Large-Scale Support Vector Learning with Structural Kernels

Aliaksei Severyn and Alessandro Moschitti

Department of Computer Science and Engineering
University of Trento
Via Sommarive 14, 38100 POVO (TN) - Italy
`{severyn,moschitti}@disi.unitn.it`

Abstract. In this paper, we present an extensive study of the cutting-plane algorithm (CPA) applied to structural kernels for advanced text classification on large datasets. In particular, we carry out a comprehensive experimentation on two interesting natural language tasks, e.g. predicate argument extraction and question answering. Our results show that (i) CPA applied to train a non-linear model with different tree kernels fully matches the accuracy of the conventional SVM algorithm while being ten times faster; (ii) by using smaller sampling sizes to approximate subgradients in CPA we can trade off accuracy for speed, yet the optimal parameters and kernels found remain optimal for the exact SVM. These results open numerous research perspectives, e.g. in natural language processing, as they show that complex structural kernels can be efficiently used in real-world applications. For example, for the first time, we could carry out extensive tests of several tree kernels on millions of training instances. As a direct benefit, we could experiment with a variant of the partial tree kernel, which we also propose in this paper.

Keywords: Structural Kernels; Support Vector Machines; Natural Language Processing.

1 Introduction

In many computer science areas such as Natural Language Processing (NLP), Bioinformatics, Data Mining and Information Retrieval, structural kernels have been widely used to capture rich syntactic information, e.g. [3,12,29,30]. In particular, in NLP, tree kernel functions can capture a representation of syntactic parse trees, which is considerably more effective than the one provided by straightforward bag-of-words models. This often results in higher accuracy, e.g. [16,18,31,11]. Unfortunately, the use of kernels forces us to solve SVM optimization problem in the dual space, which, in case of the very large datasets, makes SVM learning prohibitively expensive.

Recently, cutting plane approaches have been proposed to achieve a great speed-up in the learning of linear models on large datasets, but they still fail to provide the same training performance on non-linear learning tasks. Indeed, the

number of kernel evaluations scales quadratically with the number of examples. To overcome this bottleneck, Yu and Joachims [28] developed two approximate cutting plane algorithms (CPAs). However, their experiments were carried out using only Gaussian kernels on unstructured data.

In this paper, we study models that combine the speed of CPAs with the efficiency of SVM-light-TK, which encodes state of the art structural kernels [17,18] in SVM-light [6]. More specifically, by the means of extensive experimentation we examine the applicability of CPAs to well-known structural kernels, i.e. subtree (ST), subset tree (SST), and partial tree (PT) kernels, which provide a good sample of the efficient tree kernel technology currently available. To obtain more general results we considered two advanced and very different text categorization tasks for which syntactic information is essential: (i) Semantic Role Labeling (SRL) or predicate argument extraction [20,15] whose associated dataset contains several millions of examples; and (ii) Question Classification (QC), e.g. [13,26] from question answering domain. Some of the distinct properties of the aforementioned datasets (compared to the previous works) are:

- Features are structured and embed syntactic information: from their relationship, the classifier derives higher semantics.
- The number of features is huge (millions) and they tend to be very sparse.
- The space is discrete, the a-priori weights are very skewed, and large fragments receive exponentially lower scores than small fragments.
- There is high redundancy and inter-dependency between features.

Finally, we defined a novel kernel, unlexicalized PTK (uPTK), which is a variant of PTK. It excludes single nodes from the feature space so that structural features are better emphasized.

Our findings reveal that:

- As the theory suggests, CPAs can be successfully applied to structural spaces. Indeed, we show that CPA is at least 10 times as fast as the exact version while giving the same classification accuracy. For example, training a conventional SVM solver with tree kernels on 1 million examples requires more than seven days, while CPAs match the same accuracy in just a few hours.
- By decreasing the sampling size used in the approximation of the cutting planes, we can trade off accuracy to a small degree for even faster training time. For example, learning a model that is only 1.0 percentage point apart from the exact model reduces the training time by a factor of 50.
- Using a sample size of only 100 instances for CPAs, which takes just a couple of minutes of learning on one million of examples, we can correctly estimate the best kernel, its hyper-parameters, and the trade-off parameter. The identified set of optimal parameters can be used by more computationally expensive and accurate models (CPAs with larger sample sizes or SVM-light).
- Thanks to the gained speed-up we could efficiently compare different kernels on the large SRL datasets and establish an absolute rank, where uPTK proves to be more effective than PTK.

Our results open up new perspectives for the application of structural kernels in various NLP tasks. The experiments that could not be carried out until now on full subsets of very large corpora, e.g. training SVM-based SRL classifier such as in [22,15] using either polynomial or tree kernels on the full dataset; training SVM re-rankers for syntactic parsing [3,25] using tree kernels on the available data; and training SVM question classifiers using tree kernels and other features on large subsets of Yahoo! Answers dataset. The aforementioned motivations encouraged us to release the new Tree Kernel toolkit to the public¹.

