

On the Identification of Temporal Clauses

Georgiana Puşcaşu*, Patricio Martínez Barco, and Estela Saquete Boró

Department of Software and Computing Systems
University of Alicante, Spain
{georgie, patricio, stela}@dlsi.ua.es

Abstract. This paper describes a machine learning approach to the identification of temporal clauses by disambiguating the subordinating conjunctions used to introduce them. Temporal clauses are regularly marked by subordinators, many of which are ambiguous, being able to introduce clauses of different semantic roles. The paper also describes our work on generating an annotated corpus of sentences embedding clauses introduced by ambiguous subordinators that might have temporal value. Each such clause is annotated as temporal or non-temporal by testing whether it answers the questions *when*, *how often* or *how long* with respect to the action of its superordinate clause. Using this corpus, we then train and evaluate personalised classifiers for each ambiguous subordinator, in order to set apart temporal usages. Several classifiers are evaluated, and the best performing ones achieve an average accuracy of 89.23% across the set of ambiguous connectives.

1 Introduction

Temporality is a key dimension of natural language. Access to temporal information conveyed in text can lead to improvement in the performance of many Natural Language Processing (NLP) applications, such as Question Answering (QA), Automatic Summarisation, Topic Detection and Tracking, as well as any other NLP application involving information about temporally located events.

Natural language conveys temporal information in a wide variety of ways, including tense, aspect, narrative sequence, or expressions carrying it explicitly or implicitly. Any framework that models time and what happens or is obtained in time consists of four fundamental entities: *events*, *states*, *time expressions* and *temporal relations*. An *event* is intuitively something that happens, with a defined beginning and end [22]. *States* pertain in reality and describe conditions that are constant throughout their duration. *Temporal expressions* (TEs) are natural language phrases carrying temporal information on their own. *Temporal relations* hold between two events, between an event and a TE or between two TEs. Temporal relations can be expressed by means of verb tense, aspect, modality, as well as temporal adverbials such as: prepositional phrases (*on Monday*), adverbs of time (*then*, *weekly*) and temporal clauses (*when the war ended*). Temporal clauses, or more specifically the subordinating conjunctions that introduce them, represent an explicit way of expressing temporal relations holding between two events.

The present paper addresses the identification of temporal clauses by disambiguating the cue phrases that may introduce them. Temporal clauses are subordinate clauses

* Currently on research leave from University of Wolverhampton, United Kingdom.

defining the temporal context of the clause they are dependent on. As in the case of other dependent clauses, temporal clauses are regularly marked by cue phrases which indicate the relation between the dependent and main clauses. For the purpose of identifying temporal clauses, a set of cue phrases that normally introduce this type of clauses has been put together. In the following, we will call it the set of temporal subordinators (**STS** = {*after, as, as/so long as, as soon as, before, once, since, until/till, when, whenever, while/whilst*}). The large majority of these cue phrases are ambiguous, being able to introduce clauses showing different semantic roles. Therefore, one can not conclude, by only considering the cue phrase, that the introduced clause is temporal or not. For example, a *since*-clause can either be temporal or causal. The set of ambiguous subordinators (**SAS**) includes *as, as/so long as, since, when, while/whilst*.

This paper will therefore report on an empirical investigation of all temporal connectives, as well as on the design and evaluation of statistical models associated to each ambiguous connective, aiming to set apart the cases when the introduced clauses are temporal. The paper is structured as follows: Section 2 motivates our intentions to recognise temporal clauses and surveys related work, Section 3 explores the grammatical characteristics of temporal clauses and illustrates the ambiguity of the connectives involved in the present study. A machine learning approach to the identification of temporal clauses, as well as the development of a corpus used for training and evaluation are presented in Section 4. Section 5 describes the experiments and the results obtained by implementing and testing this approach on the developed corpus. Finally, in Section 6, conclusions are drawn and future directions of research considered.

