

Lambda Pruning: An Approximation of the String Subsequence Kernel

Alexander K. Seewald¹ and Florian Kleedorfer²

¹ GE Money Bank, CRM/Analytics
Donau-City Straße 6, A-1220 Vienna, Austria
`alex@seewald.at`

² Research Studios Austria, Smart Agent Technologies
Hasnerstraße 123, A-1160 Vienna, Austria
`florian.kleedorfer@{researchstudio.at,austria.fm}`

Abstract. The Support Vector Machine is a powerful learning algorithm which is known to work well for a variety of learning tasks. Kernels for complex data structures such as strings, lists, trees and general graphs have been developed. However, some of these kernels are of high computational complexity and are therefore not widely used. Even relatively efficient kernels may take a lot of patience to run on real-life data. We advance the state of the art for a specific kernel defined on a complex data structure, the Subsequence String Kernel (SSK) due to [9], in two ways: Firstly, by introducing an approximation technique for the SSK called Lambda Pruning (SSK-LP), which is able to reduce memory consumption and runtime by up to several orders of magnitude with little loss in accuracy. Secondly, by creating an average case time and worst case space complexity model for both SSK and SSK-LP. Estimations of runtime and memory requirements for a specific learning task can be directly computed from average string length and desired parameter settings. This combined approach has allowed us to run learning tasks which would be infeasible with the non-approximated SSK. Memory consumption for SSK-LP is constant and does not depend on string length. Models based on SSK and SSK-LP perform similarly for a set of real-life learning tasks.

1 Introduction

Historically, the perceptron classifier was one of the first linear learning machines, and although the famous paper [10], which showed that it cannot learn non-linear discriminant models (such as the XOR problem), had an adverse effect on its development, the perceptron classifier can be seen as ancestor of both Neural networks and Support Vector machines.

In both cases, a solution to solve non-linear problems is introduced: for Neural network, by combining layers of perceptrons with the back-propagation learning rule; for SVMs, by regularizing the linear problem so that it has a unique solution

(i.e. the maximum margin hyperplane) and by blowing up the dimensionality of the original space to improve linear separability.

Support Vector machines improve upon the perceptron by not only having one well-defined global optimum in concept space, but also proven convergence towards this optimum. Neural networks have many local optima, and as for the perceptron convergence is not assured.

Additionally, the kernel trick, of which we will hear more later, allows to uncouple the training method from the data, and thus allows the SVM to process not only numeric data, or data which has been appropriately transformed into numeric format, but also character sequences, trees, lists, graphs and other complex data structures directly (see [6] and [2] for two general approaches to build such kernels). This is a unique capability which few other classifiers share³. However, efficiency is an issue with non-standard kernels and experiments can take a long time and even more memory.

We tackled these issues from two directions, beginning with one well-developed kernel, the Subsequence String Kernel (SSK, [9]).

- By introducing an approximation to the SSK, SSK-LP, which takes far less memory than SSK and is usually several orders of magnitude faster.
- By creating average runtime and maximum memory consumption models for both variants. These models can be applied prior to experimentation, and compute expected runtime and memory consumption for a single kernel evaluation. As kernel evaluation runtime is one major determinant of total runtime of the SVM, especially for complex kernels, this gives a rough estimate of total runtime and the minimum amount of memory needed to successfully run the kernel on a given dataset.

It is our hope that both approaches will prove sufficient to apply the SSK more widely, and that the underlying approach will inspire others as well.

2 Background

This section aims to introduce the reader to the general concept of Support Vector Machines, the *kernel trick* and more specifically to the String Subsequence Kernel (SSK). After this, we explain our approximation method, Lambda Pruning (SSK-LP), in detail.

2.1 Kernels and Support Vector Machines

A Support Vector Machine (SVM) is a classification algorithm which learns a hyperplane that separates two sets of points that belong to different classes. The algorithm is based on viewing the linear separation of these sets as an optimization problem, the goal of which is to maximize the margin of the hyperplane,

³ E.g. Instance-based learning algorithms. Even in that case an appropriate distance measure needs to be defined, and kernels can be viewed as distance measures – see for example Section 4.3.

i.e. the distance from the hyperplane to the nearest point of each class, which is a measure of the robustness of the separation. In case of non-separability, the sum of the distances of all points which are on the wrong side of the hyperplane is bounded by a complexity parameter, usually called C .⁴

Several methods exist that solve this problem efficiently, for instance *Sequential Minimal Optimization* (SMO, [12]) or *SVM-light* [7]. A nonlinear mapping from the attribute space to a higher-dimensional feature space is usually applied prior to training, since a higher number of dimensions increases the probability that the data becomes linearly separable within the high-dimensional feature space. This space may be very high-dimensional, which would normally incur massive computational issues, as runtime and memory consumption would have to increase at least linearly with the number of dimensions used.

