

Tree Kernels for Semantic Role Labeling

Alessandro Moschitti*
University of Trento - DIT

Daniele Pighin**
University of Trento - DIT, FBK-irst

Roberto Basili†
University of Rome *Tor Vergata* - DISP

The availability of large scale data sets of manually annotated predicate argument structures has recently favored the use of Machine Learning approaches to the design of automated Semantic Role Labeling (SRL) systems. The main research in this area relates to the design choices for feature representation and for effective decompositions of the task in different learning models. Regarding the former choice, structural properties of full syntactic parses are largely employed as they represent ways to encode different principles suggested by the linking theory between syntax and semantics. The latter choice relates to several learning schemes over global views of the parses. For example, re-ranking stages operating over alternative predicate-argument sequences of the same sentence have shown to be very effective.

In this paper, we propose several kernel functions to model parse tree properties in kernel-based machines, e. g. Perceptrons or Support Vector Machines. In particular, we define different kinds of tree kernels as general approaches to feature engineering in SRL. Moreover, we extensively experiment with such kernels to investigate their contribution on individual stages of an SRL architecture both in isolation and in combination with other traditional manually-coded features. The results on boundary recognition, classification and re-ranking stages provide systematic evidence about the significant impact of tree kernels on the overall accuracy, especially when the amount of training data is small. As a conclusive result, tree kernels allow for a general and easily portable feature engineering method which is applicable to a large family of Natural Language Processing tasks.

1. Introduction

Conspicuous attention has recently been devoted to the design of systems for the automatic labeling of semantic roles (SRL) as defined in two important projects: FrameNet (Baker, Fillmore, and Lowe 1998), based on Frame Semantics, and PropBank (Palmer, Gildea, and Kingsbury 2005) inspired by Levin's verb classes. To annotate natural language sentences, such systems generally require (1) the detection of the target word embodying the predicate and (2) the detection and classification of the word sequences constituting the predicate arguments.

* Via Sommarive, 14 I-38050 Povo (TN). Email: moschitti@dit.unitn.it

** Via Sommarive, 18 I-38050 Povo (TN). Email: pighin@itc.it

† Via del Politecnico, 1 I-00133 RM. Email: basili@info.uniroma2.it

Submission received: 15 July 2006
Revised submission received: 01 May 2007
Accepted for publication: 19 June 2007

Previous work has shown that the above steps can be carried out by applying Machine Learning techniques (Carreras and Màrquez 2004, Carreras and Màrquez 2005, Litkowski 2004), for which the most important features encoding predicate argument relations are derived from (shallow or deep) syntactic information. The outcome of this research brings wide empirical evidence in favor of the linking theories between semantics and syntax, e. g. (Jackendoff 1990). Nevertheless, as no such theory provides a sound and complete treatment, the choice and design of syntactic features to represent semantic structures requires remarkable research effort and intuition.

For example, earlier studies on feature design for Semantic Role Labeling were carried out in (Gildea and Jurafsky 2002, Thompson, Levy, and Manning 2003). Since then, researchers have proposed several syntactic feature sets, where the more recent set slightly enhanced the older ones.

A careful analysis of such features reveals that most of them are syntactic tree fragments of training sentences, thus a viable way to alleviate the feature design complexity is the adoption of syntactic tree kernels (Collins and Duffy 2002). For example, in (Moschitti 2004), the predicate-argument relation is represented by means of the minimal subtree that includes both of them. The similarity between two instances is evaluated by a tree kernel function in terms of common substructures. Such approach is in line with current research on kernels for natural language learning, e. g. syntactic parsing re-ranking (Collins and Duffy 2002), relation extraction (Zelenko, Aone, and Richardella 2003) and named entity recognition (Cumby and Roth 2003, Culotta and Sorensen 2004).

Furthermore, recent work (Punyakanok et al. 2005, Haghghi, Toutanova, and Manning 2005) has shown that, to achieve high labeling accuracy, joint inference should be applied on the whole predicate argument structure. For this purpose, we need to extract features from the sentence syntactic parse tree that encodes the relationships governing complex semantic structures. This task is rather difficult since we do not exactly know which syntactic clues effectively capture the relation between the predicate and its arguments. For example, to detect the interesting context, the modeling of syntax/semantics-based features should take into account linguistic aspects like ancestor nodes or semantic dependencies (Toutanova, Markova, and Manning 2004). In this scenario, the automatic feature generation/selection carried out by tree kernels can provide useful insights on the underlying linguistic phenomena. Other advantages coming from the use of tree kernels are the following:

First, we can implement them very quickly as the feature extractor module only requires the writing of a general procedure for subtree extraction. In contrast, traditional SRL systems use more than thirty features, e. g. (Pradhan et al. 2005a), each of which requires the writing of a dedicated procedure.

Second, their combination with traditional attribute-value models produces more accurate systems, also when using the same Machine Learning algorithm in the combination, since the feature spaces are very different.

Third, we can carry out feature engineering using kernel combinations and marking strategies (Moschitti et al. 2005a, Moschitti, Pighin, and Basili 2006). This allows us to boost the SRL accuracy in a relatively simple way.

Next, tree kernels generate large tree fragment sets which constitute back-off models for important syntactic features. Using them, the learning algorithm generalizes better and produces a more accurate classifier, especially when the amount of training data is scarce.

Finally, once the learning algorithm using tree kernels has converged, we can identify the most important structured features of the generated model. One approach for

such reverse engineering process relies on the computation of the explicit feature space, at least for the highest weighted features (Kudo and Matsumoto 2003). Once the most relevant fragments are available, they can be used to design novel effective attribute-value features (which in turn can be used to design more efficient classifiers, e. g. with linear kernels) and inspire new linguistic theories.

The above points suggest that tree kernels should always be applied, at least for an initial study of the problem. Unfortunately, they suffer from two main limitations: (a) poor impact on boundary detection as, in this task, correct and incorrect arguments may share a large portion of the encoding trees (Moschitti 2004); and (b) more expensive running time and limited contribution to the overall accuracy if compared with manually derived features (Cumby and Roth 2003). Point (a) has been addressed in (Moschitti, Pighin, and Basili 2006) by showing that a marking strategy of relevant parse-tree nodes makes correct and incorrect subtrees for boundary detection quite different. Point (b) can be tackled by studying approaches to kernel engineering that allow for the design of efficient and effective kernels.

In this article, we provide a comprehensive study on the use of tree kernels for Semantic Role Labeling. For such purpose, we define tree kernels based on the composition of two different feature functions: *canonical mappings*, which map sentence-parse trees in tree structures encoding semantic information, and *feature extraction functions*, which encode the previous trees in the actual feature space. The latter functions *explode* the canonical trees in all their substructures and, in the literature, are usually referred to as *tree kernels*. For instance, in (Collins and Duffy 2002, Vishwanathan and Smola 2002, Moschitti 2006a) different tree kernels extract different types of tree fragments.

Given the heuristic nature of canonical mappings, we studied their properties by experimenting with them within Support Vector Machines and with the data set provided by CoNLL-shared tasks (Carreras and Màrquez 2005). The results show that carefully engineered tree kernels always boost the accuracy of the basic systems. Most importantly, in complex tasks such as the re-ranking of semantic role annotations, they provide an easy way to engineer new features which enhance the state-of-the-art on SRL.

In the remainder of this article, Section 2 presents traditional architectures for Semantic Role Labeling and the typical features proposed in literature. Tree kernels are formally introduced in Section 3, whereas Section 4 describes our modular architecture employing Support Vector Machines along with manually designed features, tree kernels (feature extraction functions) and their combinations. Section 5 presents our structured features (canonical mappings) inducing different kernels that we used for different SRL subtasks. The extensive experimental results obtained on the boundary recognition, role classification and re-ranking stages are presented in Section 6. Finally, Section 7 summarizes the conclusions.

2. Automatic Shallow Semantic Parsing

The recognition of semantic structures within a sentence relies on lexical and syntactic information provided by early stages of an NLP process, such as lexical analysis, part of speech tagging and syntactic parsing. The complexity of the SRL task mostly depends on two aspects: (a) the above information is generally noisy, i. e. in a real-world scenario the accuracy and reliability of NLP subsystems are generally not very high; (b) the lack of a sound and complete linguistic or cognitive theory about the links between syntax and semantics does not allow an informed, deductive approach to the problem. Nevertheless, the large amount of available lexical and syntactic information favors the

application of inductive approaches to the SRL task, which indeed is generally treated as a combination of statistical classification problems.

The next sections define the SRL task more precisely and summarize the most relevant work carried out to address the above two problems.

2.1 Problem Definition

The most well known shallow semantic theories regard two different projects: PropBank (Palmer, Gildea, and Kingsbury 2005) and FrameNet (Baker, Fillmore, and Lowe 1998). The former is based on a linguistic model inspired by Levin’s verb classes (Levin 1993), focusing on the argument structure of verbs and on the alternation patterns that describe movements of verbal arguments within a predicate structure. The latter refers to the application of Frame Semantics (Fillmore 1968) in the annotation of predicate argument structures based on frame elements (semantic roles). The above theories have been investigated in two CoNLL shared tasks (Carreras and Màrquez 2004, Carreras and Màrquez 2005) and a Senseval-3 evaluation (Litkowski 2004), respectively.

Given a sentence and a predicate word, an SRL system outputs an annotation of the sentence in which the sequences of words that make up the arguments of the predicate are properly labeled, e. g. :

[Arg0 He] **got** [Arg1 his money] [c-v back]¹

in response to the input *He got his money back*. This processing requires that: (1) the predicates within the sentence are identified and (2) the word sequences that span the boundaries of each predicate argument are delimited and assigned the proper role label.

The first sub-task can be performed either using statistical methods or hand-crafting lexical and syntactic rules. In the case of verbal predicates, it is quite easy to write simple rules matching regular expressions built on POS tags. The second task is more complex and is generally viewed as a combination of statistical classification problems: the learning algorithms are trained to recognize the extension of predicate arguments and the semantic role they play.

2.2 Models for Semantic Role Labeling

An SRL model and the resulting architecture are largely influenced by the kind of data available for the task. As an example, a model relying on a shallow syntactic parser would assign roles to chunks, whereas with a full syntactic parse of the sentence it would be straightforward to establish a correspondence between nodes of the parse tree and semantic roles. We focused on the latter as it has been shown to be more accurate by the CoNLL 2005 shared task results.

According to the deep syntactic formulation, the classifying instances are pairs of parse-tree nodes which dominate the exact span of the predicate and the target argument. Such pairs are usually represented in terms of attribute-value vectors, where the attributes describe properties of predicates, arguments and the way they are related. There is large agreement on an effective set of linguistic features (Gildea and Jurafsky 2002, Pradhan et al. 2005a) that have been employed in the vast majority of SRL systems. The most relevant features are summarized in Table 1.

¹ In PropBank notation, Arg0 and Arg1 represent the logical subject and the logical object of the target verbal predicate, respectively. C-V represents the particle of a phrasal-verb predicate.

Table 1
Standard linguistic features employed by most SRL systems.

Feature Name	Description
Predicate	Lemmatization of the predicate word
Path	Syntactic path linking the predicate and an argument, e.g. NN↑NP↑VP↓VBX
Partial path	<i>Path</i> feature limited to the branching of the argument
No-direction path	Like <i>Path</i> , but without traversal directions
Phrase type	Syntactic type of the argument node
Position	Relative position of the argument with respect to the predicate
Voice	Voice of the predicate, i. e. active or passive
Head word	Syntactic head of the argument phrase
Verb subcategorization	Production rule expanding the predicate parent node
Named entities	Classes of named entities that appear in the argument node
Head word POS	POS tag of the argument node head word (less sparse than Head word)
Verb clustering	Type of verb → direct object relation
Governing Category	Whether the candidate argument is the verb subject or object
Syntactic Frame	Position of the NPs surrounding the predicate
Verb sense	Sense information for polysemous verbs
Head word of PP	Enriched POS of prepositional argument nodes (e.g. PP-for, PP-in)
First and last word/POS	First and last words and POS tags of candidate argument phrases
Ordinal position	Absolute offset of a candidate argument within a proposition
Constituent tree distance	Distance from the predicate with respect to the parse tree
Constituent features	Description of the constituents surrounding the argument node
Temporal Cue Words	Temporal markers which are very distinctive of some roles

Once the representation for the predicate-argument pairs is available, a multi-classifier is used to recognize the correct node pairs, i. e. nodes associated with correct arguments (given a predicate), and assign them a label (which is the label of the argument). This can be achieved by training a multi-classifier on $N + 1$ classes, where the first N classes correspond to the different roles and the $(N + 1)^{th}$ is a *NARG* (non-argument) class to which non-argument nodes are assigned.