2 Motivation and Previous Work

Temporal clauses are an explicit way to express temporal relations between events. But, presently, events are not automatically identifiable according to the existing intuitive definition. Current domain-independent approaches consider as event a text unit, at a coarser-grained scale the sentence, and at a finer-grained scale the clause ([11]). An eventuality is seen as corresponding with an elementary discourse unit (EDU), that is a state of affairs in some spatio-temporal location, involving a set of participants ([2]). Researchers in discourse parsing have proposed different competing hypotheses about what constitutes an EDU. While some take the EDUs to be clauses [5], others see them as sentences [14], prosodic units [6], or intentionally defined discourse segments [4]. Considering the state-of-the-art of current NLP tools, clause splitting is feasible and good performance can be achieved ([11],[16]). We have therefore chosen the clause as elementary unit of discourse and, consequently, as the expression of one event.

Recently, the automatic recognition of temporal and event expressions in natural language has become an active area of research in computational linguistics and semantics. Therefore, a specification language for the representation of events, temporal expressions and temporal links connecting them, TimeML [17], has been developed.

Many research efforts have focused on temporal expression recognition and normalisation ([12], [22], [3], [21], [15]). The importance of the proper treatment of TEs is reflected by the relatively large number of NLP evaluation efforts centered on their identification and normalisation, such as the MUC 6 and 7 Named Entity Recognition

tasks, the ACE-2004 Event Recognition task, the Temporal Expression Recognition and Normalisation (TERN) task.

In what events are concerned, a wealth of previous work ([13], [25], [10]) has explored different knowledge sources to be used in inferring the temporal order of events. Mani and Shiffman ([11]) consider clauses as the surface realisation of events, employ clause splitting to automatically identify events, time-stamp the clauses containing temporal expressions, and finally order them using a machine learning approach. Filatova and Hovy ([3]) obtained 82% accuracy on time-stamping 172 clauses for a single event type. Other efforts in the area of event ordering include determining intra-sentence temporal relations ([9]), as well as inter-sentence temporal relations ([23]). One may therefore conclude that the identification of clauses, and especially of clauses with a temporal value, can play an important role in capturing the temporal dimension of text.

As we have seen so far, many researchers in the field of temporal information extraction start by identifying and normalising TEs, continue with time-stamping the clauses embedding TEs, and then order the events using mainly verb phrase characteristics and the appearance of certain temporal connectives (like *before*, *after*, *since*). Still, there is little in the literature on automatically detecting when a clause introduced by such connectives is temporal or not. Although the treatment of time expressions is an important first step towards the automatic handling of temporal phenomena, much temporal information is not absolute but relative. Temporal clauses, just as temporal expressions, offer an anchoring in time for the events described in the clauses they are subordinate to. Unlike TEs, they require a deeper analysis in order to be able to anchor those events on a timeline, and sometimes, when the temporal clauses serve only to temporally relate an event to another, finding an anchor is not even possible.

Temporal clauses can be used in the task of event ordering. A study of temporal connectives for the purpose of event ordering was presented by Lapata and Lascarides ([9]). The authors collected sentences containing temporal cue phrases, removed the cue phrases, and then trained a model that guessed the removed marker. Some of these cue phrases are temporally ambiguous, but since the authors were only interested in recovering the cue phrase itself, they do not address the disambiguation task. Another related study ([7]) aimed at classifying 61 different discourse connectives into five different classes. One of the classes employed in this study was *temporal* and the statistical model was based on each connective's pattern of cooccurrence with other connectives.

The present effort to identify temporal clauses can aid in marking up text according to TimeML. Among other elements to be used in the annotation of temporal information, TimeML defines *signals* as textual elements that make explicit either the relation holding between two entities, or the modality of an event, or the fact that one verb refers to two or more separate events. Temporal subordinators are included among the signals defined by TimeML. Within the task of automatically annotating *signals*, the classifiers presented in this paper can decide whether or not a certain occurrence of a subordinator is used to temporally relate two events (meaning that it should be annotated with the TimeML SIGNAL tag) or has another usage within the discourse (no SIGNAL tag).

