \sum_{i=1}^N \alpha_i \right) \geq \theta \\
 & \text{for all } \alpha \text{ with } 0 \leq \alpha \leq C \text{ and } \sum_{i=1}^N y_i \alpha_i = 0
 \end{array}$$

⇒ **Linear Problem, but infinitely many constraints**

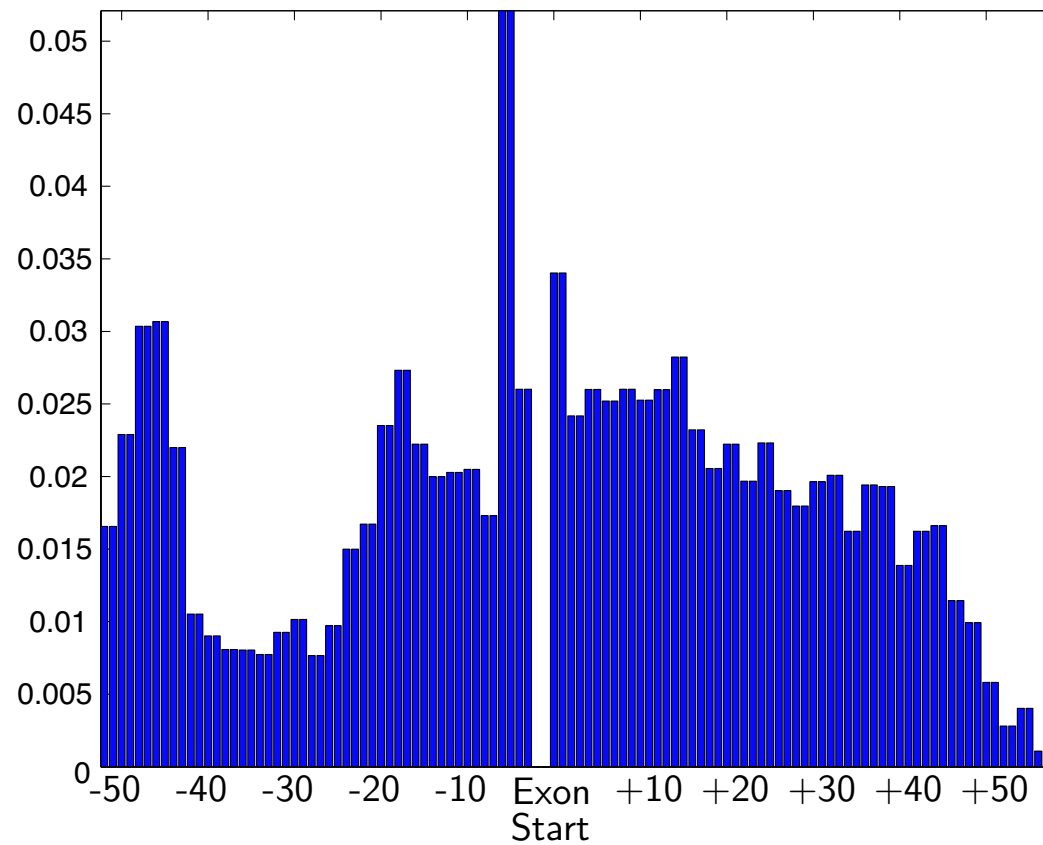
- one constraint for each α
- can be solved using Column Generation
- most violated constraints can be identified by solving SVM

COLUMN GENERATION




RESULTS ON SPLICE SITES

- Which positions in the sequence are important for discrimination:



RESULTS ON SPLICE SITES

What characterizes the positions:




- Can we use that weighting to interpret the SVM solution? \Rightarrow **Not yet!**
 - Stability of the weighting β ?
 - Which weights are significant?

Use a statistical test to investigate significance of weights

RESULTS ON SPLICE SITES

What characterizes the positions:

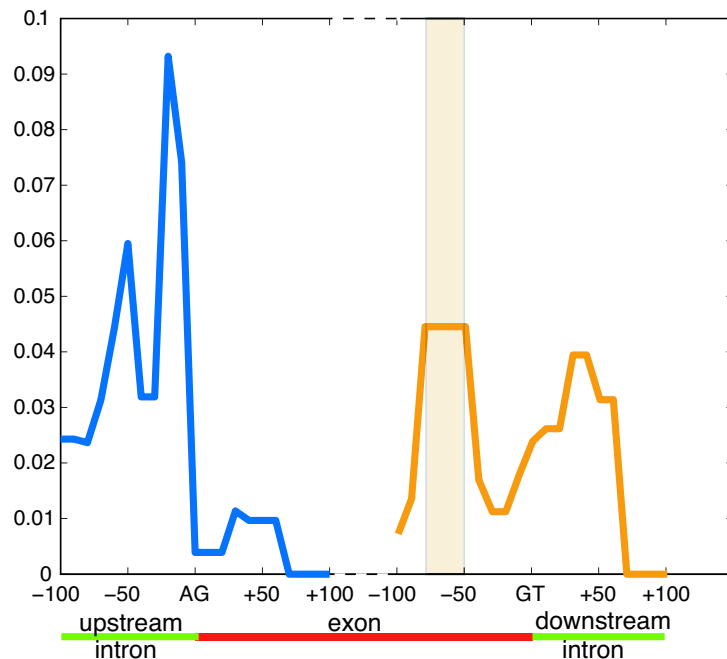


- Can we use that weighting to interpret the SVM solution? \Rightarrow **NOW !**
 - Stability of the weighting β ?
 - Which weights are significant?

Use a statistical test to investigate significance of weights

MKL TO UNDERSTAND ALTERNATIVE SPLICING

- Distinguish alternatively from constitutively spliced exons (ISMB 2005).
- Which positions are important? • Which motifs are important?



Hexamer	E-value
TTTAAA	1.8e-12
AATTTT	2.2e-10
ATTTTA	2.9e-09
CAGCAG	1.2e-08
TAATTT	8.3e-08
TTCCCC	2.1e-07

CONCLUSION

Conclusion:

- MKL learns convex combination of kernels
 - ⇒ allows for interpreting SVM result
 - ⇒ matches prior knowledge about splice sites
- simple iterative algorithm
- suitable for large scale problems ($> 100,000$ examples)

More information:

http://ida.first.fhg/~sonne/mkl_splice

Acknowledgements: Partial support from the PASCAL Network of Excellence (EU #506778), DFG grants JA379/13-2 and MU987/2-1. Thanks to Alexander Zien and K.-R. Müller.

SIGNIFICANCE ANALYSIS

- T bootstrap replicates of dataset
- train T times and obtain $\beta_{1\dots T}$ kernel weightings

- Bernoulli variable

$$X_{k,i}^t = \begin{cases} 1, & \beta_{k,i}^t > \tau \\ 0, & \text{else} \end{cases}.$$

- Testhypothesis \mathcal{H}_0 : The weight $\beta_{k,i}$ is not significant!

$$0.05 = \alpha \geq P_{\mathcal{H}_0}(\text{reject } \mathcal{H}_0) = P_{\mathcal{H}_0}(Z_{k,i} > v) = \sum_{j=v}^T \binom{T}{j} \hat{p}_0^j (1 - \hat{p}_0)^{T-j}$$

(where $Z_{k,i} = \sum_{t=1}^T X_{k,i}^t$ and $\hat{p}_0 = \#(\beta_{k,i}^t > \tau) / T \cdot M$)

STANDARD SVM OPTIMIZATION PROBLEM

This we all know

SVM Primal formulation:

$$\begin{aligned} \min \quad & \frac{1}{2} \|\mathbf{w}\|_2^2 + C \sum_{i=1}^N \xi_i \\ \text{w.r.t.} \quad & \mathbf{w} \in \mathbb{R}^k, \boldsymbol{\xi} \in \mathbb{R}_+^N, b \in \mathbb{R} \\ \text{s.t.} \quad & y_i (\mathbf{w}^\top \Phi(\mathbf{s}_i) + b) \geq 1 - \xi_i, \forall i = 1, \dots, N \end{aligned}$$

MKL OPTIMIZATION PROBLEM I

MKL Primal formulation:

$$\min \quad \frac{1}{2} \left(\sum_{j=1}^M \beta_j \|\mathbf{w}_j\|_2 \right)^2 + C \sum_{i=1}^N \xi_i$$

$$\text{w.r.t.} \quad \mathbf{w} = (\mathbf{w}_1, \dots, \mathbf{w}_M), \mathbf{w}_j \in \mathbb{R}^{k_j}, \boldsymbol{\xi} \in \mathbb{R}_+^N, \boldsymbol{\beta} \in \mathbb{R}_+^M, b \in \mathbb{R}$$

$$\text{s.t.} \quad y_i \left(\sum_{j=1}^M \beta_j \mathbf{w}_j^\top \Phi_j(\mathbf{s}_i) + b \right) \geq 1 - \xi_i, \forall i = 1, \dots, N$$

$$\sum_{j=1}^M \beta_j = 1$$

Properties: equivalent to SVM for $M = 1$; solution sparse in “blocks”; each block j corresponds to one kernel