A more efficient solution consists in dividing the labeling process in two steps: boundary detection and argument classification. A *Boundary Classifier* (BC) is a binary classifier that recognizes the tree nodes that exactly cover a predicate argument, i. e. that dominate all and only the words that belong to target arguments. Then, such nodes are classified by a *Role Multi-classifier* (RM) that assigns to each example the most appropriate label. This two-step approach (Gildea and Jurafsky 2002) has the advantage of only applying BC on all parse tree nodes. RM can ignore non-boundary nodes, resulting in a much faster classification. Other approaches have extended this solution and suggested other multi-stage classification models, e. g. (Moschitti et al. 2005b) in which a four-step hierarchical SRL architecture is described.

After node labeling has been carried out, it is possible that the output of the argument classifier does not result in a consistent annotation, as the labeling scheme may not be compatible with the underlying linguistic model. As an example, PropBank-style annotations do not allow arguments to be nested. This happens when two or more *overlapping* tree nodes, i. e. one dominating the other, are classified as positive boundaries.

The simplest solution relies on the application of heuristics that take into account the whole predicate argument structure to remove the incorrect labels, e. g. (Moschitti et al. 2005a, Tjong Kim Sang et al. 2005). A much more complex solution consists in the application of some joint inference model to the whole predicate argument structure,

as in (Pradhan et al. 2004). As an example, (Haghighi, Toutanova, and Manning 2005) associate a posterior probability with each argument node role assignment, estimate the likelihood of the alternative labeling schemes and employ a re-ranking mechanism to select the best annotation.

Additionally, the most accurate systems of the CoNLL 2005 shared task (Punyakanok et al. 2005, Pradhan et al. 2005b) use different syntactic views of the same input sentence. This allows the SRL system to recover from syntactic parser errors: for example, a propositional phrase specifying the direct object of the predicate would be attached to the verb instead of the argument. This kind of errors prevents some arguments of the proposition from being recognized, as: (1) there may not be a node of the parse tree dominating (all and only) the words of the correct sequence; (2) a badly attached tree node may invalidate other argument nodes, generating unexpected overlapping situations.

The manual design of features which capture important properties of complete predicate argument structures (also coming from different syntactic views) is quite complex. Tree kernels are a valid alternative to manual design as the next sections point out.

3. Tree Kernels

Tree kernels have been applied to reduce the feature design effort in the context of several natural language tasks, e.g. syntactic parsing re-ranking (Collins and Duffy 2002), relation extraction (Zelenko, Aone, and Richardella 2003), named entity recognition (Cumby and Roth 2003, Culotta and Sorensen 2004) and Semantic Role Labeling (Moschitti 2004).

On the one hand, these studies show that the kernel ability to generate large feature sets is useful to quickly model new and not well understood linguistic phenomena in learning machines. On the other hand, they show that sometimes it is possible to manually design features for linear kernels that produce higher accuracy and faster computation time. One of the most important causes of such mixed behavior is the inappropriate choice of kernel functions. For example, in (Moschitti, Pighin, and Basili 2006, Moschitti 2006a) several kernels have been designed and shown to produce different impacts on the training algorithms.

In the next sections, we briefly introduce the kernel trick and describe the *subtree* (ST) kernel devised in (Vishwanathan and Smola 2002), the *subset tree* (SST) kernel defined in (Collins and Duffy 2002) and the partial tree (PT) kernel proposed in (Moschitti 2006a).

3.1 Kernel Trick

The main concept underlying Machine Learning for classification tasks is the automatic learning of classification functions based on examples labeled with the class information. Such examples can be described by means of feature vectors in an n dimensional space over the real numbers, i. e. \mathfrak{R}^n . The learning algorithm uses space metrics over vectors, e.g. the scalar product, to learn an abstract representation of all instances belonging to the target class.

For example, Support Vector Machines are linear classifiers which learn a hyperplane $f(\vec{x}) = \vec{w} \times \vec{x} + b = 0$, separating positive from negative examples. \vec{x} is the feature vector representation of a classifying object o , whereas $\vec{w} \in \mathfrak{R}^n$ and $b \in \mathfrak{R}$ are parameters learned from the data by applying the *Structural Risk Minimization principle* (Vapnik

1998). The object o is mapped to \vec{x} via a feature function $\phi : \mathcal{O} \rightarrow \mathbb{R}^n$, \mathcal{O} being the set of the objects that we want to classify. o is categorized in the target class only if $f(\vec{x}) \geq 0$.

The kernel trick allows us to rewrite of the decision hyperplane as:

$$f(\vec{x}) = \left(\sum_{i=1..l} y_i \alpha_i \vec{x}_i \right) \cdot \vec{x} + b = \sum_{i=1..l} y_i \alpha_i \vec{x}_i \cdot \vec{x} + b = \sum_{i=1..l} y_i \alpha_i \phi(o_i) \cdot \phi(o) + b = 0.$$

where y_i is equal to 1 for positive examples and -1 for negative examples, $\alpha_i \in \mathbb{R}$ with $\alpha_i \geq 0$, $o_i \forall i \in \{1, \dots, l\}$ are the training instances and the product $K(o_i, o) = \langle \phi(o_i) \cdot \phi(o) \rangle$ is the kernel function associated with the mapping ϕ .

Note that we do not need to apply the mapping ϕ , we can use $K(o_i, o)$ directly. This allows us, under Mercer's conditions (Shawe-Taylor and Cristianini 2004), to define abstract kernel functions which generate implicit feature spaces. A traditional example is given by the polynomial kernel: $K_p(o_1, o_2) = (c + \vec{x}_1 \cdot \vec{x}_2)^d$, where c is a constant and d is the degree of the polynomial. This kernel generates the space of all conjunctions of feature groups up to d elements.

Additionally, we can carry out two interesting operations:

- kernel combinations, e. g. $K_1 + K_2$ or $K_1 \times K_2$, and
- feature mapping compositions, e. g.

$$K(o_1, o_2) = \langle \phi(o_1) \cdot \phi(o_2) \rangle = \langle \phi_B(\phi_A(o_1)) \cdot \phi_B(\phi_A(o_2)) \rangle.$$

Kernel combinations are very useful to integrate the knowledge provided by the manually defined features with the one automatically obtained with structural kernels; feature mapping compositions are useful methods to describe diverse kernel classes (see Section 5). In this perspective, we propose to split the mapping ϕ by defining our tree kernel as follows:

- *Canonical Mapping*, $\phi_M()$, in which a linguistic object, e. g. a syntactic parse tree, is transformed into a more meaningful structure, e. g. the subtree corresponding to a verb subcategorization frame.
- *Feature Extraction*, $\phi_S()$, which maps the canonical structure in all its fragments according to different fragment spaces S , e. g. ST, SST and PT.

For example, given the kernel $K_{ST} = \phi_{ST}(o_1) \cdot \phi_{ST}(o_2)$, we can apply a canonical mapping $\phi_M()$, obtaining $K_{ST}^M = \phi_{ST}(\phi_M(o_1)) \cdot \phi_{ST}(\phi_M(o_2)) = (\phi_{ST} \circ \phi_M)(o_1) \cdot (\phi_{ST} \circ \phi_M)(o_2)$, which is a noticeably different kernel induced by the mapping $(\phi_{ST} \circ \phi_M)$.

In the remainder of this section we start the description of our engineered kernels by defining three different feature extraction mappings based on three different kernel spaces, i. e. ST, SST and PT.

3.2 Tree kernel Spaces

The kernels that we consider represent trees in terms of their substructures (fragments). The kernel function detects if a tree subpart (common to both trees) belongs to the feature space that we intend to generate. For such purpose, the desired fragments need to be described. We consider three main characterizations: the subtrees (STs) (Vishwanathan and Smola 2002), the subset trees (SSTs) or all subtrees (Collins and Duffy 2002) and the partial trees (PTs) (Moschitti 2006a).

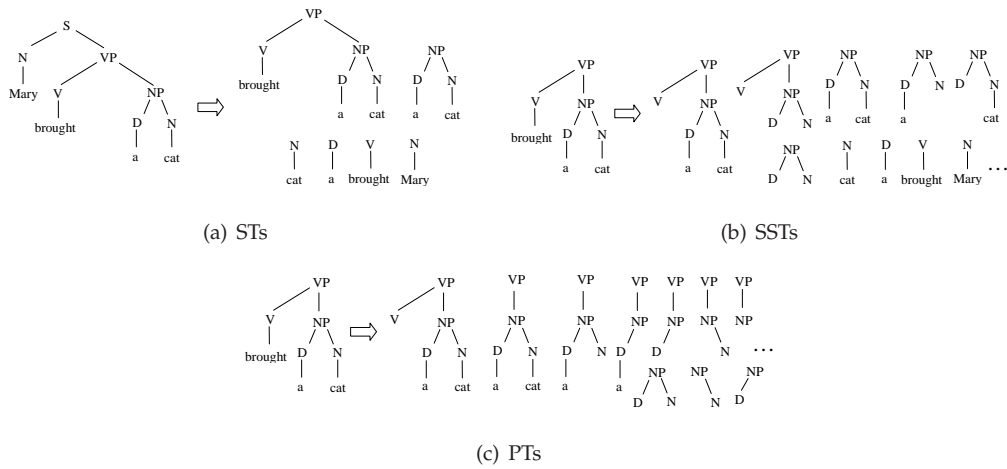


Figure 1
Example of ST (a), SST (b) and PT (c) fragments.

As we consider syntactic parse trees, each node with its children is associated with a grammar production rule, where the symbol at the left-hand side corresponds to the parent and the symbols at the right-hand side are associated with the children. The terminal symbols of the grammar are always associated with tree leaves.

A **subtree** (ST) is defined as a tree rooted in any non-terminal node along with all its descendants. For example, Figure 1(a) shows the parse tree of the sentence *Mary brought a cat* together with its six STs. A **subset tree** (SST) is a more general structure since its leaves can be non-terminal symbols. For example, Figure 1(b) shows ten SSTs (out of 17) of the subtree in Frame (a) rooted in VP. SSTs satisfy the constraint that grammatical rules cannot be broken. For example, [VP [V NP]] is an SST which has two non-terminal symbols, V and NP, as leaves. On the contrary, [VP [V]] is not an SST as it violates the production $VP \rightarrow V NP$. If we relax the constraint over the SSTs, we obtain a more general form of substructures that we call **partial trees** (PTs). These can be generated by the application of partial production rules of the grammar, consequently [VP [V]] and [VP [NP]] are valid PTs. It is worth to note that PTs consider the position of the children as, for example, [A [B][C][D]] and [A [D][C][B]] only share the single children, i.e. [A [B]], [A [C]] and [A [D]].

Figure 1(c) shows that the number of PTs derived from the same tree as before is still higher (i. e. 30 PTs). These numbers provide an intuitive quantification of different degrees of information encoded by each representation.

3.3 Feature Extraction Functions

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. In the following, we report the SubSet Tree (SST) kernel proposed in (Collins and Duffy 2002). The algorithms to efficiently compute it along with the ST and PT kernels can be found in (Moschitti 2006a).