The present study on temporal clauses is part of an on-going investigation for a methodology to provide better treatment to temporal-sensitive questions in the context of QA. It will serve to order events with respect to each other in order to be able to

answer questions like *Did event X happen before event Y?*. We envisage it will also prove useful in the retrieval of non-frequent answers that take the form of temporal clauses within the retrieved passages and that require deeper processing in order to extract the answer expected by the user. Let's suppose that the user asks *When was DaimlerChrysler formed?* and a retrieved paragraph is *DaimlerChrysler was formed when Daimler-Benz and the Chrysler Corporation merged.* If no precise date is associated to the merge of the two companies, then the question *When did Daimler-Benz and the Chrysler Corporation merge?* should be generated and answered to. As it can be noticed, the QA process will become a cyclic one, in order to provide better answers to temporally sensitive questions.

3 Grammatical Overview of Temporal Clauses

An adverbial clause of time relates the time of the situation denoted by the clause to the time of the situation expressed by the determined main clause ([18]). Semantically, temporal clauses may express time position, duration or frequency. Temporal adverbial clauses generally require a subordinator. Most common temporal adverbial clause subordinators are (according to [18]): *after, as, as/so long as, as soon as, before, once, since, until/till, when, whenever, while/whilst.*

Semantic analysis of adverbial clauses is in general complicated by the fact that many subordinators introduce clauses with different meanings, as illustrated below in the case of temporal subordinators:

- * *when* used for time and concession
 - (1) **When** I awoke one morning, I found the house in an uproar. (temporal *when*-clause)
 - (2) She paid **when** she could have entered free. (concessive *when*-clause)
- * *as* used for manner, reason and time
 - (3) The policeman stopped them **as** they were about to enter. (temporal *as*-clause)
 - (4) I went to the bank, **as** I had run out of cash. (reason *as*-clause)
 - (5) She cooks a turkey **as** her mother used to do. (similarity/comparison *as*-clause)
 - (6) **As** he grew older, he was wiser. (proportion *as*-clause)
- * *while/whilst* used for time, concession and contrast
 - (7) He looked after my dog **while** I was on vacation. (temporal *while*-clause)
 - (8) **While** I don't want to make a fuss, I feel I must protest at your interference. (concessive *while*-clause)
 - (9) **While** five minutes ago the place had presented a scene of easy revelry, it was now as somnolent and dull as the day before payday. (contrast *while*-clause)
- * *since* used for reason and time
 - (10) I've been relaxing **since** the children went away on vacation. (temporal *since*-clause)
 - (11) He took his coat, **since** it was raining. (reason *since*-clause)
- * *as long as/so long as* used for conditional and temporal clauses
 - (12) **As long as** Japan has problems with non-performing loans, the economy will not recover robustly. (temporal *as/so long as*-clause)
 - (13) I don't mind which of them wins it **so long as** Ferrari wins. (conditional *as/so long as*-clause)

The subordinators listed above are the ones we will disambiguate using the methodology described in Section 4 and evaluated in Section 5. The remaining temporal subordinators are not disambiguated, as the clauses they introduce always have a temporal value, even if those clauses may also convey other meanings:

- * *after*, apart from time, may indicate cause
(14) **After** Norma spoke, she received a standing ovation.
- * *before* may combine time with purpose, result or condition
(15) Go **before** I call the police!
- * *until/till*, apart from their main temporal meaning, may imply result
(16) She massaged her leg **until** it stopped hurting.
- * *whenever* may combine time with condition, or time with cause and condition, or time with contingency, but it is primarily used to introduce a frequency adverbial or habitual conditions
(17) **Whenever** I read I like to be alone.
- * *once* may imply, apart from time, contingency, condition and reason
(18) My family, **once** they saw the mood I was in, left me completely alone.
- * *as soon as* illustrates the proximity in time of the two situations
(19) **As soon as** I left, I burst out laughing.