This is the point where one central idea of SVMs comes into play, the so-called *kernel trick*: A closer look into the mathematical formulation of SVMs shows that the data points (vectors) are never used directly, they are only introduced into the calculations via the dot-product between two such points (see e.g. [3]). This fact gives rise to the idea of a function combining the mapping and the dot product, so that it is not necessary to compute the high-dimensional vectors explicitly. This kind of function is called a *kernel function*, or short *kernel*.

More formally, for any mapping $\phi : D \rightarrow F$ the function $K : K(x_1, x_2) = \langle \phi(x_1), \phi(x_2) \rangle$ is a kernel function. ($\langle \cdot, \cdot \rangle$ denotes the dot product). An interesting aspect is that the attribute space D from which the kernel projects data into F need not necessarily be an euclidian space, but may have any form and dimension, even infinite dimensionality.

For the Support Vector Machine algorithm to provably work it is just needed to show that the kernel function satisfies Mercer's Theorem, although simpler definitions exist for the finite-dimensional case. It is also sufficient to show that the kernel function can be written as above, i.e. that a function ϕ exists such that K can be written as a dot product of $\phi(x_1)$ and $\phi(x_2)$ (see above). Some kernels have been introduced that are not known to satisfy Mercer's theorem, but which still work well in practice. SSK-LP falls into this latter category, as do some of the kernels used within BioInformatics.⁵

2.2 String Subsequence Kernel (SSK)

The basic idea behind this type of kernel is to define the dot product of two sequences by means of the subsequences they contain. The high-dimensional feature space implied by this kernel contains all possible subsequences as features.

The method proposed in [9] considers contiguous as well as non-contiguous subsequences; the degree of contiguity determines the weight of the subsequence match in the comparison. SSK is parametrized by two values, n and λ . The length of the common subsequences to look for is n . λ is a real value $\in (0, 1)$,

⁴ The original formulation of SVMs calls this complexity parameter λ . Here, λ is a decay factor that penalizes non-contiguous subsequence matches.

⁵ E.g. kernels based on edit distance, see [2].

which is used as a decay factor to penalize non-contiguous substring matches. The kernel implicitly maps the input strings to a feature space F that has one dimension for each possible combination u of n characters.⁶ The u -coordinate $\phi_u(s)$ of a given string s is calculated by summing over all occurrences of u in s . Each occurrence of u yields a value of λ^l , where l denotes the length of that occurrence of u in s , that is, the length of u plus all the interior gaps of the occurrence. The result of the kernel function for two input strings s, t is the dot product of their feature mappings. More formally:

$$\begin{aligned} \phi_u(s) &= \sum_{\mathbf{i}: u=s[\mathbf{i}]} \lambda^{l(\mathbf{i})} \\ K_n(s, t) &= \sum_{u \in \Sigma^n} \langle \phi_u(s) \cdot \phi_u(t) \rangle = \sum_{u \in \Sigma^n} \sum_{\mathbf{i}: u=s[\mathbf{i}]} \lambda^{l(\mathbf{i})} \sum_{\mathbf{j}: u=t[\mathbf{j}]} \lambda^{l(\mathbf{j})} \\ &= \sum_{u \in \Sigma^n} \sum_{\mathbf{i}: u=s[\mathbf{i}]} \sum_{\mathbf{j}: u=t[\mathbf{j}]} \lambda^{l(\mathbf{i})+l(\mathbf{j})} \end{aligned} \tag{1}$$

At first glance this may seem like an exponential algorithm, but [9] also propose an efficient recursive formulation with dynamic programming that uses only $O(n|s||t|)$ time. However, an essential feature to obtain the reported runtime performance is using a cache for all intermediate results, which uses $O(n|s||t|)$ space.

2.3 Lambda Pruning

Equation 1 shows that the result of the kernel evaluation is a sum over different powers of λ . It is obvious that the contribution of subsequences u to the overall result is the smaller the more the subsequence match is stretched in both strings. When analyzing the recursion tree that is processed for each kernel evaluation, it becomes clear, on the other hand, that these stretched matches necessitate a vast amount of computational effort. These observations together motivate a change to the algorithm in a way that the recursion is stopped as soon as it is certain that the result of the current branch is very small. This is a trade-off between result accuracy and consumed computation time.

This behaviour can be achieved by introducing a bound to the parameters in the formula that account for the addends in the kernel computation $\lambda^{l(\mathbf{i})+l(\mathbf{j})}$. By bounding the sum of $l(\mathbf{i}) + l(\mathbf{j})$, we can stop the computation once it reaches the bound, which we have called *Maximum Lambda Exponent*, or θ .

The new kernel parameter θ is used as the upper bound for exponents of λ that may occur in the computation. This means that a subsequence match for a string u is only counted in the final result if the sum of the length of u in s and the length of u in t (both lengths including the gaps) is smaller than θ .