Given two trees T_1 and T_2 , let $\{f_1, f_2, \dots\} = \mathcal{F}$ be the set of substructures (fragments) and $I_i(n)$ be equal to 1 if f_i is rooted at node n , 0 otherwise. Collins and Duffy’s kernel is defined as

$$K(T_1, T_2) = \sum_{n_1 \in N_{T_1}} \sum_{n_2 \in N_{T_2}} \Delta(n_1, n_2), \quad (1)$$

where N_{T_1} and N_{T_2} are the sets of nodes in T_1 and T_2 , respectively and $\Delta(n_1, n_2) = \sum_{i=1}^{|\mathcal{F}|} I_i(n_1)I_i(n_2)$. The latter is equal to the number of common fragments rooted in nodes n_1 and n_2 . Δ can be computed as follows:

1. if the productions (i.e. the nodes with their direct children) at n_1 and n_2 are different then $\Delta(n_1, n_2) = 0$;
2. if the productions at n_1 and n_2 are the same, and n_1 and n_2 only have leaf children (i.e. they are pre-terminal symbols) then $\Delta(n_1, n_2) = 1$;
3. if the productions at n_1 and n_2 are the same, and n_1 and n_2 are not pre-terminals then $\Delta(n_1, n_2) = \prod_{j=1}^{nc(n_1)} (1 + \Delta(c_{n_1}^j, c_{n_2}^j))$, where $nc(n_1)$ is the number of children of n_1 and c_n^j is the j -th child of n .

Such tree kernel can be normalized and a λ factor can be added to reduce the weight of large structures (refer to (Collins and Duffy 2002) for a complete description).

3.4 Related Work

Although literature on SRL is extensive, there is almost no study on the use of tree kernels for its solution. Consequently, the reported research is mainly based on diverse natural language learning problems tackled by means of tree kernels.

In (Collins and Duffy 2002), the SST kernel was experimented with the Voted Perceptron for the parse tree re-ranking task. A combination with the original PCFG model improved the syntactic parsing. Another interesting kernel for re-ranking was defined in (Toutanova, Markova, and Manning 2004). This represents parse trees as lists of paths (leaf projection paths) from leaves to the top level of the tree. It is worth noting that the PT kernel includes tree fragments identical to such paths.

In (Kazama and Torisawa 2005), an interesting algorithm that speeds up the average running time is presented. Such algorithm looks for node pairs in which are rooted subtrees that share many substructures (*malicious nodes*) and applies a transformation to the trees rooted in such nodes to make the kernel computation faster. The results show an increase of the speed of hundreds of times with respect to the basic implementation.

In (Zelenko, Aone, and Richardella 2003), two kernels over syntactic shallow parser structures were devised for the extraction of linguistic relations, e. g. *person-affiliation*. To measure the similarity between two nodes, the *contiguous string kernel* and the *sparse string kernel* were used. In (Culotta and Sorensen 2004) such kernels were slightly generalized by providing a matching function for the node pairs. The time complexity for their computation limited the experiments on a data set of just 200 news items.

In (Shen, Sarkar, and Joshi 2003), a tree-kernel based on Lexicalized Tree Adjoining Grammar (LTAG) for the parse re-ranking task was proposed. The subtrees induced by such kernel are built using the set of elementary trees as defined by LTAG.

In (Cumby and Roth 2003), a feature description language was used to extract structured features from the syntactic shallow parse trees associated with named entities. Their experiments on named entity categorization showed that when the description language selects an adequate set of tree fragments the Voted Perceptron algorithm increases its classification accuracy. The explanation was that the complete tree fragment set contains many irrelevant features and may cause overfitting.

In (Zhang, Zhang, and Su 2006), convolution tree kernels for Relation Extraction were applied in a way similar to the one proposed in (Moschitti 2004). The combina-

tion of standard features along with several tree subparts, tailored according to their importance for the task, produced again an improvement on the state-of-the-art.

Such previous work as well as the above sections show that tree kernels can efficiently represent syntactic objects, e. g. constituent parse trees, in huge feature spaces. The next section describes our Semantic Role Labeling system adopting tree kernels within Support Vector Machines.

4. A State-of-the-Art Architecture for Semantic Role Labeling

A meaningful study of tree kernels for SRL cannot be carried out without a comparison with a state-of-the-art architecture: kernel models that improve average performing systems are a just technical exercise whose findings would have a reduced value. A state-of-the-art architecture, instead, can be used as a basic system upon which tree kernels should improve. Since kernel functions in general introduce a sensible slow down with respect to the linear approach, we also have to consider efficiency issues. These aims drove us in choosing the following components for our SRL system:

- Support Vector Machines (SVMs) as our learning algorithm; these provide both a state-of-the-art learning model (in terms of accuracy) and the possibility of using kernel functions;
- a two-stage role labeling module to improve learning and classification efficiency. This comprises:
 - a feature extractor that can represent candidate arguments using both linear and structured features;
 - a boundary classifier (BC);
 - a role multi-classifier (RM), which is obtained by applying the OVA (One vs All) approach;
- a conflict resolution module, i. e. a software component that resolves inconsistencies in the annotations using either a rule based approach or a tree kernel classifier. The latter is interesting to experiment with the classification of complete predicate argument annotations in correct and incorrect structures;
- a joint inference re-ranking module, which employs a combination of standard features and tree kernels to rank alternative candidate labeling schemes for a proposition. This module, as shown in (Gildea and Jurafsky 2002, Pradhan et al. 2004, Haghghi, Toutanova, and Manning 2005), is mandatory in order to achieve state-of-the-art accuracy.

We point out that we did not use any heuristic to filter out the nodes which are likely to be incorrect boundaries, e.g. as done in (Xue and Palmer 2004). On the one hand, this makes the learning and classification phases more complex since they involve more instances. On the other hand, our results are not biased by the quality of the heuristics, leading to more meaningful findings.

In the remainder of this section, we describe the main functional modules of our architecture for SRL and introduce some basic concepts about the use of structured features for SRL. Specific feature engineering for the above SRL subtasks is described and discussed in Section 5.

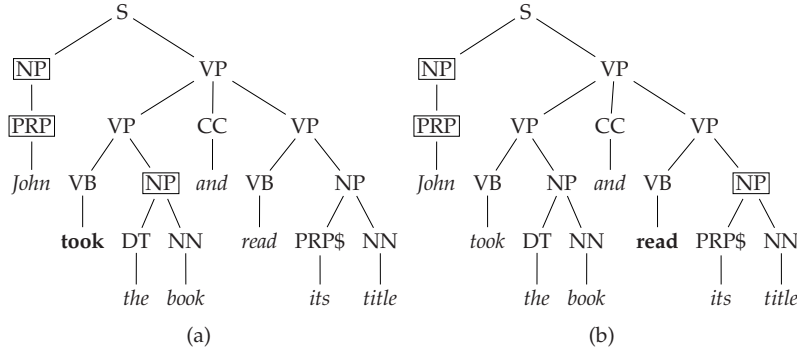


Figure 2 Positive (framed) and negative (unframed) examples of candidate argument nodes for the propositions (a) $[Arg_0 \text{ John}] \text{ took } [Arg_1 \text{ the book}] \text{ and read its title}$ and (b) $[Arg_0 \text{ John}] \text{ took the book and read } [Arg_1 \text{ its title}]$.

4.1 A Basic Two-Stage Role Labeling System

Given a sentence in natural language, our SRL system identifies all the verb predicates and their respective arguments. We divide this step in three subtasks: (a) predicate detection which can be carried out by simple heuristics based on part of speech information, (b) the detection of predicate argument boundaries, i. e. the span of their words in the sentence, and (c) the classification of the argument type, e. g. Arg_0 or Arg_M in PropBank.

The standard approach to learn both the detection and the classification of predicate arguments is summarized by the following steps:

1. Given a sentence from the *training-set*, generate a full syntactic parse-tree;
2. let \mathcal{P} and \mathcal{A} be the set of predicates and the set of parse-tree nodes (i. e. the potential arguments), respectively;
3. for each pair $\langle p, a \rangle \in \mathcal{P} \times \mathcal{A}$:
 - extract the feature representation, $\phi(p, a)$, e. g. attribute-values or tree fragments (see Section 3.1);
 - if the leaves of the subtree rooted in a correspond to all and only the words of one argument of p (i. e. a exactly covers an argument), add $\phi(p, a)$ in E^+ (positive examples), otherwise add it in E^- (negative examples).

For instance, given the example in Figure 2(a) we would consider all the pairs $\langle p, a \rangle$ where p is the node associated with the predicate *took* and a is any other tree node not overlapping with p . If the node a exactly covers the word sequences *John* or *the book*, then $\phi(p, a)$ is added to the set E^+ , otherwise it is added to E^- , as in the case of the node (NN *book*).

The E^+ and E^- sets are used to train the boundary classifier. To train the role multiclassifier, the elements of E^+ can be reorganized as positive $E_{arg_i}^+$ and negative $E_{arg_i}^-$ examples for each role type i . In this way, a binary ONE-vs-ALL classifier for each argument i can be trained. We adopted this solution following (Pradhan et al. 2005a) since it is simple and effective. In the classification phase, given an unseen sentence, all the pairs $\langle p, a \rangle$ are generated and classified by each individual role classifier C_i .

The argument label associated with the maximum among the scores provided by C_i is eventually selected.

The feature extraction function ϕ can be implemented according to different linguistic theories and intuitions. From a technical point of view, we can use ϕ to map $\langle p, a \rangle$ in feature vectors or in structures to be used in a tree kernel function. The next section describes our choices in more detail.

4.2 Linear and Structured Representation

Our feature extractor module and our learning algorithms are designed to cope with both linear and structured features, used for the different stages of the SRL process.

The standard features that we adopted are shown in Table 1. They include:

- the *Phrase Type, Predicate Word, Head Word, Governing Category, Position* and *Voice* defined in (Gildea and Jurafsky 2002);
- the *Partial Path, No Direction Path, Constituent Tree Distance, Head Word POS, First and Last Word/POS, Verb Subcategorization* and *Head Word of the Noun Phrase in the Prepositional Phrase* proposed in (Pradhan et al. 2005a);
- the *Syntactic Frame* designed in (Xue and Palmer 2004).

We indicate with structured features the basic syntactic structures extracted from the sentence-parse tree or their canonical transformation (see Section 3.1). In particular, we focus on the minimal spanning tree that includes the predicate along with all of its arguments.

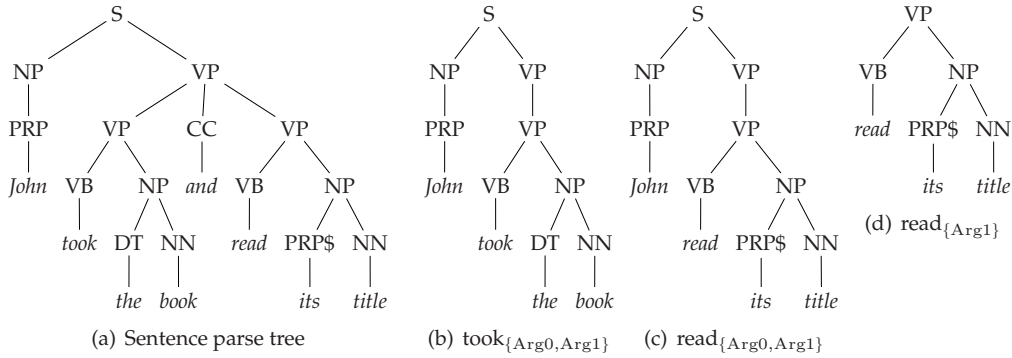
More formally, given a parse tree t , a *node set spanning tree* (NST) over a set of nodes $N_t = \{n_1, \dots, n_k\}$ is a partial tree of t that (1) is rooted at the deepest level and (2) contains all and only the nodes $n_i \in N_t$, along with their ancestors and descendants. An NST can be built as follows: for any choice of N_t , we call r the lowest common ancestor of n_1, \dots, n_k . Then, from the set of all the descendants of r , we remove all the nodes n_j that: (1) do not belong to N_t ; (2) are neither ancestors nor descendants of any node belonging to N_t .