In the remainder of this paper, Section 2 reviews the related work, while Section 3 gives careful theoretical treatment to the workings of the cutting plane approach. Section 4 introduces well-known tree kernels along with the new uPTK. The experimental analysis (Section 5) describes our experiments on SRL and QC datasets and also discusses how approximate cutting plane algorithms could be used to drive parameter selection for the exact SVM solver. Finally, in Section 6, we draw conclusions and provide directions for the further research.

2 Related Work

Since the introduction of SVMs, a number of fast algorithms that can efficiently train non-linear SVMs have been proposed: for example, decomposition method along with working set selection [6]. Decomposition methods work directly in the dual space and perform well on moderately large datasets but their performance degrades when the number of training examples reaches a level of millions of examples.

Recently, a number of efficient algorithms using cutting planes to train conventional SVM classifiers have been proposed. For example, SVM^{perf} [7] is based on a cutting plane algorithm and exhibits linear computational complexity in the number of examples when linear kernels are used. To improve the convergence rate of the underlying cutting plane algorithm, Franc and Sonnenburg [5] developed the optimized cutting plane algorithm (OCAS) that achieves speed-up factor of 29 over SVM^{perf} . Alternatively, another promising approach based on stochastic gradient descent and projection steps called Pegasos [24] has shown promising performance for linear kernels in binary classification tasks.

While the aforementioned algorithms deliver state of the art performance with respect to accuracy and training time, they scale well only when linear kernels are used. To overcome this bottleneck, an idea to use approximate cutting planes with random sampling was employed by Yu and Joachims [28]. At each iteration step an exact cutting plane is replaced by its approximation that is built from a small subset of examples sampled from the training dataset.

Another approximate technique to speed up the computation time of tree kernels by comparing only a sparse subset of relevant subtrees is discussed in [23].

Among the basis pursuit approaches that find a set of basis vectors to construct sparse kernel SVMs are works by Keerthi [9] and Joachims [8]. The former uses a greedy technique to select vectors, while the latter extracts basis vectors

¹ Available at <http://projects.disi.unitn.it/iKernels>

as a part of the cutting-plane optimization process. Not only the trained sparse model speeds up the classification time, but it also improves the training complexity. Even though the achieved scaling behavior is roughly linear, the method in [8] is limited to the use of only Gaussian kernels.

Following a different line of research, a number of methods that exploit low-rank approximation of a kernel matrix have been proposed in [4,27]. However, the experiments were carried out only on fairly small datasets (thousands of examples).

The approach we choose to employ in this paper is different in a sense that it is generally applicable to the learning of any non-linear discriminant function, structural kernels in particular, and can efficiently handle datasets with hundreds of thousands examples.

3 Cutting Plane Algorithm for Structural Kernels

In this section, we illustrate the cutting plane approach. Previous works focus on structural SVMs whereas the topic of this paper is binary classification, thus we present a re-elaborated version of the algorithm tailored for the binary classification task starting off with the derivation of the dual formulation. This results in an easier discussion of the sampling approach to compute the approximate cutting planes.

3.1 Cutting-Plane Algorithm (Dual)

Below is an equivalent formulation of an SVM optimization problem, known as a 1-slack reformulation [7], used to derive the cutting-plane algorithm:

$$\begin{aligned} & \underset{\mathbf{w}, \xi \geq 0}{\text{minimize}} && \frac{1}{2} \|\mathbf{w}\|^2 + C\xi \\ & \text{subject to} && \forall \mathbf{c} \in \{0, 1\}^n : \\ & && \frac{1}{n} \mathbf{w} \cdot \sum_{i=1}^n c_i y_i \mathbf{x}_i \geq \frac{1}{n} \sum_{i=1}^n c_i - \xi, \end{aligned} \tag{1}$$

where each constraint can be viewed as an average sum of a subset of constraints of the form: $y_i(\mathbf{w} \cdot \mathbf{x}_i) \geq 1 - \xi_i$ that are selected by a vector $\mathbf{c} = (c_1, \dots, c_n) \in \{0, 1\}^n$. While the total number of constraints is 2^n , there is only a single slack variable ξ shared across all the constraints. For more details on the 1-slack formulation, an interested reader should refer to [7].

To derive the dual formulation, the Lagrangian of the primal problem (1) is computed as:

$$L_P = \frac{1}{2} \|\mathbf{w}\|^2 + C\xi - \sum_{j=1}^{|S|} \alpha_j \left(\frac{1}{n} \sum_{k=1}^n c_{kj} (y_k \mathbf{w} \cdot \mathbf{x}_k - 1) + \xi \right), \tag{2}$$

where S denotes the set of constraints in (1) and $|S| = 2^n$, c_{kj} denotes the k^{th} component of a vector $\mathbf{c}_j$ that corresponds to the j^{th} constraint in (1), and $\alpha_j \geq 0$ are the Lagrange multipliers.