⁶ I.e. the kernel implicitly applies the feature mapping $\phi : D \rightarrow F$.

$$K_{n,\theta}(s,t) = \sum_{u \in \Sigma^n} \sum_{\mathbf{i}: u=s[\mathbf{i}]} \sum_{\mathbf{j}: u=t[\mathbf{j}]} f(l(\mathbf{i}), l(\mathbf{j}))$$

$$f(x,y) = \begin{cases} 0 & \text{if } x+y > \theta \\ \lambda^{x+y} & \text{otherwise} \end{cases}$$

$\theta \dots$ maximum lambda exponent

The influence that θ has on the result is, qualitatively speaking, control of subsequence match relaxation. The definition set of θ is $x \in \mathbb{N}, 2n \leq x < \infty$. $\bar{\theta} = \theta - 2n$ specifies a slack value for match relaxation defining how many interior gaps each subsequence match may have in s and t together.

Using $\theta = 2n$ ($\bar{\theta} = 0$) makes SSK-LP identical to an n-grams kernel, as only contiguous substring matches are allowed. Values slightly above $2n$ ($\bar{\theta} > 0$) create a 'fuzzy' n-grams kernel, while using $\theta \in [4n, 8n]$ ($\bar{\theta} \in [2n, 6n]$) usually gives quite a good approximation of SSK, depending on the type of data used. We have chosen to set a default value of $\theta = 3n$ ($\bar{\theta} = n$), which seems a good compromise for a variety of learning tasks.

The recursive computation⁷ of SSK has to be adapted in order to conform to this definition of lambda pruning. An additional parameter m is added to the kernel function. Each time a value of λ^x is multiplied with the recursion result m is decremented by x and passed to the next recursion level; if the condition $m < 2i$ is true, the recursion ends prematurely as this means that i characters would need to be matched in each of string s and t , incurring a total weight of λ^{2i} . Thus the recursion depth is limited by m .

Definition 1. Recursive computation of SSK with Lambda Pruning (SSK-LP)

$$K'_{0,m}(s,t) = 1, \quad \text{for all } s,t$$

$$K'_{i,m}(s,t) = 0, \quad \text{if } \min(|s|, |t|) < i$$

$$K_{i,m}(s,t) = 0, \quad \text{if } \min(|s|, |t|) < i$$

$$K'_{i,m}(s,t) = 0, \quad \text{if } m < 2i$$

$$K'_{i,m}(sx,t) = \lambda K'_{i,m-1}(s,t) + K''_{i,m}(sx,t), \quad i = 1, \dots, n-1$$

$$K''_{i,m}(sx,tu) = \lambda^{|u|} K''_{i,m-|u|}(sx,t), \quad \exists k : u_k = x$$

$$K''_{i,m}(sx,tx) = \lambda(K''_{i,m-1}(sx,t) + \lambda K'_{i-1,m-2}(s,t))$$

$$K_{n,m}(sx,t) = K_{n,m}(s,t) + \sum_{j:t_j=x} K'_{n-1,m-2}(s,t[1:j-1])\lambda^2$$

In the case of Lambda pruning, we did not use a cache of intermediate results. Although this would have improved runtime slightly, it would also have drastically increased memory consumption beyond that needed for the original SSK.⁸

⁷ see *Efficient Computation of SSK* in [9].

⁸ As there are usually many ways to enter the same recursion tree, but at different entry points, the cached entry could no longer be properly depth-limited as there is

3 Modeling Time, Space and Error

This section describes the average case runtime model, the worst case memory consumption model and an investigation into the SSK-LP approximation error.

3.1 Experimental setup

There are several parameters of the SSK which directly influence runtime and memory consumption. Most prominent among these is the string length $strLen$ of the input strings (assuming $strLen \approx |s| \approx |t|$) and the common subsequence length n . λ has no effect on runtime and memory consumption, but features prominently in the approximation error versus SSK-LP and was therefore also included. For SSK-LP, we also have θ as an essential parameter for the trade-off between runtime and approximation accuracy.

Five values for n were considered (1,2,3,4,5), as well as for λ (0.1, 0.25, 0.5, 0.75, 0.9). θ defaults to $3 * n$ which yields a good trade-off between speed and accuracy, but we used $6 * n$, $9 * n$ and $12 * n$ as well.

Rather than doing a theoretical worst-case analysis on runtime, we opted for a systematic experiment with real-life input data in an average-case runtime analysis, which will later be compared to spot measurements on input data from several different sources. Random strings are somewhat of a best case for SSK, while the worst case would be when both strings consist only of the same character, and both are unlikely to appear in practice.

As real-life corpus, we used an earlier text-mining dataset from [14], consisting of several tens of thousand samples of MEDLINE abstracts, as string corpus. Each MEDLINE entry corresponds to a single string in this corpus. From this corpus, we sampled about one hundred strings of specified lengths, namely $strLen = 1500, 750, 375, \text{ and } 188$. We first sampled strings with exactly the desired length, then those one character longer than the desired length, then those one character shorter than the desired length and so on until enough samples were collected, so the lengths of the collected strings are as close to the target as possible.⁹ Each set of string samples was then paired with a randomly shuffled version of the same set. Thus we obtained pairs of strings with approximately the desired length.

We tested all combinations of these parameter values in extensive experiments on a single machine¹⁰, and measured time in two different ways to make sure that the overall CPU load was accounted for, as the experiment did not use a dedicated machine. The reported execution times are relative to this platform,

only one value per subtree, unless we add the current depth to the cache index. This however would increase memory consumption for the full cache by a factor of $n + 1$.