Since predicate arguments are associated with tree nodes, we can define the *predicate argument spanning tree* (AST_n) of a predicate argument node set $A_p = \{a_1, \dots, a_n\}$ as the NST over these nodes and the predicate node, i. e. the node exactly covering the predicate p^2 . An AST_n corresponds to the *minimal* parse subtree whose leaves are all and only the word sequences belonging to the arguments and the predicate. For example, Figure 3(a) shows the parse tree of the sentence: *John took the book and read its title.* $took_{\{ARG_0, ARG_1\}}$ and $read_{\{ARG_0, ARG_1\}}$ are two AST_n structures associated with the two predicates *took* and *read*, respectively, and are shown in Frames (b) and (c) of Figure 3.

For each predicate, only one NST is a valid AST_n . Careful manipulations of an AST_n can be employed for those tasks that require a representation of the whole predicate argument structure, e. g. overlap resolution or proposition re-ranking.

It is worth to note that the *predicate argument feature*, or PAF in (Moschitti 2004) is a canonical transformation of the AST_n in the subtree including the predicate p and only one of its arguments. For the sake of uniform notation, PAF will be referred to as AST_1 (*argument spanning tree*), the subscript 1 stressing the fact that the structure only

² The AST_n of a predicate p and its argument nodes $\{a_1, \dots, a_n\}$, will also be referred to as $p_{\{a_1, \dots, a_n\}}$.

**Figure 3**

A sentence parse tree (a), the correct AST_n s associated with two different predicates (b,c) and a correct AST_1 relative to the argument Arg1 *its title* of the predicate *read*.

encompasses one of the predicate arguments. An example AST_1 is shown in Figure 3(d). Manipulations of an AST_1 structure can lead to interesting tree kernels for local learning tasks, such as boundary detection and argument classification.

Regardless of the adopted feature space, our multiclassification approach suffers from the problem of selecting both boundaries and argument roles independently of the whole structures. Thus, it is possible that (a) two labeled nodes refer to the same arguments (node overlaps) and (b) invalid role sequences are generated, e. g. $\text{Arg0}, \text{Arg0}, \text{Arg0}, \dots$. Next, we describe our approach to solve such problems.

4.3 Conflict Resolution

We call a *conflict*, or *ambiguity*, or *overlap* resolution a stage of the SRL process which resolves annotation conflicts that invalidate the underlying linguistic model. This happens, for example, when both a node and one of its descendants are classified as positive boundaries, i. e. they received a role label. We say that such nodes are *overlapping* as their leaf (i. e. word) sequences overlap. Since this situation is not allowed by the PropBank annotation definition, we need a method to select the most appropriate word sequence. Our system architecture can employ one of three different *disambiguation* strategies:

- a basic solution which, given two overlapping nodes, randomly selects one to be removed;
- the following heuristics:
 1. the node causing the major number of overlaps is removed, e. g. a node which dominates two nodes labeled as arguments;
 2. core arguments (i. e. arguments associated with the subcategorization frame of the target verb) are always preferred over adjuncts (i. e. arguments that are not specific to verbs or verb senses);
 3. in case the above two rules do not eliminate all conflicts, the nodes located deeper in the tree are discarded;
- a tree-kernel based overlap resolution strategy consisting of an SVM trained to recognize non-clashing configurations that often correspond to correct propositions.

The latter approach consists of:

(1) a software module that generates all the possible non-overlapping configurations of nodes. These are built using the output of the local node classifiers by generating all the permutations of argument nodes of a predicate and removing the configurations that contain at least one overlap;

(2) an SVM trained on such non overlapping configurations, where the positive examples are correct predicate argument structures (although eventually not complete) and negative ones are not. At testing time, we classify all the alternative non-clashing configurations. In case more than one structure is selected as correct, we choose the one associated with the highest SVM score.

These disambiguation modules can be invoked after either the BC or the RM classification. The different information available after each phase can be used to design different kinds of features. For example, the knowledge of the candidate role of an argument node can be a key issue in the design of effective conflict resolution methodologies, e. g. by eliminating $\text{Arg}X, \text{Arg}X, \text{Arg}X, \dots$ sequences. Such different approaches are discussed in Section 5.2.

The next section describes a more advanced approach that can eliminate overlaps and choose the most correct annotation for a proposition among a set of alternative labeling schemes.

4.4 A Joint Model for Re-ranking

The heuristics considered in the previous sections only act when a conflict is detected. In a real situation, many incorrect annotations are generated with no overlaps. To deal with such cases, we need a re-ranking module based on a joint BC and RM model as suggested in (Haghighi, Toutanova, and Manning 2005). Such model is based on (1) an algorithm to evaluate the most likely labeling schemes for a given predicate and (2) a re-ranker that sorts the labeling schemes according to their *correctness*.

Step 1 uses the probabilities associated with each possible annotation of parse tree nodes, hence requiring a probabilistic output from BC and RM. As the SVM learning algorithm produces metric values, we applied Platt's algorithm (Platt 1999) to convert them into probabilities, as already proposed in (Pradhan et al. 2005c). These posterior probabilities are then combined to generate the N labelings that maximize a likelihood measure. Step 2 requires the training of an automatic re-ranker. This can be designed using a binary classifier that, given two annotations, decides which one is more accurate. We modeled such classifier by means of three different kernels based on standard features, structured features and their combination.

4.4.1 Evaluation of the N best annotations. First, we converted the output of each node-classifier into a posterior probability conditioned by its output scores (Platt 1999). This method uses a parametric model to fit onto a sigmoid distribution the posterior probability $P(y = 1, f)$, where f is the output of the classifier and the parameters are dynamically adapted to give the best probability output³. Second, we select the N most likely sequences of node labelings. Given a predicate, the likelihood of a labeling scheme

³ We actually implemented the pseudo-code proposed in (Lin, Lin, and Weng 2003) which, with respect to Platt's original formulation, is theoretically demonstrated to converge and avoids some numerical difficulties that may arise.

(or *state*) s for the K candidate argument nodes is given by:

$$p(s) = \prod_{i=1}^K p'_i(l), \quad p'_i(l) = \begin{cases} p_i(l_i)p_i(ARG) & \text{if } l_i \neq \text{NARG} \\ (1 - p_i(ARG))^2 & \text{otherwise} \end{cases} \quad (2)$$

where $p_i(l)$ is the probability of node i being assigned the label l , and $p'_i(l)$ is the same probability weighted by the probability $p_i(ARG)$ of the node being an argument. If $l = \text{NARG}$ (not an argument) then both terms evaluate to $(1 - p_i(ARG))$ and the likelihood of the NARG label assignment is given by $(1 - p_i(ARG))^2$.

To select the N states associated with the highest probability, we cannot evaluate the likelihood of all possible states since they are an exponential number. In order to reduce the search space we (a) limit the number of possible labelings of each node to M and (b) avoid traversing all the states by applying a Viterbi algorithm to search for the most likely labeling schemes. From each state we generate the states in which a candidate argument is assigned different labels. This operation is bound to output at most N states which are generated by traversing a maximum of $N \times M$ states. Therefore, in the worst case scenario the number of traversed states is $V = N \times M \times K$, K being the number of candidate argument nodes in the tree.

During the search we also enforce overlap resolution policies. Indeed, for any given state in which a node n_j is assigned a label $l \neq \text{NARG}$, we generate all and only the states in which all the nodes that are dominated by n_j are assigned the NARG label.

4.4.2 Modeling an Automatic Re-ranker. The Viterbi algorithm generates the N most likely annotations for the proposition associated with a predicate p . These can be used to build annotation pairs, $\langle s_i, s_j \rangle$, which, in turn, are used to train a binary classifier that decides if s_i is more accurate than s_j . Each candidate proposition s_i can be described by a structured feature t_i and a vector of standard features v_i . As a whole, an example e_i is described by the tuple $\langle t_i^1, t_i^2, v_i^1, v_i^2 \rangle$, where t_i^1 and v_i^1 refer to the first candidate annotation, while t_i^2 and v_i^2 refer to the second one. Given such data, we can define the following re-ranking kernels:

$$\begin{aligned} K_{tr}(e_1, e_2) &= K_t(t_1^1, t_1^1) + K_t(t_1^2, t_2^2) - K_t(t_1^1, t_2^2) - K_t(t_1^2, t_2^1) \\ K_{pr}(e_1, e_2) &= K_p(v_1^1, v_2^1) + K_p(v_1^2, v_2^2) - K_p(v_1^1, v_2^2) - K_p(v_1^2, v_2^1), \end{aligned}$$

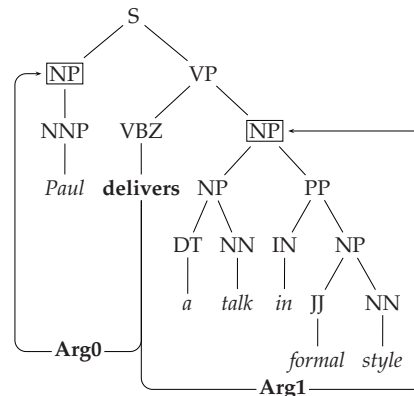
where K_t is one of the tree kernel functions defined in Section 3 and K_p is a polynomial kernel applied to the feature vectors. The final kernel that we use is the following combination:

$$K(e_1, e_2) = \frac{K_{tr}(e_1, e_2)}{|K_{tr}(e_1, e_2)|} + \frac{K_{pr}(e_1, e_2)}{|K_{pr}(e_1, e_2)|}.$$

The above sections have shown how our SRL architecture exploits tree kernel functions to a large extent. In the next section, we describe in more detail our structured features and the engineering methods applied for the different subtasks of the SRL process.

5. Structured Feature Engineering

Structured features are an effective alternative to *standard* features under many aspects. An important advantage is that the target feature space can be utterly changed even

**Figure 4**

Parse tree of the example proposition [Arg0 Paul] delivers [Arg1 a talk in formal style].

by small modifications of the applied kernel function. This can be exploited to identify features relevant to learning problems lacking a clear and sound linguistic or cognitive justification.

As shown in Section 3.1, a kernel function is a scalar product $\phi(o_1) \cdot \phi(o_2)$, where ϕ is a mapping in an Euclidean space, and o_1 and o_2 are the target data, e. g. parse trees. To make the engineering process easier, we decompose ϕ in a *canonical mapping*, ϕ_M , and a *feature extraction* function, ϕ_S , over the set of incoming parse trees. ϕ_M transforms a tree into a canonical structure equivalent to an entire class of input parses while ϕ_S shatters an input tree into its subparts (e. g. subtrees, subset trees or partial trees as described in Section 3). A large number of different feature spaces can thus be explored by suitable combinations $\phi = \phi_S \circ \phi_M$ of mappings.

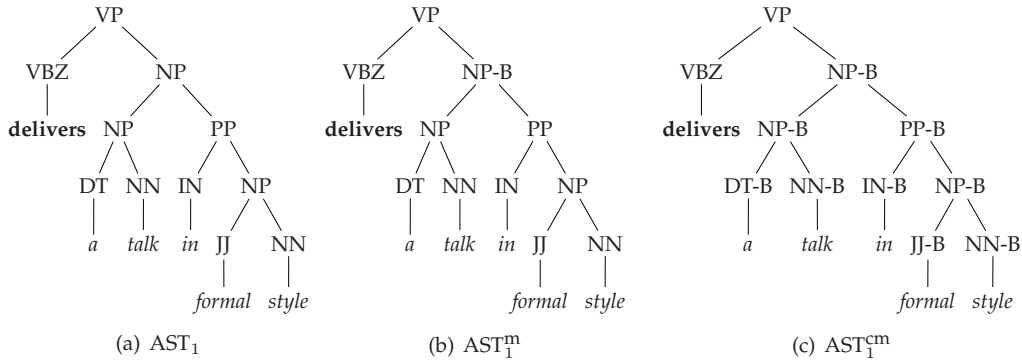
We study different canonical mappings to capture syntactic/semantic aspects useful for SRL. In particular, we define structured features for the different phases of the SRL process, i. e. Boundary Detection, Argument Classification, Conflict Resolution and Proposition Re-ranking.