Using the fact that both gradients of L_P with respect to w and ξ vanish:

$$\begin{aligned}\frac{\delta L_P}{\delta w} &= w - \sum_{j=1}^{|S|} \alpha_j \left(\frac{1}{n} \sum_{k=1}^n c_{kj} y_k x_k \right) = 0 \\ \frac{\delta L_P}{\delta \xi} &= \frac{1}{n} \sum_{j=1}^{|S|} \alpha_j - C = 0,\end{aligned}\tag{3}$$

and by substituting variables from (3) into (2), we obtain the dual Lagrangian:

$$L_D = \sum_{i=1}^{|S|} \alpha_i h^{(i)} - \frac{1}{2} \sum_{i=1}^{|S|} \sum_{j=1}^{|S|} \alpha_i \alpha_j g^{(i)} \cdot g^{(j)}\tag{4}$$

where $h^{(i)} = \frac{1}{n} \sum_{k=1}^n c_{ki}$ and $g^{(i)} = -\frac{1}{n} \sum_{k=1}^n c_{ki} y_k x_k$ are respectively the bias and the subgradient that define a cutting plane model $(h^{(i)}, g^{(i)})$. Now we can state the dual variant of the optimization problem (1):

$$\begin{aligned}\text{maximize}_{\alpha \geq 0} \quad & h^T \alpha - \frac{1}{2} \alpha^T H \alpha \\ \text{subject to} \quad & \alpha^T \mathbf{1} \leq C,\end{aligned}\tag{5}$$

where $H_{ij} = g^{(i)} \cdot g^{(j)}$. Using the first equation from (3), we get the connection between the primal and dual variables:

$$w = \sum_{j=1}^{|S|} \alpha_j \left(\frac{1}{n} \sum_{k=1}^n c_{kj} y_k x_k \right) = - \sum_{j=1}^{|S|} \alpha_j g^{(j)},\tag{6}$$

Now we are ready to present the CPA method (Algorithm 1). It starts with an empty set of constraints S and computes the optimal solution (line 5) to the unconstrained problem (1). Next, the algorithm finds the most violated constraint (lines 9-11) by forming a binary vector c that defines the maximum sum of violated constraints. This requires the computation of the following quantity for each training example:

$$w \cdot \phi(x_i) = - \sum_{j=1}^{|S|} \alpha_j g^{(j)} \phi(x_i) = \sum_{k=1}^n \left(\sum_{j=1}^{|S|} \frac{1}{n} \alpha_j c_{kj} y_k \right) K(x_i, x_k),\tag{7}$$

where $\phi(x)$ is a mapping from the input to the feature space and $K(x_i, x_k) = \phi(x_i) \cdot \phi(x_k)$ is a kernel. We then compute the bias $h^{(t)}$ and the subgradient $g^{(t)}$ (lines 12 and 13 respectively) to build a new cutting plane which defines a half-space $h^{(t)} + w \cdot g^{(t)} \leq \xi$ corresponding to the constraint that is most violated by the current solution. Then it is included in the set of active constraints S (line 14). This way, a series of successively tightening approximations to the original

Algorithm 1. Cutting Plane Algorithm (dual) with uniform sampling

```

1: Input:  $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n), C, \epsilon$ 
2:  $S \leftarrow \emptyset; t = 0;$ 
3: repeat
4:   Update matrix  $H$  with a new constraint
5:    $\alpha \leftarrow \text{optimize QP problem (5)}$ 
6:    $\xi = \frac{1}{C}(\mathbf{h}^T \alpha - \frac{1}{2} \alpha^T H \alpha)$ 
7:    $\mathbf{w} = -\sum_{j=1}^{|S|} \alpha_j \mathbf{g}^{(j)}$ 
8:   Sample  $r$  examples from the training set
   /* find the most violated constraint (cutting plane) */
9:   for  $i = 1$  to  $r$  do
10:   $c_i \leftarrow \begin{cases} 1 & y_i(\mathbf{w} \cdot \phi(\mathbf{x}_i)) \leq 1 \\ 0 & \text{otherwise} \end{cases}$ 
11: end for
12:  $h^{(t)} = \frac{1}{r} \sum_{i=1}^r c_i$ 
13:  $\mathbf{g}^{(t)} = -\frac{1}{r} \sum_{i=1}^r c_i y_i \phi(\mathbf{x}_i)$ 
   /* add a constraint to the active set */
14:  $S \leftarrow S \cup \{(h^{(t)}, \mathbf{g}^{(t)})\}$ 
15:  $t = t + 1$ 
16: until  $h^{(t)} + \mathbf{w} \cdot \mathbf{g}^{(t)} \leq \xi + \epsilon$ 
17: return  $\mathbf{w}, \xi$ 

```

problem is constructed. The algorithm stops when no constraints are violated by more than ϵ , which is formalized by the criteria in line 16.

The analysis of the inner product given by (7) reveals that, since it needs to be computed for each training example, it requires the time $O(n^2 + Tn)$ after total of T iterations. Similarly, as we add a cutting plane to S at each iteration t , a new column is added to the matrix H (Algorithm 1, line 4) requiring the computation of

$$H_{it} = \mathbf{g}^{(i)} \cdot \mathbf{g}^{(t)} = \frac{1}{n^2} \sum_{k=1}^n \sum_{l=1}^n c_{ki} c_{lt} y_k y_l K(\mathbf{x}_k, \mathbf{x}_l) \quad (8)$$

which takes $O(Tn^2)$. Thus, the obtained $O(n^2)$ scaling behavior makes cutting plane training no better than conventional decomposition methods.