⁹ The actual string lengths were 1499.60 ± 0.49 , 749.18 ± 3.13 , 373.91 ± 7.34 , and 187.58 ± 3.09 , respectively.

¹⁰ Athlon64 4000+ with 4GB of main memory, which was sufficient to keep even the largest working set in memory, running the pure64 version of Debian, kernel 2.6.11., with Sun Java 1.5. (64bit version) and a CVS version of WEKA from May 2005 with SMO.java V1.12.




Fig. 1. Runtime for SSK, averaged over λ .

and therefore will need to be calibrated for other platforms, but should give a rough overview of the speed of string kernel evaluation.

3.2 Average Case Runtime

Figure 1 shows the runtime for SSK, while Figure 2 shows the runtime for SSK-LP and the four different values for θ . All runtimes are averaged over different values of λ to reduce variance.¹¹ In each case we fitted double ridge-regression models which are shown below, and as lines in each graph. Each symbol corresponds to a specific average runtime measurement, averaged both over the approximately 100 string sample pairs as well as over all different values of λ .

For SSK, we obtained

$$runTime_{SSK} = (strLen^2 * n)^{0.726} * e^{-5.3165}$$

(for $n = 1$)¹² and

$$runTime_{SSK} = (strLen^2 * n)^{1.0402} * e^{-8.897}$$

¹¹ Average relative standard deviation is 0.44% for SSK and 1.03% for SSK-LP, the latter for execution times of more than 100ms. Values below this had high variance as we approach the accuracy limits of runtime measurement (1ms per sample).

¹² $e \approx 2.718281828 = \text{Euler's number}$, since we used natural logarithms.

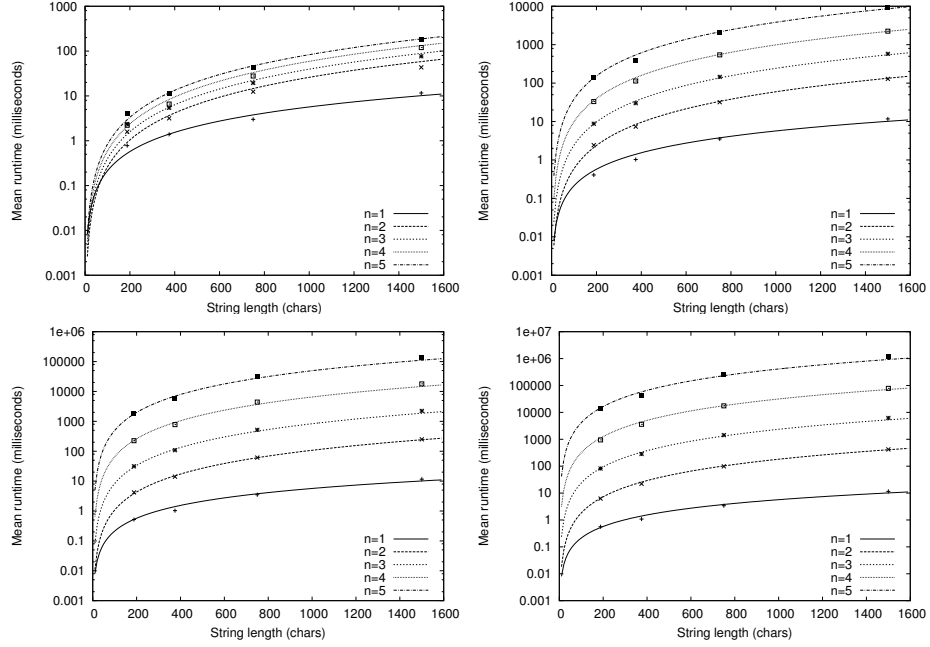


Fig. 2. Runtime for SSK-LP, averaged over λ . Top left $\theta = 3 * n$, Top right $\theta = 6 * n$. Bottom left: $\theta = 9 * n$, Bottom right: $\theta = 12 * n$.

(for $n > 1$). The output is the estimated time in milliseconds per kernel evaluation. Some small optimizations of the first inner loop are responsible for the better runtime of $n = 1$, and for $n > 1$ we see that the average case runtime is very close to the worst case runtime of $O(strLen^2 * n)$ which was given in [9].

For SSK-LP, we similarly obtained

$$runTime_{SSK_LP} = strLen^{1.425} * e^{-8.1063}$$

(for $n = 1$) and

$$runTime_{SSK_LP} = \frac{\theta^{1.0549}}{n!^{1.8839} * (\theta - n)!^{1.039} * \theta^{1.3331}} * strLen^{1.9978} * e^{-10.4875}$$

(for $n > 1$). The left part is similar to a binominal coefficient of θ and n (except for the additional factors of $\frac{1}{n^{1.8839}}$ and $\frac{1}{\theta^{1.3331}}$), which we have reason to believe to feature prominently in the runtime of SSK-LP due to an earlier analysis. The correlation coefficient of this model is 0.9986, relative mean squared error is 0.1662 (5.2037%), which reinforces our belief that this model, albeit being rather more complex than the one for SSK, is valid. Also, figure 2 shows graphically that this function fits the observations rather well.