5.1 Structures for Boundary Detection and Argument Classification

The AST_1 or PAF structures, already mentioned in Section 4.2, have shown to be very effective for argument classification but not for boundary detection. The reason is that two nodes that encode correct and incorrect boundaries may generate very similar AST_1 s and, consequently, have many fragments in common. To solve this problem, we specify the node that exactly covers the target argument node by simply marking it (or marking all its descendants) with the label B , denoting the boundary property.

For example, Figure 4 shows the parse tree of the sentence *Paul delivers a talk in formal style*, highlighting the predicate with its two arguments, i. e. Arg0 and Arg1. Figure 5 shows the AST_1 , AST_1^m and AST_1^{cm} , i. e. the basic structure, the structure with the marked argument node and the completely marked structure, respectively.

To understand the usefulness of node-marking strategies, we can examine Figure 6. This reports the case in which a correct and an incorrect argument node are chosen by also showing the corresponding AST_1 and AST_1^m representations (frames (a) and (b)). Frame (c) shows that the number of common fragments of two AST_1 structures

**Figure 5**

AST₁ (a), AST₁^m (b) and AST₁^{cm} (c) structures relative to the argument Arg1 *a talk in formal style* of the predicate *delivers* of the example parse tree shown in Figure 4.

is 14. This is much larger than the number of common AST₁^m fragments, i. e. only 3 substructures (Frame (d)).

Additionally, since the type of a target argument strongly depends on the type and number of the other predicate arguments⁴ (Toutanova, Haghghi, and Manning 2005, Punyakanok et al. 2005), we should extract features from the whole predicate argument structure. In contrast, AST₁s completely neglect the information (i. e. the tree portions) related to non-target arguments.

One way to use this further information with tree kernels is to use the minimum subtree that spans all the predicate argument structures, i. e. the AST_{*n*} defined in Section 4.2. However, AST_{*n*}s pose two problems:

First, we cannot use them for the boundary detection task since we do not know the predicate argument structure yet. We can derive the AST_{*n*} (its approximation) from the nodes selected by a boundary classifier, i. e. the nodes that correspond to potential arguments. Such approximated AST_{*n*}s can be easily used in the argument classification stage.

Second, an AST_{*n*} is the same for all the arguments in a proposition, thus we need a way to differentiate it for each target argument. Again, we can mark the target argument node as shown in the previous section. We refer to this subtree as a *marked target* AST_{*n*} (AST_{*n*}^{mt}). However, for large arguments (i. e. spread on a large part of the sentence tree) the substructures' likelihood of being part of different arguments is quite high.

To address this problem, we can mark all the nodes that descend from the target argument node. We refer to this structure as a *completely marked target* AST_{*n*} (AST_{*n*}^{cm}). AST_{*n*}^{cm}s may be seen as AST₁s enriched with new information coming from the other arguments (i. e. the non-marked subtrees). Note that if we only consider the AST₁ subtree from a AST_{*n*}^{cm} we obtain AST₁^{cm}.

⁴ This is true at least for core arguments.




Figure 6 AST₁s (a) and AST₁^ms (b) extracted for the same target argument with their respective common fragment spaces (c,b).

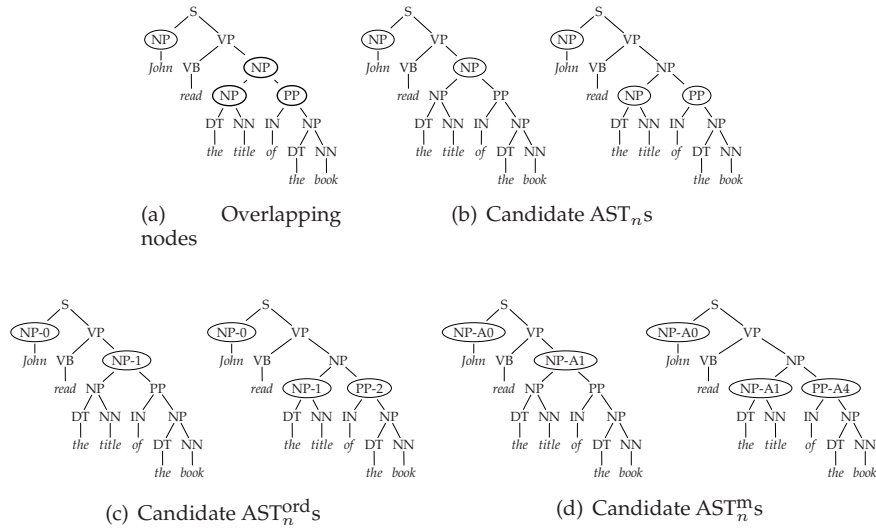
5.2 Structured Features for Conflict Resolution

This section describes structured features employed by the tree kernel based conflict resolution module of the SRL architecture described in Section 4.3. Such subtask is performed by means of:

1. a first annotation of potential arguments using a high recall boundary classifier and, eventually, the role information provided by a role multiclassifier (RM);
2. an AST_n classification step aiming at selecting, among the substructures that do not contain overlaps, those that are more likely to encode the correct argument set.

The set of argument nodes recognized by BC can be associated with a subtree of the corresponding sentence parse, which can be classified using tree kernel functions. These should evaluate whether a subtree encodes a correct predicate argument structure or not. As it encodes features from the whole predicate argument structure, the AST_n that we introduced in Section 4.2 is a structure that can be employed for this task.

Let A_p be the set of potential argument nodes for the predicate p output by BC; the classifier examples are built as follows: (1) we look for node pairs $\langle n_1, n_2 \rangle \in A_p \times A_p$ where n_1 is the ancestor of n_2 or *vice versa*; (2) we create two node sets $A_1 = A - \{n_1\}$ and $A_2 = A - \{n_2\}$ and classify the two NSTs associated with A_1 and A_2 with the tree

**Figure 7**

An overlap situation (a) and the candidate solutions resulting from the employment of the different marking strategies.

kernel classifier to select the most correct set of argument boundaries. This procedure can be generalized to a set of overlapping nodes O with more than 2 elements, as we simply need to generate all and only the permutations of A 's nodes that do not contain overlapping pairs.

Figure 7 shows a working example of such a multi-stage classifier. In (a), the BC labels as potential arguments four nodes (circled), three of which are overlapping (in bold circles). The overlap resolution algorithm proposes two solutions (b) of which only one is correct. In fact, according to the second solution, the propositional phrase *of the book* would incorrectly be attached to the verbal predicate, i. e. in contrast with the parse tree. The AST_n classifier, applied to the two NSTs, should detect this inconsistency and provide the correct output. Figure 7 also highlights a critical problem the AST_n classifier has to deal with: as the two NSTs are perfectly identical, it is not possible to distinguish between them using only their fragments.

The solution to engineer novel features is to simply add the boundary information provided by BC to the NSTs. We mark with a progressive number the phrase type corresponding to an argument node, starting from the leftmost argument. We call the resulting structure an Ordinal Predicate Argument Spanning Tree (AST_n^{ord}). For example, in the first NST of Figure 7(c), we mark as NP-0 and NP-1 the first and second argument nodes, whereas in the second NST, we have a hypothesis of three arguments on three nodes that we transform as NP-0, NP-1 and PP-2.

This simple modification enables the tree kernel to generate features useful to distinguish between two identical parse trees associated with different argument structures. For example, for the first NST the fragments [NP-1 [NP PP]], [NP [DT NN]] and [PP [IN NP]] are generated. They do not match any longer with the fragments of the second NST [NP-0 [NP PP]], [NP-1 [DT NN]] and [PP-2 [IN NP]].

We also experimented with another structure, the *marked predicate argument spanning tree* (AST_n^m), in which each argument node is marked with a role label assigned by a

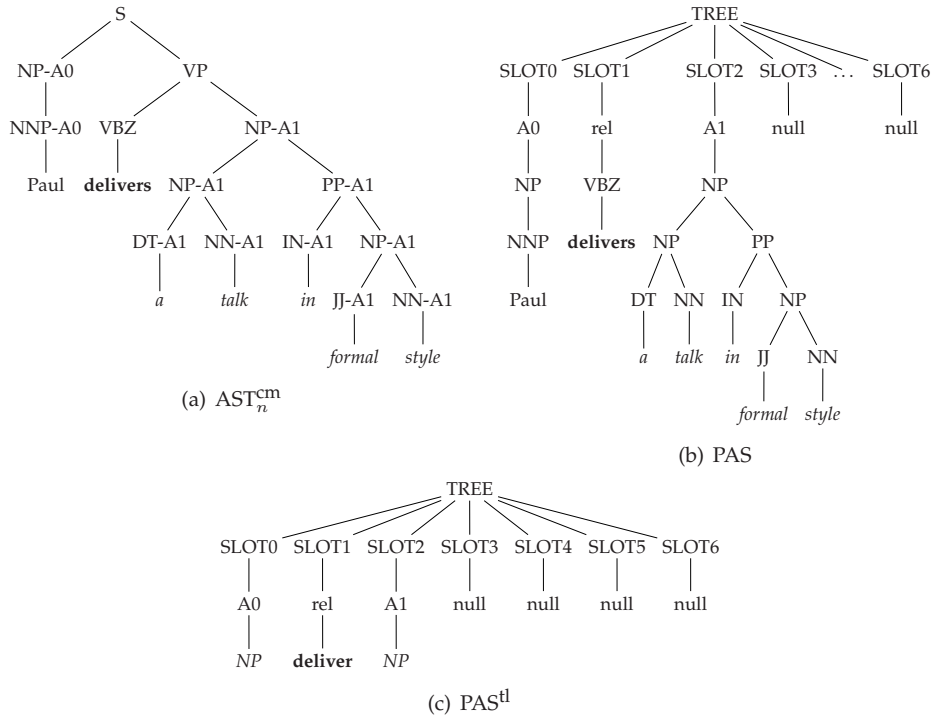


Figure 8
Different representations of a same proposition.

role multi-classifier (RM). Of course, this model requires a RM to classify all the nodes recognized by BC first. An example AST_n^m is shown in Figure 7(d).

5.3 Structures for Proposition Re-ranking

In Section 4.4, we presented our re-ranking mechanism which is inspired by the joint inference model described in (Haghighi, Toutanova, and Manning 2005). Designing structured features for the re-ranking classifier is complex under many aspects. Unlike the other structures that we have discussed so far, the defined mappings should: (1) preserve as much information as possible about the whole predicate argument structure; (2) focus the learning algorithm on the whole structure; (3) be able to identify those small differences that distinguish more or less accurate labeling schemes. Among the possible solutions that we have explored, three are especially interesting in terms of accuracy improvement or linguistic properties, and are described hereinafter.

The AST_n^{cm} (*completely marked AST_n* , see Figure 8(a)) is an AST_n in which each argument node label is enriched with the role assigned to the node by RM. The labels of the descendants of each argument node are modified accordingly, down to pre-terminal nodes. The AST_n^{cmt} is a variant of AST_n^{cm} in which only the target is marked. Marking a node descendant is meant to force substructures matching only among homogeneous argument types. This representation should provide rich syntactic and lexical information about the parse tree encoding the predicate argument structure.

The PAS (*predicate argument structure*, see Figure 8(b)) is a completely different structure that preserves the parse subtrees associated with each argument node while

discarding the intra-argument syntactic parse information. Indeed, the syntactic links between the argument nodes are represented as a dummy 1-level tree, which appears in by any PAS and therefore does not influence the evaluation of similarity between pairs of structures. Such structure accommodates the predicate and all the arguments of an annotation in a sequence of 7 slots⁵. Each slot is attached an argument label which in turn is attached the subtree rooted in the argument node. The predicate is represented by means of a pre-terminal node labeled *rel* to which the lemmatization of the predicate word is attached as a leaf node. In general, a proposition consists of m arguments, with $m \leq 6$, where m varies according to the predicate and the context. To guarantee that predicate structures with a different number of arguments are matched in the SST kernel function, we attach a dummy descendant marked *null* to the slots not filled by an argument.