MKL OPTIMIZATION PROBLEM II

Dual Formulation (Bach, Lanckriet, Jordan 2004):

$$\begin{aligned}
 \min \quad & \frac{1}{2}\gamma^2 - \sum_{i=1}^N \alpha_i \\
 \text{w.r.t.} \quad & \gamma \in \mathbb{R}, \boldsymbol{\alpha} \in \mathbb{R}^N \\
 \text{s.t.} \quad & 0 \leq \boldsymbol{\alpha} \leq C, \sum_{i=1}^N \alpha_i y_i = 0 \\
 & \underbrace{\sum_{r=1}^N \sum_{s=1}^N \alpha_r \alpha_s y_r y_s K_j(\mathbf{s}_r, \mathbf{s}_s)}_{=: S_j(\boldsymbol{\alpha})} - \gamma^2 \leq 0, \quad \forall j = 1, \dots, M
 \end{aligned}$$

“partial Lagrangian:”

$$L := \frac{1}{2}\gamma^2 - \sum_{i=1}^N \alpha_i + \sum_{j=1}^M \beta_j (S_j(\boldsymbol{\alpha}) - \gamma^2)$$

MKL OPTIMIZATION PROBLEM II

Reformulation as Semi-Infinite Linear Program:

$$\begin{aligned} & \max_{\beta} \min_{\alpha} \sum_{j=1}^M \beta_j \left(\frac{1}{2} S_j(\alpha) - \sum_{i=1}^N \alpha_i \right) \\ & \text{s.t. } 0 \leq \alpha \leq C, \sum_{i=1}^N \alpha_i y_i = 0, \sum_{j=1}^M \beta_j = 1 \end{aligned}$$

$$\begin{aligned} & \max \quad \theta \\ & \text{w.r.t.} \quad \theta \in \mathbb{R}, \beta \in \mathbb{R}_+^M \text{ with } \sum_{j=1}^M \beta_j = 1 \\ & \text{s.t.} \quad \sum_{j=1}^M \beta_j \left(\frac{1}{2} S_j(\alpha) - \sum_{i=1}^N \alpha_i \right) \geq \theta \\ & \quad \text{for all } \alpha \text{ with } 0 \leq \alpha \leq C \text{ and } \sum_{i=1}^N y_i \alpha_i = 0 \end{aligned}$$

⇒ **Linear, but infinitely many constraints**

THE SEMI-INFINITE LINEAR PROGRAM I

$$\begin{aligned}
 & \max && \theta \\
 & \text{w.r.t.} && \theta \in \mathbb{R}, \boldsymbol{\beta} \in \mathbb{R}_+^M \text{ with } \sum_{j=1}^M \beta_j = 1 \\
 & \text{s.t.} && \sum_{j=1}^M \beta_j \left(\frac{1}{2} S_j(\boldsymbol{\alpha}) - \sum_{i=1}^N \alpha_i \right) \geq \theta \\
 & && \text{for all } \boldsymbol{\alpha} \text{ with } 0 \leq \boldsymbol{\alpha} \leq C \text{ and } \sum_{i=1}^N y_i \alpha_i = 0
 \end{aligned}$$

Properties:

- optimize a convex combination
- infinitely many constraints
- quite easy to identify violated constraints

THE SEMI-INFINITE LINEAR PROGRAM II

$$\begin{aligned}
 & \max && \theta \\
 & \text{w.r.t.} && \theta \in \mathbb{R}, \boldsymbol{\beta} \in \mathbb{R}_+^M \text{ with } \sum_{j=1}^M \beta_j = 1 \\
 & \text{s.t.} && \sum_{j=1}^M \beta_j \left(\frac{1}{2} S_j(\boldsymbol{\alpha}) - \sum_{i=1}^N \alpha_i \right) \geq \theta \\
 & && \text{for all } \boldsymbol{\alpha} \text{ with } 0 \leq \boldsymbol{\alpha} \leq C \text{ and } \sum_{i=1}^N y_i \alpha_i = 0
 \end{aligned}$$

Solving the SILP:

- Column Generation
 - fast, but no convergence rate
- Use Boosting like techniques: Arc-GV or AdaBoost*
 - known convergence rate $\mathcal{O}(\log(M)/\varepsilon^2)$
- SMO like algorithm
 - consider suboptimal SVM solutions: empirically 3-5 times faster

SOLVING THE SILP: COLUMN GENERATION I

$$\begin{aligned}
 & \max && \theta \\
 & \text{w.r.t.} && \theta \in \mathbb{R}, \boldsymbol{\beta} \in \mathbb{R}_+^M \text{ with } \sum_{j=1}^M \beta_j = 1 \\
 & \text{s.t.} && \sum_{j=1}^M \beta_j \left(\frac{1}{2} S_j(\boldsymbol{\alpha}) - \sum_{i=1}^N \alpha_i \right) \geq \theta \\
 & && \text{for all } \boldsymbol{\alpha} \text{ with } 0 \leq \boldsymbol{\alpha} \leq C \text{ and } \sum_{i=1}^N y_i \alpha_i = 0
 \end{aligned}$$