4 Methodology

4.1 Creating an Annotated Corpus

For the purpose of identifying temporal clauses, we have used the Susanne Corpus ([20]), a freely available corpus developed at Oxford University consisting of 14,299 clauses. Figure 1 illustrates the distribution of all temporal subordinators in Susanne Corpus, derived by counting all the clauses introduced by each subordinator $t \in \text{STS}$ (for the ambiguous subordinators no distinction was made between temporal/non-temporal usages). All **STS** subordinators account for 859 clauses in the Susanne Corpus.

For each subordinator $s \in \text{SAS}$, we have extracted all the sentences including subordinate clauses initiated by **s** (either **s** was the first word in a clause, or it was preceded only by coordinating conjunctions or modifying adverbs such as *just*, *even*, *especially*). This extraction methodology automatically excludes the cases when subordinators like *since* or *as* occupy the first position in a sentence and play the role of a preposition (*As a detective, I always pay close attention to details.*).

Out of all the levels of annotation embedded in the Susanne Corpus, we have preserved only clause and sentence boundaries. Afterwards, each clause introduced by **s** was annotated with an extra attribute ($\text{TEMPORAL} = \text{"YES"/"NO"}$) showing its temporal nature, at the same time indicating which clause it is subordinate to. The annotation was made by simply testing whether or not the subordinate clause can answer any of the questions *when*, *how often* or *how long* with respect to the action of its superordinate clause.

Due to the fact that there were only 9 occurrences of *as/so long as* in the Susanne Corpus, we have extracted from the Reuters Corpus [19] 50 more sentences including clauses introduced by any of the two connectives. We have then split the selected sentences into clauses and annotated each occurrence of the connective as temporal or non-temporal.

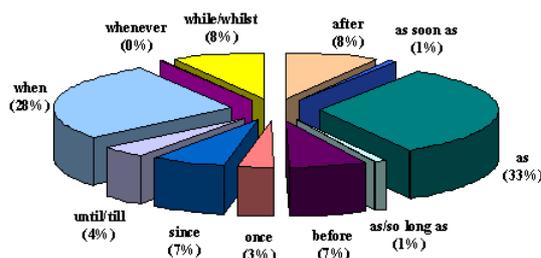


Fig. 1. Distribution of temporal subordinators in Susanne Corpus

The extracted sentences were then parsed using Conexor’s FDG Parser ([24]), with the aim of a realistic evaluation, independent of the manually attached POS-labels present in Susanne Corpus. The Conexor parser gives information on a word’s POS, morphological lemma and it’s functional dependency on surrounding words.

4.2 A Machine Learning Approach

Machine learning has been successfully employed in solving many NLP tasks. There has been much recent interest in machine learning approaches to discourse parsing. One example of employing machine learning in the disambiguation of discourse markers is provided by Hutchinson ([8]). The author aims at acquiring the meaning of a wide set of discourse markers (140) and classifying them along three dimensions: polarity, veridicality and type (i.e. causal, temporal or additive). Though, the temporal class of discourse markers used for training purposes included most subordinators able to introduce temporal clauses, with no attempt being made to set apart their non-temporal usages. Also, the author has excluded from his experiments discourse markers which showed a high degree of ambiguity across classes.

The machine learning method we selected to apply to the problem discussed in this paper is *memory based learning* (MBL). MBL is a supervised inductive learning algorithm for solving classification tasks. It is founded on the hypothesis that the extrapolation of behaviour from stored representations of earlier experience to new situations, based on the similarity of the old and the new situation, is of key importance. The MBL algorithm we used for learning, and then classifying, is k-nearest neighbours. For the purposes of the work described in this paper, the default learning algorithm of the software package called TiMBL [1] was employed (k-nearest neighbours, information gain weighting; k = 1 - due to the reduced quantity of training data available). The evaluation was performed with the *leave-one-out* approach (similar to 10-fold cross-validation), a reliable way of testing the real error of a classifier. The underlying idea is that every instance in turn is selected once as a test item, and the classifier is trained on all remaining instances.

Each training/test instance has been characterized by features pertaining to the classes described in the following section.