The PAS^{tl} (*type-only, lemmatized PAS*, see Figure 8(c)) is a specialization of the PAS that only focuses on the syntax of the predicate argument structure, i. e. the type and relative position of each argument, minimizing the amount of lexical and syntactic information derived from the parse tree. The differences with the PAS are that: (1) each slot is attached to a pre-terminal node representing the argument type and a terminal node whose label indicates the syntactic type of the argument; (2) the predicate word is lemmatized.

The next section presents the experiments to evaluate the effectiveness of the proposed canonical structures in SRL.

6. Experiments

The experiments aim to measure the contribution and the effectiveness of our proposed kernel engineering models and of the diverse structured features that we designed (Section 5). In this perspective, the role of feature extraction functions is not fundamental since the study carried out in (Moschitti 2006a) strongly suggests that the SST (Collins and Duffy 2002) kernel produces higher accuracy than the PT kernel when dealing with constituent parse trees, which are adopted in our study⁶. We then selected the SST kernel and designed the following experiments:

- (a) a study of canonical functions based on node marking for boundary detection and argument classification, i. e. AST_1^m , (Section 6.2). Moreover, as the standard features have shown to be effective, we combined them with AST_1^m based kernels on the boundary detection and classification tasks (Section 6.2);
- (b) we varied the amount of training data to demonstrate the higher generalization ability of tree kernels. (Section 6.3);
- (c) given the promising results of kernel engineering, we also applied it to solve a more complex task, i. e. conflict resolution in SRL annotations, (see Section 6.4). As this involves the complete predicate argument structure, we could test advanced canonical functions generating AST_n , AST_n^{ord} and AST_n^m ;
- (d) previous work has shown that re-ranking is very important to boost the accuracy of SRL. Therefore, we tested advanced canonical mappings, i. e. those based on AST_n^{cm} , PAS and PAS^{tl} , on such task (Section 6.5).

⁵ We assume that predicate argument structures cannot be composed by more than 6 arguments, which is generally true.

⁶ Of course the PT kernel may be much more accurate to process PAS and PAS^{tl} since these are not simply constituent parse trees. Nevertheless, a study on the PT kernel potential is beyond the purpose of this article.

Table 2

Number of arguments (Arguments) and of unrecoverable arguments (Unrecoverable) due to parse tree errors in sections 2, 3 and 24 of the Penn TreeBank/PropBank.

Sec.	Arguments	Unrecoverable
2	198,373	454 (0.23%)
3	147,193	347 (0.24%)
24	139,454	731 (0.52%)

6.1 General Setup

The empirical evaluations were mostly carried out within the setting defined in the CoNLL-2005 Shared Task (Carreras and Màrquez 2005). As a target dataset, we used the PropBank⁷ and the automatic Charniak parse trees of the sentences of Penn TreeBank 2 corpus⁸ (Marcus, Santorini, and Marcinkiewicz 1993) from the CoNLL 2005 Shared Task data⁹. We employed the SVM-light-TK software available at <http://ai-nlp.info.uniroma2.it/moschitti/> which encodes fast tree kernel evaluation (Moschitti 2006b) and combinations between multiple feature vectors and trees in the SVM-light software (Joachims 1999). We used the default regularization parameter (option `-c`) and $\lambda = 0.4$ (see (Moschitti 2004)).

6.2 Testing Canonical Functions based on Node Marking

In these experiments, we measured the impact of node marking strategies on boundary detection (BD) and the complete SRL task, i.e. BD and role classification (RC). We employed a configuration of the architecture described in Section 4 and previously adopted in (Moschitti et al. 2005b) in which the simple conflict resolution heuristic is applied. The results were derived within the CoNLL setting by means of the related evaluator.

More in detail, in the BD experiments, we used the first million instances from the Penn TreeBank sections 2-6 for training¹⁰ and Section 24 for testing. Our classification model applied to the above data replicates the results obtained in the CoNLL 2005 shared task, i.e. the highest accuracy in BD among the systems using only one parse tree and one learning algorithm. For the complete SRL task, we used the previous BC and all the available data, i.e. the sections from 2 to 21, for training the role multiclassifier.

It is worth mentioning that, as the automatic parse trees contain errors, some arguments cannot be associated with any covering node thus we cannot extract a tree representation for them. In particular, Table 2 shows the number of arguments (Column 2) for sections 2, 3 and 24 as well as the number of arguments that we could not take into account (Unrecoverable) due to the lack of a parse tree nodes exactly covering their word spans. Note how Section 24 of the Penn TreeBank (which is not part of the Charniak training set) is much more affected by this problem.

⁷ <http://www.cis.upenn.edu/~ace>

⁸ <http://www.cis.upenn.edu/~treebank>

⁹ <http://www.lsi.upc.edu/~srlconll/>

¹⁰ This was the most expensive process in terms of training time, requiring more than one week.

Table 3

Comparison between different models on Boundary Detection and the complete Semantic Role Labeling tasks. The training set is constituted by the first 1 million instances from sections 02-06 for the boundary classifier and all arguments from sections 02-21 for the role multiclassifier (253,129 instances). The performance is measured against section 24 (149,140 instances).

Kernels	Boundary Detection			Semantic Role Labeling		
	P	R	F1	P	R	F1
AST ₁	75.75%	71.68%	73.66	64.71%	61.71%	63.17
AST ₁ ^m	77.32%	74.80%	76.04	66.58%	64.87%	65.71
Poly	82.18%	79.19%	80.66	75.86%	72.60%	73.81
Poly+AST ₁	81.74%	80.71%	81.22	74.23%	73.62%	73.92
Poly+AST ₁ ^m	81.64%	80.73%	81.18	74.36%	73.87%	74.11

Given the above setting, the impact of node marking can be measured by comparing the AST₁ and the AST₁^m based kernels. The results are reported in the rows AST₁ and AST₁^m of Table 3. Columns 2, 3 and 4 show their Precision, Recall and F_1 measure on BD whereas 5, 6 and 7 report the performance on SRL. We note that marking the argument node simplifies the generalization process as it improves both tasks of about 3.5 and 2.5 absolute percent points, respectively.

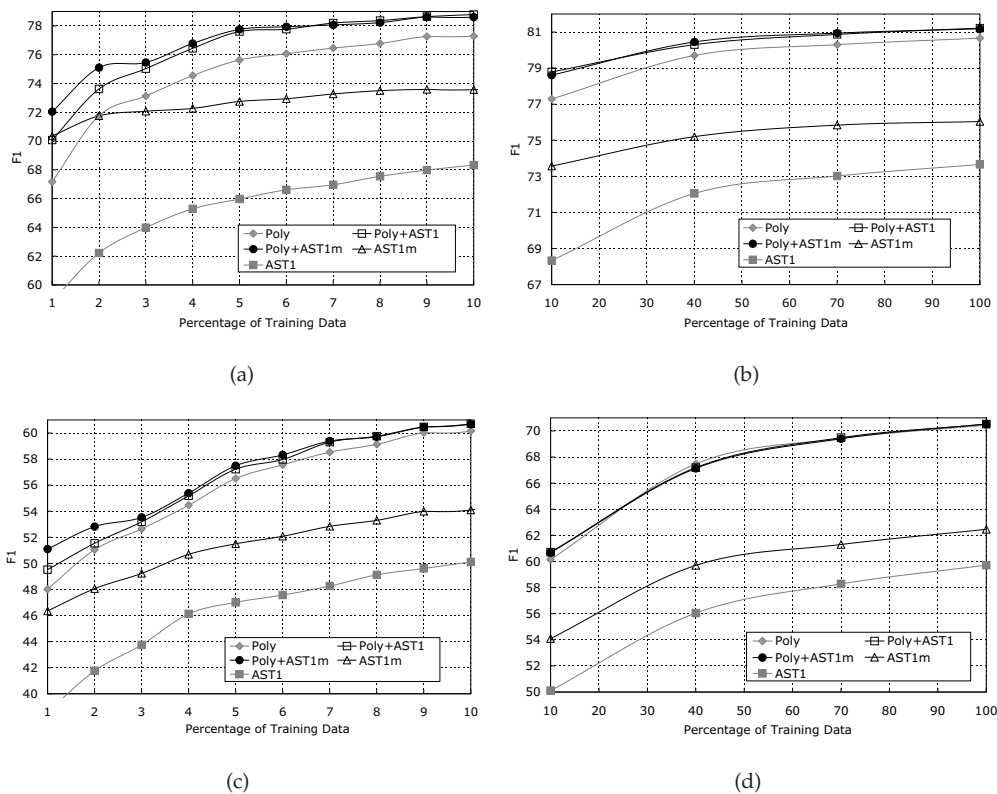
However, Row Poly shows that the polynomial kernel using state-of-the-art features (Moschitti et al. 2005b) outperforms AST₁^m by about 4.5 percent points in BD and 8 points in the SRL task. The main reason is that the employed tree structures do not explicitly encode very important features like the passive voice or predicate position. In (Moschitti 2004), these are shown to be very effective especially when used in polynomial kernels. Of course, it is possible to engineer trees including such and other standard features with a canonical mapping, but the aim here is to provide new interesting representations rather than abiding to the simple exercise of representing already designed features within tree kernel functions. In other words, we follow the idea presented in (Moschitti 2004), where tree kernels were suggested as a means to derive new features rather than generate a stand-alone feature set.

Rows Poly+AST₁ and Poly+AST₁^m investigate this possibility by presenting the combination of polynomial and tree kernels. Unfortunately, the results on both BD and SRL do not show enough improvement to justify the use of tree kernels, e.g. Poly+AST₁^m improves Poly of only 0.52 in BD and 0.3 in SRL. The small improvement is intuitively due to the use of (1) a state-of-the-art model as a baseline and (2) a very large amount of training data which decreases the contribution of tree features. In the next section an analysis in terms of training data will shed some light on the role of tree kernels for BD and RC in SRL.

6.3 The Role of Tree Kernels for Boundary Detection and Argument Classification

The previous section has shown that if a state-of-the-art model¹¹ is adopted, then the tree kernel contribution is marginal. On the contrary, if a non state-of-the-art model is

¹¹ The adopted model is the same used in (Moschitti et al. 2005b) which is the most accurate among the systems that use a single learning model, a single source of syntactic information and no accurate inference mechanism. If tree kernels improved such basic model they would likely improve the accuracy of more complex systems as well.

**Figure 9**

Learning curves for BD in frames (a) and (b) and the SRL task in frames (c) and (d), where 100% of data corresponds to 1 million candidate argument nodes for boundary detection and 64,000 argument nodes for role classification.

adopted tree kernels can play a relevant role. To verify this hypothesis, we tested the polynomial kernel over the standard feature vector proposed in (Gildea and Jurafsky 2002) obtaining an F1 of 67.3, which is comparable with the AST_1^m model, i.e. 65.71. Moreover, a kernel combination produced a relevant improvement of both models reaching an F1 of 70.4.

Thus, the role of tree kernels relates to the design of features for novel linguistic tasks for which the optimal data representation has not been developed yet. For example, although SRL has been studied for many years and many effective features have been designed, representations for languages like Arabic are still not very well understood and raise challenges in the design of effective predicate argument descriptions.

However, the above hypothesis on the usefulness of tree kernels is not completely satisfactory as the huge feature space produced by them should play a more important role in predicate argument representation. For example, the many fragments extracted by an AST_1 provide a very promising back-off model for the Path feature, which should improve the generalization process of SVMs.

As back-off models show their advantages when the amount of training data is small, we experimented with Poly, AST_1 , AST_1^m , Poly+ AST_1 and Poly+ AST_1^m and different bins of training data, starting from a very small set, i.e. 10,000 instances (1%) until 1

Table 4

Boundary detection accuracy (F_1) on gold parse trees and ambiguous structures employing the different conflict resolution methodologies described in Section 4.3.