To address this limitation, we employ the approach of Yu and Joachims [28] to construct approximate cuts by sampling r examples from the training set. They suggest two strategies to sample examples, namely uniform and importance sampling (the pseudocode of the algorithm using uniform sampling is presented in Algorithm 1). These two strategies derive constant-time and linear-time algorithms. The former uniformly samples r examples from the training set to approximate the cut. Thus, we approximate a subgradient $\mathbf{g}^{(t)}$ with only r examples, which replaces the number of expensive kernel evaluations in (7) over n by a more tractable: $\sum_{i,j=1}^r K(\mathbf{x}_i, \mathbf{x}_j)$ (lines 9-13). The importance sampling acts in a more targeted way as it looks through the whole dataset to compute two cutting planes, one to be used in the optimization problem (line 5), and the

other for termination criterion (line 16). The training complexity reduces from $O(n^2)$ to $O(T^2r^2)$, when the uniform sampling algorithm is used, to $O(Tnr)$ for the importance sampling.

4 Tree Kernels

The main idea underlying tree kernels is to compute the number of common substructures between two trees T_1 and T_2 without explicitly considering the whole fragment space. Let $\mathcal{F} = \{f_1, f_2, \dots, f_{|\mathcal{F}|}\}$ be the set of tree fragments and $\chi_i(n)$ an indicator function equal to 1 if the target f_i is rooted at node n and equal to 0 otherwise. A tree kernel function over T_1 and T_2 is defined as

$$TK(T_1, T_2) = \sum_{n_1 \in N_{T_1}} \sum_{n_2 \in N_{T_2}} \Delta(n_1, n_2),$$

where N_{T_1} and N_{T_2} are the sets of nodes in T_1 and T_2 , respectively, and

$$\Delta(n_1, n_2) = \sum_{i=1}^{|\mathcal{F}|} \chi_i(n_1) \chi_i(n_2).$$

The Δ function is equal to the number of common fragments rooted in nodes n_1 and n_2 and thus depends on the fragment type.

4.1 Fragment Types

In [17], we pointed out that there are three main categories of fragments: the subtree (ST), the subset tree (SST) and the partial tree (PT) fragments corresponding to three different kernels. STs are fragments rooted in any node of a tree along with all its descendants. The SSTs are more general structures since, given the root node of an SST, not all its descendants (with respect to the referring tree) have to be included, i.e. the SST leaves can be non-terminal symbols. PT fragments are still more general since their nodes can contain a subset of the children of the original trees, i.e. partial sequences.

For example, Figure 1 illustrates the syntactic parse tree of the sentence **Autism is a disease** on the left along with some of the possible fragments on the right of the arrow. ST kernel generates complete structures like [D a] or [NP [D a] [N disease]]. SST kernel can generate more structures, e.g. [NP [D] [N disease]] whereas PT kernel can also separate children in the fragments, e.g. [NP [N disease]], and generate the individual tree nodes as features, e.g. **Autism** or **VBZ**.

[30] provided a version of SST kernel, which also generates leaves, i.e. words, as features, hereafter, SST-bow. However, such lexical features, when the data is very sparse, tend to cause overfitting. Thus, we give the definition of a variant of PTK, namely, the unlexicalized partial tree kernel (uPTK), which does not include lexicals and individual nodes in the feature space. This will promote the importance of structural information.

Fig. 1. A tree for the sentence “Autism is a disease” (top left) with some of its partial tree fragments (PTFs)

4.2 Unlexicalized Partial Tree Kernel (uPTK)

The algorithm for the uPTK computation straightforwardly follows from the definition of the Δ function of PTK provided in [17]. Given two nodes n_1 and n_2 in the corresponding two trees T_1 and T_2 , Δ is evaluated as follows:

1. if the node labels of n_1 and n_2 are different then $\Delta(n_1, n_2) = 0$;
2. else $\Delta(n_1, n_2) = \mu \left(\lambda^2 + \sum_{\mathbf{I}_1, \mathbf{I}_2, l(\mathbf{I}_1)=l(\mathbf{I}_2)} \lambda^{d(\mathbf{I}_1)+d(\mathbf{I}_2)} \prod_{j=1}^{l(\mathbf{I}_1)} \Delta(c_{n_1}(\mathbf{I}_{1j}), c_{n_2}(\mathbf{I}_{2j})) \right)$,

where: (a) $\mathbf{I}_1 = \langle h_1, h_2, h_3, \dots \rangle$ and $\mathbf{I}_2 = \langle k_1, k_2, k_3, \dots \rangle$ are index sequences associated with the ordered child sequences c_{n_1} of n_1 and c_{n_2} of n_2 , respectively; (b) $\mathbf{I}_{1j}$ and $\mathbf{I}_{2j}$ point to the j -th child in the corresponding sequence; (c) $l(\cdot)$ returns the sequence length, i.e. the number of children; (d) $d(\mathbf{I}_1) = \mathbf{I}_{1l(\mathbf{I}_1)} - \mathbf{I}_{11} + 1$ and $d(\mathbf{I}_2) + 1 = \mathbf{I}_{2l(\mathbf{I}_2)} - \mathbf{I}_{21} + 1$; and (e) μ and λ are two decay factors for the size of the tree and for the length of the child subsequences with respect to the original sequence, i.e. we account for gaps.