Fig. 3. Memory consumption for SSK, averaged over λ .

3.3 Worst Case Memory Consumption

Figure 3 shows the memory consumption for SSK. Memory consumption was averaged over all values of λ , which reduces variance.¹³ For each set of string pair samples, only the highest measured memory consumption was used. A regression model has been fitted to the data. Memory consumption is dominated by the cache for intermediate results, and can be computed as $Mem_{SSK} = 38.5888 * (strLen + 1)^2 * (n + 1) + 1,655,638.1147$ (in bytes, round up final result).

The worst case memory consumption of SSK-LP, Mem_{SSK-LP} , is 4,259,840 bytes, independent of all tested parameter settings.

3.4 Approximation Error of SSK-LP vs. SSK

Figure 4 shows the approximation error of SSK-LP vs. SSK in dependence on θ (X axis), n (Y axis) and $lambda$ (lines within each subgraph). As can be seen, for λ of 0.1 and 0.25, agreement is very good over all parameter settings. In some cases agreement is also good for $\lambda = 0.5$. For λ of 0.75 and 0.9, the approximation is much worse, since in this case non-contiguous subsequence matches have a much higher effect on the final value. Figure 5 shows the runtime of SSK vs.

¹³ As expected, values were practically identical for all five different λ values, given any set of other parameter values.

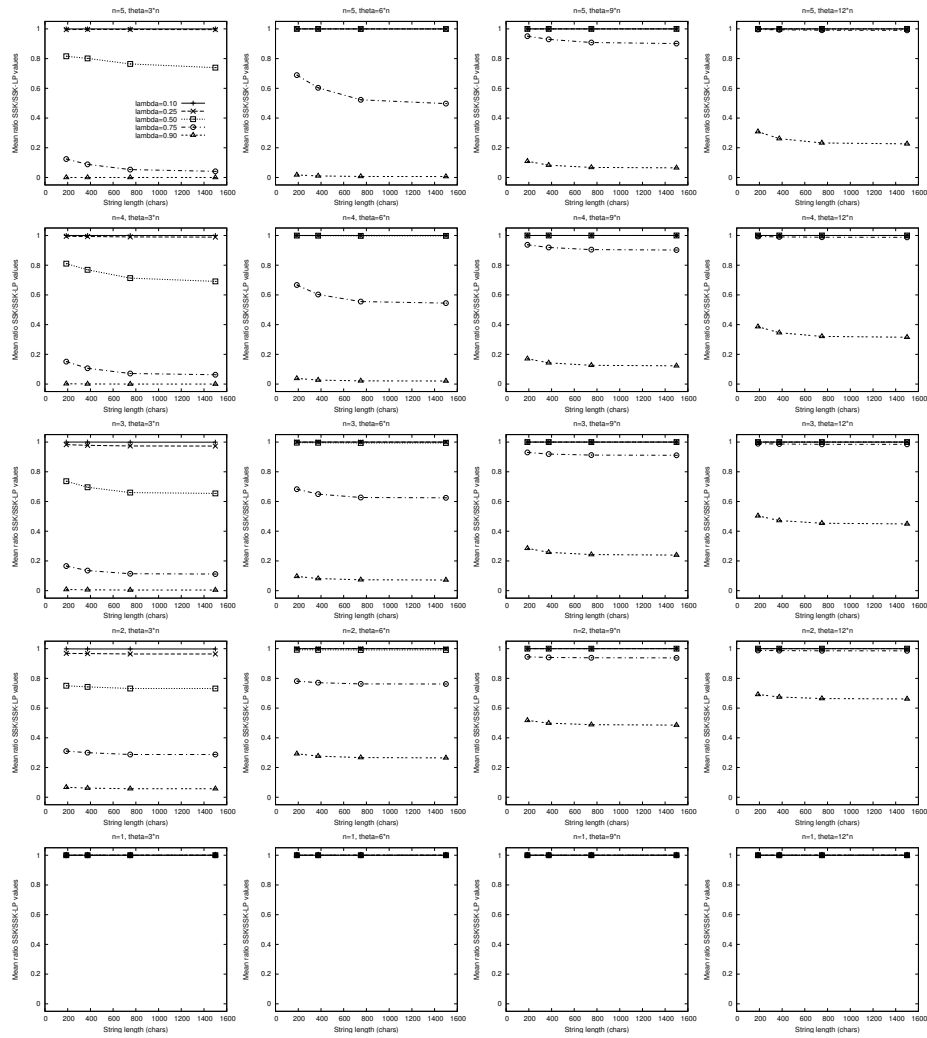


Fig. 4. Ratio of SSK-LP values to SSK values (1.0 = perfect agreement).