RND	HEU	AST_n^{ord}
73.13	71.50	91.11

million (100%) of instances. The results on the BD classifiers and the complete SRL task are very interesting and are illustrated by Figure 9. We note that:

First, Graph (a) shows that with only 1% of data, i.e. 640 arguments as positive examples, the F_1 on BD of the AST_1^m kernel is surprisingly about 3 points higher than the one obtained by the polynomial kernel (Poly), i. e. the state-of-the-art. When AST_1^m is combined with Poly the improvement reaches 5 absolute percent points. This suggests that tree kernels should always be used when small training data sets are available.

Second, although the performance of AST_1 is very lower than all the other models, its combination with Poly produces results similar to Poly+ AST_1^m , especially when the size of training data increases. This, in agreement with the back-off property, indicates that the number of tree fragments is more relevant than their quality.

Third, Graph (b) shows that as we increase training data, the advantage of using tree kernels decreases. This is rather intuitive as (i) in general increasing training data allows weaker machine learning/representation models to fill the gap with those more accurate, and (ii) if the hypothesis that tree kernels provide back-off models is true, a lot of training data makes them less critical, e.g. the probability to find the Path feature of a test instance in the training set becomes high.

Finally, graphs (c) and (d) show learning curves¹² similar to graphs (a) and (b), but with a reduced impact of tree kernels on the Poly model. This is due to the reduced impact of AST_1^m on role classification. Such finding is in agreement with the results in (Moschitti 2004), which show that for argument classification the SCF structure (a variant of the AST_n^m), is more effective. Thus a comparison between learning curves of Poly and SCF on RC may show a behavior similar to Poly and AST_1^m for BD.

6.4 Conflict Resolution Results

In these experiments, we are interested in (1) the evaluation of the accuracy of our tree kernel based conflict resolution strategy and (2) studying the most appropriate structured features for the task.

A first evaluation was carried out over gold Penn TreeBank parses and PropBank annotations. We compared the alternative conflict resolution strategies implemented by our architecture (see Section 4.3), namely the random (RND), the heuristic (HEU) and a tree kernel based disambiguator working with AST_n^{ord} structures. The disambiguators were run on the output of BC, i. e. without any information about the candidate arguments roles. BC was trained on sections 2 to 7 with a high-recall linear kernel. We applied it to classify sections 8 to 21 and obtained 2,988 NSTs containing at least an overlapping node. These structures generated 3,624 positive NSTs (i. e. correct structures)

¹² Note that using all training data, all the models reach lower F_1 s than the respective values shown in Table 3. This happens since the data for training the role multiclassifier is restricted to the first million instances, i.e. about 64,000 out of the total 253,129 arguments.

Table 5

SRL accuracy on different PropBank target sections in terms of F_1 measure of the different structured features employed for conflict resolution.

Target section	AST_n	AST_n^{ord}	AST_n^{m}
21	73.7	77.3	78.7
23	68.9	71.2	72.1

and 4,461 negative NSTs (incorrect structures) in which no overlap is present. We used them to train the AST_n^{ord} classifier. The F_1 measure on the boundary detection task was evaluated on the 385 overlapping annotations of Section 23, consisting of 642 argument and 15,408 non-argument nodes.

The outcome of this experiment is summarized in Table 4. We note that: (1) the RND disambiguator (slightly) outperforms the HEU. This suggests that the heuristics that we implemented were inappropriate for solving the problem. It also underlines how difficult it is to explicitly choose the aspects that are relevant for a complex, non local task such as overlap resolution. (2) The AST_n^{ord} classifier outperforms the other strategies by about 20% points, i. e. 91.11 vs 73.13 and 71.50. This datum along with the previous one is a good demonstration of how tree kernels can be effectively exploited to describe phenomena whose relevant features are largely unknown or difficult to be explicitly represented. It should be noted that a more accurate baseline can be evaluated by using the Viterbi-style search (see Section 4.4.1). However, the experiments in Section 6.5 show that the heuristics produce the same accuracy (at least when the complete task is carried out).

The above experiments suggested that tree kernels are promising methods to resolve annotation conflicts, thus, we tried to select the most representative structured features (i. e. AST_n , AST_n^{ord} or AST_n^{m}) also when automatic parse trees are used. We trained BC on sections 2-8, whereas, to achieve a very accurate argument classifier, we trained a role multi-classifier (RM) on sections 2-21. Then, we trained the AST_n , AST_n^{ord} and AST_n^{m} classifiers on the output of BC. To test BC, RM and the tree kernel classifiers, we ran two evaluations on Section 23 and Section 21¹³.

Table 5 shows the F_1 measure of the different tree kernels (Columns 2, 3 and 4) for conflict resolution over the NSTs of sections 21 and 23. Several points should be remarked:

(1) the general performance is by far lower than the one achieved on gold trees, as shown in Table 4. This datum and the gap of about 6% points between sections 21 and 23 confirm the impact of parsing accuracy on the subtasks of the SRL process;

(2) the ordinal numbering of arguments (AST_n^{ord}) and the role type information (AST_n^{m}) provide tree kernels with more meaningful fragments since they improve the basic model of about 4%;

(3) the deeper semantic information generated by the argument labels provides useful clues to select correct predicate argument structures, since the AST_n^{m} model improves AST_n^{ord} performance on both sections.

¹³ As Section 21 of the Penn TreeBank is part of the Charniak parser training set, the performance derived on its parse trees represents an upper bound for our classifiers, i. e. the results using a nearly ideal syntactic parser and role multiclassifier.

Table 6

Number of propositions, alternative annotations (as output by the Viterbi algorithm) and pair comparisons (i. e. re-ranker input examples) for the PropBank sections used for the experiments.

	Section 12	Section 23	Section 24
Propositions	4,899	5,267	3,248
Alternatives	24,494	26,325	16,240
Comparisons	74,650	81,162	48,582

6.5 Proposition Re-ranking Results

In these experiments, Section 23 was used for testing our proposition re-ranking. We employed a BC trained on sections from 2 to 8, whereas RM was trained on sections 2-12¹⁴. In order to provide a probabilistic interpretation of the SVM output (see Section 4.4.1), we evaluated each classifier distribution parameter based on its output on Section 12. For computational complexity reasons, we decided to consider the 5 most likely labelings for each node and the 5 first alternatives output by the Viterbi algorithm (i. e. $M = 5$ and $N = 5$).

With this set up, we evaluated the accuracy lower and upper bounds of our system. As our baseline, we consider the accuracy of a re-ranker that always chooses the first alternative output by the Viterbi algorithm, i. e. the most likely according to the joint inference model. This accuracy has been measured in 75.91 F_1 percent points; this is practically identical to the 75.89 obtained by applying heuristics to remove overlaps generated by BC.

This does not depend on the bad quality of the 5 top labelings. Indeed, we selected the best alternative produced by the Viterbi algorithm according to the gold standard score, and we obtained an F_1 of 84.76 for $N = 5$. Thus, the critical aspect resides in the selection of the best annotations, which should be carried out by an automatic re-ranker.

We learned such classifier with the data described in Table 6. Rows 2 and 3 show the number of distinct propositions and alternative annotations output by the Viterbi algorithm for each of the employed sections. In Row 3, the number of pair comparisons, i. e. the number of training/test examples for the classifier, is shown.

By using such data, we carried out a complete SRL experiment which is summarized in Table 7. First, we compared the accuracy of the AST_n^{cm} , PAS and PAS^{tl} -classifiers trained on Section 24 (in Row 3, columns 2, 3 and 4) and discovered that the latter structure produces a noticeable F_1 improvement, i. e. 78.15 vs 76.47 and 76.77, whereas the accuracy gap between the PAS and the AST_n^{cm} -classifiers is very small, i. e. 76.77 vs 76.47 percent points. We selected the most interesting structured feature, i. e. the PAS^{tl} , and extended it with the local (to each argument node) standard features commonly employed for the boundary detection and argument classification tasks, as in (Haghighi, Toutanova, and Manning 2005). This richer kernel ($PAS^{tl}+STD$, Column 5) was compared with the PAS^{tl} one. The comparison was performed on 2 different training sets (rows 2 and 3): in both cases, the introduction of the standard features produced a performance decrement, most notably in the case of Section 12, i. e. 82.07 vs 75.06. Our best re-ranking kernel, i. e. the PAS^{tl} , was then employed in a larger experiment, using

¹⁴ In these experiments we did not use tree kernels for BC and RM as we wanted to measure the impact of tree kernels only on the re-ranking stage.

Table 7

Summary of the proposition re-ranking experiments with different training sets.

Training Section	AST _n ^{cm}	PAS	PAS ^{tl}	PAS ^{tl} +STD
12	-	-	78.27	77.61
24	76.47	76.77	78.15	77.77
12+24	-	-	78.44	-

both sections 12 and 24 for testing (Row 4), achieving an F_1 measure of 78.44. We note that:

First, the accuracy of the AST_n^{cm} and PAS classifiers is very similar, i. e. 76.77 vs 76.47. This datum suggests that the intra-argument syntactic information is not critical for the re-ranking task, as including it or not in the learning algorithm does not lead to noticeable differences.

Second, the PAS^{tl} kernel is much more effective than those based on AST_n^{cm} and PAS, which are always outperformed. This may be due to the fact that two AST_n^{cm}s (or PASs) always share a large number of substructures, since most alternative annotations tend to be very alike and the small differences among them only affect a small part of the encoding syntactic information; on the other hand, the small amount of local parsing information encoded in the PAS^{tl}s enables a good generalization process;

Finally, the introduction of the standard, local standard features in our re-ranking model caused a performance loss of about 0.5 points on both sections 12 and 24. This fact, which is in contrast with what has been shown in (Haghighi, Toutanova, and Manning 2005), might be the consequence of the small training sets that we employed. Indeed, local standard features tend to be very sparse and their effectiveness should be evaluated against a larger data set.

7. Discussions and Conclusions

The design of automatic systems for the labeling of semantic roles requires the solution of complex problems. Among other issues, feature engineering is made difficult by the structured nature of the data, i. e. features should represent information expressed by automatically generated parse trees. This raises two main problems: (1) the modeling of effective features, partially solved for some subtasks in previous works and (2) the implementation of the software for the extraction of a large number of such features.

A system completely (or largely) based on tree kernels alleviates both problems as (1) kernel functions automatically generate features and (2) only a procedure for the extraction of subtrees is needed. Although, some of the manually designed features seem to be superior to those derived with tree kernels, their combination still seems worth applying. Moreover, tree kernels provide a back-off model that greatly outperforms state-of-the-art SRL models when the amount of training data is small.

To demonstrate the above points, we carried out a comprehensive study on the use of tree kernels for semantic role labeling by designing several canonical mappings. These correspond to the application of innovative tree kernel engineering techniques tailored on different stages of an SRL process. The experiments with such methods and Support Vector Machines on the data set provided by the CoNLL 2005-shared task (Carreras and Màrquez 2005) show that:

First, tree kernels are a valid support to manually designed features for many stages of the SRL process. We have shown that our improved tree kernel, i.e. the one based on

AST_1^m , highly improves accuracy in both boundary detection and the SRL task when the amount of training data is small, e.g. 5 absolute percent points over a state-of-the-art boundary classifier. In the case of argument classification the improvement is less evident but still consistent, about 3%.

Second, appropriately engineered tree kernels can replace standard features in many SRL subtasks. For example, in complex tasks such as conflict resolution or re-ranking, they provide an easy way to build new features that would be difficult to describe explicitly. More generally, tree kernels can be used to combine different sources of information for the design of complex learning models.

Third, in the specific re-ranking task, our structured features show a noticeable improvement over our baseline, i. e. about 2.5 percent points. This could be increased considering that we have not been able to fully exploit the potential of our re-ranking model, whose theoretical upper bound is 6 percent points away. Still, although we only used a small fraction of the available training data (i. e. only 2 sections out of 22 were used to train the re-ranker) our system's accuracy is in line with state-of-the-art systems (Carreras and Màrquez 2005) that do not employ tree kernels.