The uPTK, can be obtained by removing λ^2 from the equation in the step 2. An efficient algorithm for the computation of PTK is given in [17]. This evaluates Δ by summing the contribution of tree structures coming from different types of sequences, e.g. those composed by p children such as:

$$\Delta(n_1, n_2) = \mu \left(\lambda^2 + \sum_{p=1}^{lm} \Delta_p(c_{n_1}, c_{n_2}) \right), \quad (9)$$

where Δ_p evaluates the number of common subtrees rooted in subsequences of exactly p children (of n_1 and n_2) and $lm = \min\{l(c_{n_1}), l(c_{n_2})\}$. It is easy to verify that we can use the recursive computation of Δ_p proposed in [17] by simply removing λ^2 from Eq. 9.

5 Experiments

In these experiments, we study the impact of the cutting plane algorithms (CPAs), reviewed in Section 3, on learning complex text classification tasks in

structural feature spaces. For this purpose, we compare the accuracy and the learning time of CPAs, according to different sample size against the conventional SVMs.

In the second set of experiments, we investigate the possibility of using fast parameter and kernel selection with CPA for conventional SVM. For this purpose, we carried out experiments with different classifiers on two different domains.

5.1 Experimental Setup

We integrated two approximate cutting plane algorithms using sampling [28] with SVM-light-TK [17]. For brevity, in this section we will refer to the algorithm that uses uniform sampling as uSVM, importance sampling as iSVM, and SVM-light-TK as SVM. While the implementation of sampling algorithms uses MOSEK to optimize quadratic problem, SVM is based on SVM-light 5.0 solver. As the stopping criteria of the algorithms, we fix the precision parameter ϵ at 0.001.

We experimented with five different kernels: the ST, SST, SST-bow, PT, uPT kernels described in Section 4, which are also normalized in the related kernel space. All the experiments that do not involve parameter tuning use the default trade-off parameter (i.e. 1 for normalized kernels) and the default λ fixed at 0.4.

As a measure of classification accuracy we use the harmonic average of the Precision and Recall, i.e. F_1 -score. All the experiments were run on machines equipped with Intel[®] Xeon[®] 2.33GHz CPUs carrying 6Gb of RAM under Linux 2.6.18 kernel.

5.2 Data

We used two different natural language datasets corresponding to two different tasks: Semantic Role Labeling (SRL) and Question Answering.

The first consists of the Penn Treebank texts [14], PropBank annotation [20] and automatic Charniak parse trees [2] as provided by the CoNLL 2005 evaluation campaign [1]. In particular, we tackle the task of identification of the argument boundaries (i.e. the exact sequence of words compounding an argument). This corresponds to the classification of parse tree nodes in correct or not correct boundaries². For this purpose, we train a binary *Boundary Classifier* (BC) using the AST subtree defined in [15], i.e. the minimal subtree, extracted from the sentence parse tree, including the predicate and the target argument nodes. To test the learned model, we extract two sections, namely *sec23* and *sec24*, that contain 234,416 and 149,140 examples respectively. The models are trained on two subsets of 100,000 and 1,000,000 examples. The proportion of positive examples in the whole corpus is roughly 5%. The dataset along with the exact structural representation is available at <http://danielepiggin.net/cms/research/MixedFeaturesForSRL>.

² In the automatic trees some boundary may not correspond to any node. In this case, we choose the lower node dominating all the argument words.

Fig. 2. F_1 -score as a function of the sampling size on SRL dataset (1 million examples). Horizontal axis at top is the training time of the uSVM algorithm.

The second corpus, is a Question Classification (QC) dataset, whose testset comes from the TREC 10 - Question Answering evaluation campaign whereas the training set³ was developed in [13]. The task consists in selecting the most appropriate type of the answer from a set of given possibilities. The *coarse grained* question taxonomy [30,13] consists of six non overlapping classes: Abbreviations (ABBR), Descriptions (DESC, e.g. definitions or explanations), Entity (ENTY, e.g. animal, body or color), Human (HUM, e.g. group or individual), Location (LOC, e.g. cities or countries) and Numeric (NUM, e.g. amounts or dates). For each question, we used the full parse tree as its representation (similarly to [30,17,19]). This is automatically extracted by means of the Stanford parser⁴ [10]. We actually have only 5,483 questions in our training set, due to parsing issues with a few of them. The testset is constituted by 500 questions and the size of the categories varies from one thousands to few hundreds.