SSK-LP in the same layout as Figure 4. It shows that SSK-LP is faster in the lower triangle (for slightly less than half of the parameter settings tested). A more precise estimate of the applicability of SSK-LP to a given problem can always be computed directly, via the models of memory consumption and runtime we presented earlier.




Fig. 5. Runtime of SSK-LP and SSK (in milliseconds) according to the runtime models.

4 Real-Life Performance

In this section we report on experiments with real-life data. We have focussed on three different problems: domain recognition from biological research papers, email spam recognition from sender email address, and the use of SSK as a similarity measure for redundancy clustering – a linguistic task related to sentence entailment.

We used the default settings of $n = 3$ and $\lambda = 0.5$ for SSK, and $\theta = 9$ for SSK-LP, unless otherwise noted. The SVM cost parameter was set to $C = 1$. Table 1 gives a short overview of all results.

Sec.	SSK				SSK-LP			
	est.RT/eval	est.RT total	runtime	result	est.RT/eval	est.RT total	runtime	result
4.1.	1731ms	399.5d	n/a	n/a	89ms	20.5d	21.6d	97.5%
4.2.	0.36ms	2h 17m	16h 50m	84.32%	0.03ms	11.4m	3h 19m	83.98%
4.3.	17.76ms	17.7m	13.95m	see text	1.09ms	1.09m	1.03m	see text

Table 1. This table shows all real-life performance results at one glance. All variants are run with default settings $n = 3, \lambda = 0.5$, and $\theta = 3 * n = 9$ (except 4.2. which used feature space normalization)

4.1 Text mining

For this experiment, we wanted to check the suitability of SSK for standard text mining. We chose the domain dataset from [14]. The learning task was text classification, i.e. classifying bibliographic entries from MEDLINE as belonging to one of four biological domains: Archaea, Bacteria, Eukaryota or Virus. We used the 5% sample and a two-fold CV for accuracy estimation. The results are thus comparable to those reported in [14], Table 1, Acc.CV. The dataset contains 5156 examples. Average length of the input strings is 1497.21 ± 486.26 .

According to our runtime model, SSK is expected to take around 1731ms for one kernel evaluation. To generate the full kernel matrix for a dataset with n examples, at least $\frac{n(n-1)}{2}$ kernel evaluations are needed. Additionally, a two-fold CV must compute a half-sized kernel matrix twice, so the total number of kernel evaluations is about $\frac{3n^2}{4}$. Running SSK on this dataset once would have taken more than a year, and was therefore not considered.

SSK-LP, on the other hand, should take around 89ms for one kernel evaluation, which reduces the expected runtime to 20.5 days. We ran SSK-LP on this data, and the actual runtime was 518h (21.6d), which agrees well with our estimate.

The performance SSK-LP at 97.5% accuracy is competitive to a linear SVM with the word vector as input (i.e. one attribute for each word that appears in training data) at 97.7%. Concerning runtime, it is not competitive: the linear SVM takes around two minutes for the same task. This agrees well with [7], who argued that for most text classification tasks the relatively simple word vector representation is already sufficient.

4.2 Spam filtering

For the second task, we chose a string classification task. Rather than a classic text mining task, where it is feasible to segment the string into smaller constituents by tokenization, this is no longer easily possible for string classification. [16] describes extensive experiments on a large corpus consisting of about 90,000 ham and spam mails. Based on the anecdotal observation that it is often possible to recognize spam mails via a small number of string features (namely, sender address, sender name and the subject), we were interested to find out whether SSK could learn to do the same.

We randomly chose a small number of 3,902 samples of recipient email addresses from this corpus (roughly half from spam and half from ham mails). Sender email addresses were chosen since these are the hardest to tokenize; sender name and subject would have been amenable to a simple tokenization approach and was thus considered unsuitable. The string length in this case is only 25.56 ± 11.08 , so both SSK and SSK-LP can be run on this data. Our runtime models give 0.36ms for SSK, and 0.03ms for SSK-LP, which is below the applicability of the runtime model – the shortest runtime that could be measured was 1ms. This explains why the agreement in that case is quite bad. SSK took 16h 50min to run (est. 2h 17m), while SSK-LP took 3h 19min (est. 11.4m). The speedup factor of about an order of magnitude for SSK-LP over SSK can still be approximately deduced from the differences in estimated runtimes. We used feature space normalization¹⁴ here, which approximately doubles runtime for kernel evaluation.

The accuracy was again estimated via two-fold cross-validation. The accuracies of SSK at 84.32% and SSK-LP at 83.98% were very similar. Systematic parameter optimization of c , n and λ for SSK, including switching off feature space normalization had only negligible influence on the result. The default parameter settings already give the second highest accuracy of 83.39%. SSK-LP with $\theta = 2 * n = 6$ (equivalent to 3-grams) using feature space normalization offers an accuracy of 82.4% and took 57min to run.

4.3 Redundancy clustering

Here, we describe previously unpublished work on redundancy clustering within the BioMinT project (for context see e.g. [11],[14],[15]). The main focus of BioMinT was the extraction of new knowledge from biological research papers via innovative text mining approaches. One subtask was to automatically determine and remove redundant sentences from a result set to give a succinct presentation of the found results to the user. On the one hand, this is somewhat related to summarization approaches; on the other hand, to the known linguistic concept of sentence entailment. Entailment is considered a very hard task, and there are few approaches which perform better than a simple baseline approach.