Finally, although the study carried out in this article is quite comprehensive, several issues should be considered in more depth in the future:

(a) The tree feature extraction functions ST, SST and PT should be studied in combination with the proposed canonical mappings. For example, as the PT kernel seems more suitable for the processing of dependency information, it would be interesting to apply it in an architecture using such kind of syntactic parse trees, e.g. (Chen and Rambow 2003). In particular, the combination of different extraction functions on different syntactic views may lead to very good results.

(b) Once the set of the most promising kernels is established, it would be interesting to use all the available CoNLL 2005 data. This would allow us to estimate the potential of our approach by comparing with previous work on a fairer basis.

(c) The use of fast tree kernels (Moschitti 2006a) along with the proposed tree representations makes the learning and classification much faster, so that the overall running time is comparable with polynomial kernels. However, when used with SVMs their running time on very large datasets (e. g. millions of instances) becomes prohibitive. Exploiting tree kernel derived features in a more efficient way, e.g. by selecting the most relevant fragments and using them in an explicit space, is thus an interesting line of future research. Note that such fragments would be the product of a reverse engineer process useful to derive linguistic insights on semantic role theory.

(d) As CoNLL 2005 (Punyakanok et al. 2005) has shown that multiple parse trees provide the most important boost to the accuracy of SRL systems, we would like to extend our model to work with multiple syntactic views of each input sentence.

Acknowledgments

This article is the result of research on kernel methods for Semantic Role Labeling which started in 2003 and went under the review of several program committees of different scientific communities, from which it highly benefitted. In this perspective, we would like to thank the reviewers of the SRL special issue as well as those of the ACL, CoNLL, EACL, ECAI, ECML, HLT-NAACL and ICML conferences. We are indebted to Silvia Quarteroni for her help in reviewing the English formulation of an earlier version of this article.

References

Baker, Collin F., Charles J. Fillmore, and John B. Lowe. 1998. The Berkeley FrameNet project. In *COLING-ACL '98: Proceedings of the Conference, held at the University of Montréal*, pages 86–90.

- Carreras, Xavier and Lluís Màrquez. 2004. Introduction to the CoNLL-2004 Shared Task: Semantic Role Labeling. In Hwee Tou Ng and Ellen Riloff, editors, *HLT-NAACL 2004 Workshop: Eighth Conference on Computational Natural Language Learning (CoNLL-2004)*, pages 89–97, Boston, Massachusetts, USA, May 6 - May 7.
- Carreras, Xavier and Lluís Màrquez. 2005. Introduction to the CoNLL-2005 Shared Task: Semantic Role Labeling. In *Proceedings of the Ninth Conference on Computational Natural Language Learning (CoNLL-2005)*, pages 152–164, Ann Arbor, Michigan, June.
- Chen, John and Owen Rambow. 2003. Use of Deep Linguistic Features for the Recognition and Labeling of Semantic Arguments. In *Proceedings of the 2003 Conference on Empirical Methods in Natural Language Processing*, pages 41–48.
- Collins, Michael and Nigel Duffy. 2002. New Ranking Algorithms for Parsing and Tagging: Kernels over Discrete structures, and the voted perceptron. In *ACL02*, pages 263–270.
- Culotta, Aron and Jeffrey Sorensen. 2004. Dependency Tree Kernels for Relation Extraction. In *ACL04*, pages 423–429, Barcelona, Spain.
- Cumby, Chad and Dan Roth. 2003. Kernel Methods for Relational Learning. In *Proceedings of ICML 2003*, pages 107–114, Washington, DC, USA.
- Fillmore, Charles J. 1968. The Case for Case. In Emmon Bach and Robert T. Harms, editors, *Universals in Linguistic Theory*. Holt, Rinehart, and Winston, New York, pages 1–210.
- Gildea, Daniel and Daniel Jurafsky. 2002. Automatic Labeling of Semantic Roles. *Computational Linguistics*, 28(3):245–288.
- Haghighi, Aria, Kristina Toutanova, and Christopher Manning. 2005. A Joint Model for Semantic Role Labeling. In *Proceedings of the Ninth Conference on Computational Natural Language Learning (CoNLL-2005)*, pages 173–176, Ann Arbor, Michigan, June.
- Jackendoff, Ray. 1990. *Semantic Structures, Current Studies in Linguistics Series*. Cambridge, Massachusetts: The MIT Press.
- Joachims, Thorsten. 1999. Making large-scale SVM learning practical. In B. Schölkopf, C. Burges, and A. Smola, editors, *Advances in Kernel Methods - Support Vector Learning*, pages 169–184.
- Kazama, Jun'ichi and Kentaro Torisawa. 2005. Speeding up Training with Tree Kernels for Node Relation Labeling. In *Proceedings of EMNLP 2005*, pages 137–144, Toronto, Canada.
- Kudo, Taku and Yuji Matsumoto. 2003. Fast Methods for Kernel-Based Text Analysis. In Erhard Hinrichs and Dan Roth, editors, *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics*, pages 24–31.
- Levin, Beth. 1993. *English Verb Classes and Alternations*. The University of Chicago Press.
- Lin, H.-T., C.-J. Lin, and R.C. Weng. 2003. A Note on Platt's Probabilistic Outputs for Support Vector Machines. Technical report, National Taiwan University.
- Litkowski, Kenneth. 2004. Senseval-3 task: Automatic labeling of semantic roles. In Rada Mihalcea and Phil Edmonds, editors, *Senseval-3: Third International Workshop on the Evaluation of Systems for the Semantic Analysis of Text*, pages 9–12, Barcelona, Spain, July.
- Marcus, M. P., B. Santorini, and M. A. Marcinkiewicz. 1993. Building a Large Annotated Corpus of English: The Penn Treebank. *Computational Linguistics*, 19:313–330.
- Moschitti, Alessandro. 2004. A study on Convolution Kernels for Shallow Semantic Parsing. In *proceedings of the 42th Conference on Association for Computational Linguistic (ACL-2004)*, pages 335–342, Barcelona, Spain.
- Moschitti, Alessandro. 2006a. Efficient Convolution Kernels for Dependency and Constituent Syntactic Trees. In *Proceedings of The 17th European Conference on Machine Learning*, pages 318–329, Berlin, Germany.
- Moschitti, Alessandro. 2006b. Making Tree Kernels Practical for Natural Language Learning. In *Proceedings of 11th Conference of the European Chapter of the Association for Computational Linguistics (EACL2006)*, pages 113–120.
- Moschitti, Alessandro, Bonaventura Coppola, Daniele Pighin, and Roberto Basili. 2005a. Engineering of Syntactic Features for Shallow Semantic Parsing. In *Proceedings of the ACL Workshop on Feature Engineering for Machine Learning in Natural Language Processing*, pages 48–56, Ann Arbor, Michigan, June.
- Moschitti, Alessandro, Ana-Maria Giuglea, Bonaventura Coppola, and Roberto Basili. 2005b. Hierarchical Semantic Role Labeling. In *Proceedings of the Ninth Conference on Computational Natural Language Learning (CoNLL-2005)*, pages 201–204, Ann Arbor, Michigan, June.
- Moschitti, Alessandro, Daniele Pighin, and Roberto Basili. 2006. Tree Kernel Engineering in Semantic Role Labeling Systems. In *Proceedings of the Workshop on Learning Structured Information in Natural Language Applications, EACL 2006*, pages 49–56, Trento, Italy, April.

- European Chapter of the Association for Computational Linguistics.
- Palmer, Martha, Daniel Gildea, and Paul Kingsbury. 2005. The Proposition Bank: an Annotated Corpus of Semantic Roles. *Computational Linguistics*, 31(1):71–106.
- Platt, J. 1999. Probabilistic Outputs for Support Vector Machines and Comparison to Regularized Likelihood Methods. In A.J. Smola, P. Bartlett, B. Schoelkopf, and D. Schuurmans, editors, *Advances in Large Margin Classifiers*, pages 61–74. MIT Press.
- Pradhan, Sameer, Kadri Hacioglu, Valerie Krugler, Wayne Ward, James H. Martin, and Daniel Jurafsky. 2005a. Support Vector Learning for Semantic Argument Classification. *Machine Learning*, 60:1-3:11–39.
- Pradhan, Sameer, Kadri Hacioglu, Wayne Ward, James H. Martin, and Daniel Jurafsky. 2005b. Semantic Role Chunking Combining Complementary Syntactic Views. In *Proceedings of the Ninth Conference on Computational Natural Language Learning (CoNLL-2005)*, pages 217–220, Ann Arbor, Michigan, June.
- Pradhan, Sameer, Wayne Ward, Kadri Hacioglu, James Martin, and Daniel Jurafsky. 2005c. Semantic Role Labeling using Different Syntactic Views. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL'05)*, pages 581–588, Ann Arbor, Michigan, June.
- Pradhan, Sameer S, Wayne H Ward, Kadri Hacioglu, James H Martin, and Dan Jurafsky. 2004. Shallow Semantic Parsing using Support Vector Machines. In Susan Dumais, Daniel Marcu, and Salim Roukos, editors, *HLT-NAACL 2004: Main Proceedings*, pages 233–240, Boston, Massachusetts, USA, May 2 - May 7.
- Punyakankok, Vasin, Peter Koomen, Dan Roth, and Wen-tau Yih. 2005. Generalized Inference with Multiple Semantic Role Labeling Systems. In *Proceedings of the Ninth Conference on Computational Natural Language Learning (CoNLL-2005)*, pages 181–184, Ann Arbor, Michigan, June.
- Shawe-Taylor, John and Nello Cristianini. 2004. *Kernel Methods for Pattern Analysis*. Cambridge University Press.
- Shen, Libin, Anoop Sarkar, and Aravind k. Joshi. 2003. Using LTAG Based Features in Parse Reranking. In *Empirical Methods for Natural Language Processing (EMNLP)*, pages 89–96, Sapporo, Japan.
- Thompson, Cynthia A., Roger Levy, and Christopher Manning. 2003. A Generative Model for Semantic Role Labeling. In *14th European Conference on Machine Learning*, pages 397–408.
- Tjong Kim Sang, Erik, Sander Canisius, Antal van den Bosch, and Toine Bogers. 2005. Applying Spelling Error Correction Techniques for Improving Semantic Role Labelling. In *Proceedings of the Ninth Conference on Computational Natural Language Learning (CoNLL-2005)*, pages 229–232, Ann Arbor, Michigan, June.
- Toutanova, Kristina, Aria Haghighi, and Christopher Manning. 2005. Joint Learning Improves Semantic Role Labeling. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL'05)*, pages 589–596, Ann Arbor, Michigan, June.
- Toutanova, Kristina, Penka Markova, and Christopher Manning. 2004. The Leaf Path Projection View of Parse Trees: Exploring String Kernels for HPSG Parse Selection. In Dekang Lin and Dekai Wu, editors, *Proceedings of EMNLP 2004*, pages 166–173, Barcelona, Spain, July.
- Vapnik, Vladimir N. 1998. *Statistical Learning Theory*. John Wiley and Sons.
- Vishwanathan, S.V.N. and A.J. Smola. 2002. Fast Kernels on Strings and Trees. In *Proceedings of Neural Information Processing Systems*, pages 569–576.
- Xue, Nianwen and Martha Palmer. 2004. Calibrating Features for Semantic Role Labeling. In Dekang Lin and Dekai Wu, editors, *Proceedings of EMNLP 2004*, pages 88–94, Barcelona, Spain, July.
- Zelenko, D., C. Aone, and A. Richardella. 2003. Kernel Methods for Relation Extraction. *Journal of Machine Learning Research*, pages 1083–1106.
- Zhang, Min, Jie Zhang, and Jian Su. 2006. Exploring Syntactic Features for Relation Extraction using a Convolution tree kernel. In *Proceedings of the Human Language Technology Conference of the NAACL, Main Conference*, pages 288–295, New York City, USA, June.