5.3 Accuracy and Efficiency vs. Sampling Size

In these experiments, we test the trade-off between speed and accuracy of CPAs. We use uSVM, since it is faster than iSVM, and compare it against SVM on the SRL task by training on 1 million examples and testing on the two usual testing sections, i.e. *sec23* and *sec24*. We used the SST kernel since it has been indicated as the most accurate in similar tasks, e.g. [17]. Figure 2 plots the F_1 -score for different values of the sampling size. The dashed horizontal lines denote the accuracy achieved by the exact SVM. The training time for the uSVM algorithm is plotted along the top horizontal axis.

³ <http://l2r.cs.uiuc.edu/cogcomp/Data/QA/QC/>

⁴ <http://nlp.stanford.edu/software/lex-parser.shtml>

Fig. 3. Number of support vectors (displayed beside each data point) as a function of the sampling size for the 1-slack algorithm

We note that:

- changing the sample size allows us to tune the trade-off between accuracy and speed. Indeed, as we increase the sampling size, the F_1 -score grows until it reaches the accuracy of the exact SVM (at 5,000). In this case, the uSVM produces the same accuracy of SVM while being 10 times faster.
- with a sample size of 2,500 the accuracy of the uSVM is only 1.0 percentage point apart from the exact model whereas the training time savings are of a factor over 50. This corresponds to a training time smaller than 4 hours for uSVM vs. 7.5 days for SVM.
- finally, we note that our reported F_1 -score for boundary classification is state-of-the-art only if tree kernels are used, e.g. [21].

5.4 Producing a Sparse Model

Another interesting dimension to compare sampling algorithms with the exact SVM solver would be to evaluate the sparseness of the produced solutions. As we have already mentioned in Section 3.1, uSVM and iSVM employ the 1-slack formulation that produces very sparse models [7]. This becomes especially important when very large datasets are used, as the improved sparsity can greatly reduce the classification time. Figure 3 is different from Figure 2, as it displays the classification results for sec23 subset of SRL dataset along with the number of support vectors (plotted beside a data point) learned by each model.

Indeed, the number of support vectors grows, as we increase the sampling size in an attempt to match the accuracy of the exact SVM. The model learned by the exact SVM contains 61,881 support vectors, while uSVM model produces 41,714 support vectors. We would like to have more experimental data to make

Fig. 4. F_1 -score (left) and training time (right) as a function of margin parameter C on the 100,000 subset of SRL dataset (C is on a logarithmic scale)

Fig. 5. F_1 -score as a function of margin parameter C on the DESC (left) and NUM (right) categories of Question Classification dataset (C is on a logarithmic scale)

the comparison more sound but the slow training of the exact SVM makes this endeavor almost infeasible (just the training of the exact SVM on 1 million of examples takes over 7.5 days!)

5.5 Fast Parameterization

On very large datasets, finding the best set of parameters for an exact solver, e.g. SVM-light, using structural kernels becomes prohibitively expensive, thus greatly limiting the possibility to find the best performing model. Hence, we test the idea of fast parameter selection for the exact SVM using CPAs and small samples.

We chose the trade-off parameter, C , as our target parameter and we select a subset of 100,000 examples from SRL data to train uSVM, iSVM, and SVM with $C \in \{0.01, 0.1, 1, 10, 100, 1000\}$. We could not use 1 million dataset since SVM prevents to carry out experiments within a tractable time frame. The left plot of Figure 4 shows that F_1 of the three models has the same behavior according to C . Thus, we can select the optimum value according to the fast method (e.g. with a sample size 1000) to estimate the best value of C of SVM.

Moreover, it is interesting to observe the plot on the right of Figure 4. This shows the training time on the previous subset with respect to C values. It reveals

Table 1. F_1 measured on two testsets, *sec23* and *sec24*, for five different kernels: ST, SST, SST-bow, PT, and uPT kernels, trained on 1 million instances. The best results are shown in bold. The training time is given in minutes. The bottom row is the performance of SVM trained on only **100k** examples.

Sample size	ST			SST			SST-bow			PT			uPT		
	F_1		time	F_1		time	F_1		time	F_1		time	F_1		time
	<i>sec23</i>	<i>sec24</i>		<i>sec23</i>	<i>sec24</i>		<i>sec23</i>	<i>sec24</i>		<i>sec23</i>	<i>sec24</i>		<i>sec23</i>	<i>sec24</i>	
100	6.8	6.5	3.9	74.9	72.6	1.0	74.3	73.0	2.2	71.9	70.7	3.7	73.3	71.4	3.7
250	14.2	13.0	14.4	78.4	76.2	4.1	78.5	76.3	6.1	74.6	73.2	13.7	76.5	74.6	14.2
500	20.3	18.7	46.6	80.2	77.2	14.0	79.5	77.3	16.9	76.3	74.4	45.0	78.3	76.3	47.6
1000	23.5	21.2	143.7	82.0	79.7	45.3	81.3	79.0	55.9	78.2	76.2	158	79.6	76.9	158.8
SVM	12.6	10.94	213.8	80.78	78.56	37.5	80.38	78.13	42.2	74.39	73.47	89.1	77.54	75.87	100.4

that the use of sampling algorithms, particularly uSVM, provides substantial speed-up in the training time, especially when large values for C are used. Not surprisingly, uSVM is a preferable choice, since it has a constant-time scaling behavior. These results show that the proposed algorithms can provide fast and reliable parameter search for the best model.