Within BioMinT, a corpus of redundant sentences was provided by one of our partners. We received this information as a set of sentence groups. Within each sentence group there was a primary sentence that provided all of the information, associated with a set of secondary sentences that provided only information that was already present in the primary sentence. All secondary sentences were considered redundant. To make the task more feasible, we mapped this dataset to pairwise similarity, without distinguishing between primary sentence and secondary sentences. The similarity between two sentences was determined via several methods, and all pairs of sentences with similarity values above a certain threshold were considered redundant while all other pairs were considered to be not redundant. Thus we obtained 59,810 sentence pairs, of which 2,500 were

¹⁴ Using kernel $K'(x, y) = \frac{K(x, y)}{K(x, x)K(y, y)}$, which ensures $K'(x, x)$ is equal to 1.




Fig. 6. This figure shows a ROC curve of the two similarity measures (see text). The X axis corresponds to the proportion of non-redundant sentences which were correctly classified, and the Y axis corresponds to the proportion of redundant sentences which were correctly classified by the system. The unbroken line from (0,1) to (1,0) corresponds to the average performance of a random similarity measure.

marked as redundant. Instead of fixing the threshold arbitrarily and reporting a single result, we have chosen to visualize the results as a ROC curve, where the performance at arbitrary thresholds is observable at one glance. The string length for this learning task was 165.68 ± 61.27 , giving single evaluation times of 17.76ms for SSK and 1.09ms for SSK-LP.

We have selected four ways to determine similarity between arbitrary sentences (S1,S2):

- *Simple*, the number of common words between S1 and S2, divided by the length (in words) of the shorter sentence. This is a simple baseline approach.
- *SSK-def*, which determines the similarity between S1 and S2 by computing an evaluation of the unnormalized string kernel (i.e. SSK), which is normally used as a kernel for the learning algorithm family of Support Vector Machines. We chose the default settings of $\lambda = 0.5$ and $n = 3$ here.
- *SSK-opt* is the same as *SSK-def* but with $n = 6$ which was found to be the optimal setting for n here.
- *SSK-LP*, which determines the similarity between S1 and S2 via SSK-LP. In this case, we used $\lambda = 0.5$, $n = 6$ and $\theta = 2 * n = 12$, so this is equivalent to a simpler 6-grams approach.

Figure 6 shows the results. The unbroken line visualizes the average performance of a random similarity measure (i.e. where the values of $\text{sim}(S1, S2)$ are chosen randomly). The higher above this line a measure is, the better it performs. As can be seen, SSK_{def} already performs much better than the simpler baseline approach. Parameter optimization for n improves on this result. While the equivalent 6-grams approach ($SSK-LP$) performs competitively, we did not optimize λ here which might have improved performance for SSK beyond $SSK-opt$.

5 Related Work

[8] propose a different kind of string kernel, the so-called spectrum or mismatch kernel. The restriction to n -grams (i.e., contiguous subsequences) as features allows a more efficient implementation. They do not show that their kernel satisfies Mercer’s theorem, but [2] demonstrate this. As SSK-LP with $\theta = 2 * n$ is equivalent to using n -gram features, this specific variant must therefore satisfy Mercer’s theorem as well.

[1] introduce convolution kernels for natural language processing, i.e. a tree kernel based on the ideas of [6]. They quote (unproved) an average runtime that is linear in the tree size, but worst case performance is on the order of square of tree size. This kernel is shown to satisfy Mercer’s theorem. Features are all common subtrees, so again no gaps are considered.

[9] is another work that can be viewed as based on [6], and gives the original implementation of SSK as well as some experimental results. This work has already been amply discussed throughout this paper.

A recent work is [2] on Rational Kernels, which unifies several approaches to build kernels on complex data structures, and also shows that edit distance with a alphabet of size greater than 1 does not satisfy Mercer’s theorem. They show that a single algorithm is sufficient to implement all rational kernels, albeit via a finite automaton approach – a programming language by any other name – and thus cannot claim to simplify the implementation of new kernels.

[13] offer a comprehensive evaluation of several approaches to compute string kernels, namely trie-based and full dynamic programming (similar to SSK), and propose a sparse dynamic programming approach. While runtime results are given, no analytical model is obtained. Additionally, their proposed new method depends on alphabet size and the number of gaps that are permitted as well as string length, so the results are not directly comparable to our work. Their runtime complexity is $O(n|M|\log|s|)$, where $|M|$ is the size of the alphabet. However, while the worst-case runtime will in some cases be better than that for SSK and SSK-LP, the obtained actual runtime might still be too slow for reasonably sized real-life datasets. As case in point, they have focussed only on single kernel evaluations rather than running a SVM based on these kernels.