- solved by taking set of most violated constraints into account
- most violated constraints given by SVM solution for fixed $\boldsymbol{\beta}$

$$\sum_{j=1}^M \beta_j \left(\frac{1}{2} S_j(\boldsymbol{\alpha}) - \sum_{i=1}^N \alpha_i \right) = \frac{1}{2} \sum_{r=1}^N \sum_{s=1}^N \alpha_r \alpha_s y_r y_s \sum_{j=1}^M \beta_j k_j(\mathbf{s}_r, \mathbf{s}_s) - \sum_{i=1}^N \alpha_i,$$

- iteratively find most violated constraints, solve linear program with current constraints, . . . , till convergence to the global optimum

SOLVING THE SILP: COLUMN GENERATION II

$$D^0 = 1, \theta^1 = 0, \beta_k^1 = \frac{1}{M} \text{ for } k = 1, \dots, M$$

for $t = 1, 2, \dots$ **do**

 obtain SVM's α^k with kernel

$$\mathbf{k}^t(\mathbf{s}_i, \mathbf{s}_j) := \sum_{k=1}^M \beta_k^t \mathbf{k}_k(\mathbf{s}_i, \mathbf{s}_j)$$

for $k = 1, \dots, M$ **do**

$$D_k^t = \frac{1}{2} \sum_{r,s} \alpha_r^t \alpha_s^t y_r y_s \mathbf{k}_k(\mathbf{s}_r, \mathbf{s}_s) - \sum_r \alpha_r^t$$

end for

$$D^t = \sum_{k=1}^M \beta_k^t D_k^t$$

$$(\boldsymbol{\beta}^{t+1}, \theta^{t+1}) = \operatorname{argmax} \theta$$

$$\text{w.r.t. } \boldsymbol{\beta} \in \mathbb{R}_+^M, \theta \in \mathbb{R} \text{ with } \sum_k \beta_k = 1$$

$$\text{s.t. } \sum_{k=1}^M \beta_k D_k^r \geq \theta \text{ for } r = 1, \dots, t$$

if $|1 - \frac{\theta^{t+1}}{D^t}| \leq \epsilon$ **then break**

end for

SOLVING THE SILP: BOOSTING I

Boosting

Primal

$$\begin{aligned} & \max_{\alpha, \rho} \quad \rho \\ & \text{subject to} \quad y_n \sum_{j=1}^J \alpha_j h_j(\mathbf{s}_n) \geq \rho \\ & \quad \alpha_j \geq 0 \\ & \quad \sum_j \alpha_j = 1 \end{aligned}$$

Dual

$$\begin{aligned} & \min_{\mathbf{d}, \delta} \quad \delta \\ & \text{subject to} \quad \sum_{n=1}^N d_n y_n h_j(\mathbf{s}_n) \leq \delta \\ & \quad d_n \geq 0 \\ & \quad \sum_n d_n = 1 \end{aligned}$$

- max w.r.t. $-\delta$; d_n corresponds to β_n ; h_j corresponds to expression using a certain α
- Arc-GV and AdaBoost*: ρ and δ are optimized

The number of hypotheses can be very large or infinite!

SOLVING THE SILP: BOOSTING I

An Arc-GV like Algorithm

$$D^0 = 1, \rho^1 = \tau_k^1 = 0, \beta_k^1 = \frac{1}{M} \text{ for } k = 1, \dots, M$$

for $t = 1, 2, \dots$ **do**

obtain SVM's α^k with kernel

$$k^t(\mathbf{s}_i, \mathbf{s}_j) := \sum_{k=1}^M \beta_k^t k_k(\mathbf{s}_i, \mathbf{s}_j)$$

for $k = 1, \dots, M$ **do**

$$D_k^t = \frac{1}{2} \sum_{r,s} \alpha_r^t \alpha_s^t y_r y_s k_k(\mathbf{s}_r, \mathbf{s}_s) - \sum_r \alpha_r^t$$

end for

$$D^t = \sum_{k=1}^M \beta_k^t D_k^t$$

$$\gamma_t = \operatorname{argmin}_{\gamma \in [0,1]} \sum_{k=1}^M \beta_k^t \exp \left\{ \gamma (D_k - \rho^t) \right\}$$

for $k = 1, \dots, M$ **do**

$$\tau_k^{t+1} = \tau_k^t + \gamma_t D_k^t$$

$$\beta_k^{t+1} = \beta_k^t \exp(\gamma_t D_k^t) / \left(\sum_{k'} \beta_{k'}^t \exp(\gamma_t D_{k'}^t) \right)$$

end for

$$\rho^{t+1} = \max_k \tau_k^{t+1} / \sum_{r=1}^t \gamma_r$$

if $\left| 1 - \frac{\rho^{t+1}}{D^t} \right| \leq \epsilon$ **then break**

end for