To generalize the findings above we carried out the same experiment on the second dataset. We ran tests for a range of values on all six categories of the QC dataset. Since each category has only 5500 examples, here, however, the main concern is not the training time, but the ability of uSVM and iSVM to match the behavior of the exact SVM with respect to the parameter C . For brevity, we provide the results only for DESC and NUM categories. Figures 5 shows that both sampling algorithms match well the behavior of the exact SVM for DESC and NUM categories of QC dataset.

5.6 Kernel Testing and Selection

The previous section has shown that sampling algorithms can be used for fast selection of the parameter C . Here we investigate if the same approach can be applied for the efficient selection of the best kernel. The aim is to check if a very fast uSVM algorithm, i.e. using a small sample size, can identify the best performing kernel. Thus, we ran uSVM algorithm on 1 million subset of SRL dataset by varying the sample size and using five different structural kernels: ST, SST, SST-bow, PT and uPT kernels. The results reported in Table 1 show that:

- uSVM has a consistent behavior across different sample sizes for all kernel functions. This suggests that a small sample, e.g. 100 training examples, can be used to select the most suitable kernel for a given task in a matter of a couple of minutes.
- the bottom row of the Table 1 reports the results of SVM trained on a smaller subset of 100k examples (only one experiment with SVM on a subset of 1 million examples takes over 7 days), which demonstrates that the results obtained with uSVM are consistent with the exact solver.

Fig. 6. F_1 -score as a function of λ decay factor (length of the child sequences) for SST kernel. Training carried out by uSVM on SRL (1million examples) for samples sizes 100, 250 and 500.

- as already happened in similar tasks, SST kernel appears to provide the best performance with respect to the training time and accuracy whereas the ST kernel, as expected, cannot generate enough features to characterize correct or incorrect boundary. Indeed, its F_1 is very low even when large amount of data is used.
- surprisingly the SST-bow kernel which simply adds bow to SST is less accurate. This suggests that adding words can reduce the generalization ability of the tree kernels, while the syntactic information is very important like in SRL. This aspect is confirmed by the F_1 of uPT, which is higher than PT. Indeed, the former does not generate *pure* lexical features, i.e. words, whereas the latter does.
- finally, since we can quickly pick the most appropriate kernel, we can also perform a fast tuning of kernel hyper-parameters. For this task, we experiment with λ , which is one of the most important factors defining tree kernel performance. Figure 6 displays the behavior of F_1 with respect to the range of λ values. The plot shows that F_1 varies considerably so several values should be tested. We carried out these experiment in a few minutes but using SVM we would have required several weeks.

6 Conclusions

In this paper, we have studied the cutting plane technique for training SVMs and we have shown that it is also effective with structural kernels. The experiments on Semantic Role Labeling and Question Classification show very promising results with respect to accuracy and efficiency. Our major achievement is a speed-up factor of over 10 compared to the exact SVM, while obtaining the same precise

solution. In addition, the proposed method gives the flexibility to train very fast models by trading off accuracy for the training time.

The main idea promoted in the paper is that CPA with sampling, while being as accurate as exact SVM, provides the possibility to quickly select optimal kernels or model parameters. Unlike other approximate techniques, which use small subsets of original data to train a model, CPA uses smaller samples only to approximate the subgradient while working on the entire dataset. Our experiments show that parameters/kernels selected by CPA are also optimal for the exact SVM. To the best of our knowledge, this is the first attempt to demonstrate this fact.