6 Conclusion

The work at hand presents *Lambda Pruning*, a novel approach to approximating the String Subsequence Kernel (SSK) by Lodhi et al.[9]. The resulting kernel, SSK-LP for short, is implemented in Java and available as a part of the Weka data mining platform¹⁵. Our runtime and space complexity models allow to compute the expected runtime and memory consumption for both variants, and would even allow an automatic choice between the variants based on input data and given parameter settings.

We have investigated both variants on several learning tasks and noted that the string kernels are most useful for string classification tasks, where a tokenization is not easily apparent, and as similarity measure for redundancy clustering. The given runtime models usually agree well with actual runtime and have proven useful to determine the relative speedup of SSK-LP versus SSK as well as to exclude some parameter settings because of excessive runtime. Our work is therefore an important step towards the practical applicability of string kernels for real-life learning tasks, and this approximation approach may prove useful for other complex kernels as well.

We should note that in our real-life experiments, n-gram approaches were found to be competitive, and their implementations can be made more efficient. More work is needed to improve SSK to take adequate advantage of the capability to consider non-contiguous subsequences as well as improving efficiency further.

Being the first algorithm for learning from sequences in Weka, this opens a new perspective for its users; moreover, the reduced complexity of SSK-LP makes the Support Vector Machine approach a viable alternative for text mining even on small devices with little main memory.

7 Acknowledgements

This research was partially supported by the European Commission as project no. QLRI-CT-2002-02770 (*BioMinT*) under the RTD programme *Quality of Life and Management of Living Resources*. The Studio Smart Agent Technologies is supported by the Austrian Federal Ministry of Economics and Labour. We gratefully acknowledge the support of the University of Manchester in providing the redundancy dataset.

References

1. M. Collins and N. Duffy. *Convolution kernels for natural language*. In T. G. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems*, volume 14, Cambridge, MA, 2002. MIT Press.

¹⁵ Check out CVS release from www.cs.waikato.ac.nz/~ml/weka. Use `weka.classifiers.functions.SMO` with kernel `weka.classifiers.functions.supportVector.StringKernel`. -P 0 = SSK, -P 1 = SSK-LP.

2. C. Cortes, P. Haffner and M. Mohri: *Rational Kernels: Theory and Algorithms*, Journal of Machine Learning Research 5 (2004) 1035–1062.
3. N. Cristianini, J. Shawe-Taylor: *An Introduction to Support Vector Machines*, Cambridge University Press, 2000.
4. E. Frank, I.H. Witten: *Data Mining - Practical Machine Learning Tools and Techniques (2nd ed)*. Morgan Kaufmann / Elsevier, 2005.
<http://www.cs.waikato.ac.nz/~ml/weka/>
5. T. Gartner, P. A. Flach, and S. Wrobel: *On Graph Kernels: Hardness Results and Efficient Alternatives*, Sixteenth Annual Conference on Computational Learning Theory and The Seventh Kernel Workshop (COLT-2003).
6. D. Haussler. *Convolution kernels on discrete structures*. Technical Report UCSC-CRL -99-10, Baskin School of Engineering, University of California, Santa Cruz, 1999.
7. T. Joachims *Learning to Classify Text Using Support Vector Machines*, Kluwer Academic Publishers, May 2002.
8. C. Leslie, E. Eskin, and W. S. Noble: *The spectrum kernel : A string kernel for SVM protein classification*. In Proceedings of the Pacific Symposium on Biocomputing, pages 564–575, 2002.
9. H. Lodhi, C. Saunders, J. Sahwe-Taylor, N. Christianini, C. Watkins: *Text Classification using String Kernels*, Journal of Machine Learning Research 2 (2002), p.419–444.
10. M. Minsky, S.A. Papert: *Perceptrons: An Introduction to Computational Geometry*. MIT Press, Cambridge, MA, expanded edition, 1988/1969.
11. V. Pillet, M. Zehnder, A.K. Seewald, A-L Veuthey, and J. Petrak. *GPSDB: a new database for synonyms expansion of gene and protein names*. Bioinformatics 2005 21: 1743-1744.
12. J. Platt: *Sequential Minimal Optimization: A Fast Algorithm for Training Support Vector Machines*
13. J. Rousu, and J. Shawe-Taylor. *Efficient Computation of Gapped Substring Kernels on Large Alphabets*, Journal of Machine Learning Research 6 (2005), pp.1323–1344.
14. A.K. Seewald: *Recognizing Domain and Species from MEDLINE Proteomics Publications. Workshop on Data Mining and Text Mining for Bioinformatics, 14th European Conference on Machine Learning (ECML-2003), Dubrovnik-Cavtat, Croatia, 2003*.
15. A.K. Seewald: *Ranking for Medical Annotation: Investigating Performance, Local Search and Homonymy Recognition*. Proceedings of the Symposium on Knowledge Exploration in Life Science Informatics (KELSI 2004), Milano, Italy.
16. A.K. Seewald: *An Evaluation of Naive Bayes Variants in Content-Based Learning for Spam Filtering*. Austrian Research Institute for Artificial Intelligence, Technical report, TR-2005-20, Vienna, 2005.