References

1. Carreras, X., Màrquez, L.: Introduction to the CoNLL-2005 Shared Task: Semantic Role Labeling. In: Proceedings of the 9th Conference on Natural Language Learning, CoNLL-2005, Ann Arbor, MI, USA (2005)
2. Charniak, E.: A maximum-entropy-inspired parser. In: ANLP, pp. 132–139 (2000)
3. Collins, M., Duffy, N.: New ranking algorithms for parsing and tagging: Kernels over discrete structures, and the voted perceptron. In: ACL, pp. 263–270 (2002)
4. Fine, S., Scheinberg, K.: Efficient svm training using low-rank kernel representations. *Journal of Machine Learning Research* 2, 243–264 (2001)
5. Franc, V., Sonnenburg, S.: Optimized cutting plane algorithm for support vector machines. In: Cohen, W.W., McCallum, A., Roweis, S.T. (eds.) *ICML. ACM International Conference Proceeding Series*, vol. 307, pp. 320–327. ACM, New York (2008)
6. Joachims, T.: Making large-scale SVM learning practical. In: Schölkopf, B., Burges, C., Smola, A. (eds.) *Advances in Kernel Methods - Support Vector Learning*, ch. 11, pp. 169–184. MIT Press, Cambridge (1999)
7. Joachims, T.: Training linear SVMs in linear time. In: *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pp. 217–226 (2006)
8. Joachims, T., Yu, C.N.J.: Sparse kernel svms via cutting-plane training. *Machine Learning* 76(2-3), 179–193 (2009); *European Conference on Machine Learning (ECML) Special Issue*
9. Keerthi, S.S., Chapelle, O., Decoste, D., Bennett, P., Parrado-hernndez, E.: Building support vector machines with reduced classifier complexity. *Journal of Machine Learning Research* 8, 2006 (2001)
10. Klein, D., Manning, C.D.: Fast exact inference with a factored model for natural language parsing. In: Becker, S., Thrun, S., Obermayer, K. (eds.) *NIPS*, pp. 3–10. MIT Press, Cambridge (2002)
11. Kudo, T., Matsumoto, Y.: Fast methods for kernel-based text analysis. In: *Proceedings of ACL'03* (2003)
12. Leslie, C., Eskin, E., Cohen, A., Weston, J., Noble, W.S.: Mismatch string kernels for discriminative protein classification. *Bioinformatics* 20(4), 467–476 (2004)
13. Li, X., Roth, D.: Learning question classifiers: the role of semantic information. *Natural Language Engineering* 12(3), 229–249 (2006)
14. Marcus, M., Santorini, B., Marcinkiewicz, M.: Building a large annotated corpus of English: the Penn Treebank. *Computational Linguistics* 19(2), 313–330 (1993)
15. Moschitti, A., Pighin, D., Basili, R.: Tree kernels for semantic role labeling. *Computational Linguistics* 34(2), 193–224 (2008)

16. Moschitti, A., Zanzotto, F.: Fast and effective kernels for relational learning from texts. In: Ghahramani, Z. (ed.) *Proceedings of the 24th Annual International Conference on Machine Learning, ICML 2007* (2007)
17. Moschitti, A.: Making tree kernels practical for natural language learning. In: *EACL. The Association for Computer Linguistics* (2006)
18. Moschitti, A.: Kernel methods, syntax and semantics for relational text categorization. In: *Proceeding of CIKM '08, NY, USA* (2008)
19. Moschitti, A., Quarteroni, S., Basili, R., Manandhar, S.: Exploiting syntactic and shallow semantic kernels for question/answer classification. In: *Proceedings of ACL'07* (2007)
20. Palmer, M., Kingsbury, P., Gildea, D.: The proposition bank: An annotated corpus of semantic roles. *Computational Linguistics* 31(1), 71–106 (2005)
21. Pighin, D., Moschitti, A.: Efficient linearization of tree kernel functions. In: *Proceedings of the Thirteenth Conference on Computational Natural Language Learning (CoNLL-2009)*, pp. 30–38. Association for Computational Linguistics, Boulder (June 2009), <http://www.aclweb.org/anthology/W09-1106>
22. Pradhan, S., Hacıoglu, K., Krugler, V., Ward, W., Martin, J.H., Jurafsky, D.: Support vector learning for semantic argument classification. *Mach. Learn.* 60(1-3), 11–39 (2005)
23. Rieck, K., Krueger, T., Brefeld, U., Mueller, K.R.: Approximate tree kernels. *Journal of Machine Learning Research* 11, 555–580 (2010)
24. Shalev-Shwartz, S., Singer, Y., Srebro, N.: Pegasos: Primal estimated sub-gradient solver for SVM. In: Ghahramani, Z. (ed.) *ICML. International Conference Proceeding Series*, vol. 227, pp. 807–814. ACM, New York (2007)
25. Shen, L., Joshi, A.K.: An SVM-based voting algorithm with application to parse reranking. In: Daelemans, W., Osborne, M. (eds.) *Proceedings of CoNLL HLT-NAACL 2003*, pp. 9–16 (2003), <http://www.aclweb.org/anthology/W03-0402.pdf>
26. Surdeanu, M., Ciaramita, M., Zaragoza, H.: Learning to rank answers on large online QA collections. In: *Proceedings of ACL-08, HLT, Columbus, Ohio* (2008), <http://www.aclweb.org/anthology/P/P08/P08-1082>
27. Williams, C., Seeger, M.: Using the nystrm method to speed up kernel machines. In: *Advances in Neural Information Processing Systems*, vol. 13, pp. 682–688. MIT Press, Cambridge (2001)
28. Yu, C.N.J., Joachims, T.: Training structural svms with kernels using sampled cuts. In: *ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, pp. 794–802 (2008)
29. Zaki, M.J.: Efficiently mining frequent trees in a forest. In: *KDD '02: Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 71–80. ACM Press, New York (2002)
30. Zhang, D., Lee, W.S.: Question classification using support vector machines. In: *SIGIR*, pp. 26–32. ACM, New York (2003)
31. Zhang, M., Zhang, J., Su, J.: Exploring Syntactic Features for Relation Extraction using a Convolution tree kernel. In: *Proceedings of NAACL, New York City, USA*, pp. 288–295 (2006), <http://www.aclweb.org/anthology/N/N06/N06-1037>