4.3 Feature Description

For the purpose of identifying temporal clauses, several classes of features have been designed:

- [I] **Collocation features** encode information (the word and its part of speech) about the words situated within a window of two preceding and two following words with respect to the investigated subordinator. The motivation behind including the surrounding words as features lies in the fact that, many times, a word's meaning can be inferred from its nearby context. The morphological information of the context words is also useful in predicting the usage of a subordinator.
- [II] **Verb features** The verb phrase of the subordinate clause (SubVP) and the verb phrase of the main clause (MainVP) are identified using a set of grammatical rules, and then classified along the following dimensions:
- * MODALITY: **future** (*will, shall, be going to*), **obligation/necessity** (*must, should, have (got) to, ought to, need to, be supposed to*), **permission/possibility/ability** (*can, could, may, might*);
 - * ASPECT: simple, progressive, perfective, perfective progressive;
 - * TENSE: **present, past**;
 - * VOICE: **active, passive**;
 - * POSITIVENESS: **affirmative, negative**
 - * TENSE SIGNATURE: this feature conveys the representation normally used with verb phrases, that combines tense, modality and aspect (for example, it has the value **Future Simple** in the case of future modality and simple aspect, **Present Progressive** in the case of present tense and progressive aspect). It has been introduced to verify if it produces better results than the combination of simple features characterising the verb phrases.
- [III] **Verb connection features** This class includes:
- * MainVP-SubVP: a feature that encodes the tense signatures of the two verb phrases and was included because there are many regularities manifested by the main-subordinate clause pairs corresponding to certain semantic roles (for example in the case of *when*-clauses, the correspondence **Past Tense Simple - Past Tense Simple** signals a temporal use)
 - * SAME LEMMA: a feature indicating whether the two VPs lemmas are identical (this may indicate contrastive, therefore non-temporal, clauses, as in *During school, Sue liked Chemistry while John liked Maths.*)
- [IV] **Cooccurrence features** are used to indicate whether or not, within the span covered by each feature, certain subordinator-specific phrases appear, thus pointing to a certain semantic role. The possible spans covered by these features are the same clause and the main clause span. In the case of *as*, the same clause span feature indicates whether *if* or *though* or *to whether* follow *as*, pointing to a non-temporal usage. The feature corresponding to the main clause span illustrates the presence within this span of:
- * *so, same, as, such*, in the case of *as* (indicating non-temporal usage)
 - * *then, in that case*, for *as/so long as* (indicating non-temporal usage)
 - * *rather, however, therefore, how*, in the case of *since* (indicating non-temporal usage)
 - * *then, always, never, often, usually, every*, in the case of *when* (indicating temporal usage)
 - * *yet, besides, on the other hand, instead, nevertheless, moreover*, in the case of *while/whilst* (indicating non-temporal usage)
- [V] **Structural feature** denotes the position of the subordinate clause with respect to the matrix (before, after or embedded), also indicating the presence/absence of punctuation signs between the two clauses.

[VI] *FDG-based feature* contains information provided by the Conexor FDG-parser that predicts the type of relation holding between the subordinate and matrix clauses. This information is normally attached by the parser to the verb phrase of the subordinate clause.

The classes of features described so far were defined so that their values can be automatically extracted from any text analysed with Conexor.

5 Experiments

In order to assess the impact of the feature classes defined above, we have evaluated several feature combinations using the ML method and settings described in Section 4.2. As baseline for each connective we have considered a classifier that assigns to all instances the class most commonly observed among the annotated examples. Twelve different models have been evaluated in order to compare the relevance of various feature classes to the classification of each temporal connective. The evaluated models are described in detail in the following:

- * **MainVP (Tense Signature only)** This model is trained using only the tense signature of the main clause's verb phrase.
- * **MainVP (All features)** For the main clause's VP, the five characteristics included in the verb features class (modality, aspect, tense, voice, positiveness) have been selected.
- * **SubVP (Tense Signature only)** The model is trained using only the tense signature of the subordinate clause's VP.
- * **SubVP (All features)** The five simple features of the VP corresponding to the subordinate clause are used for training.
- * **BothVP (MainVP + SubVP)** All features characterising the two verb phrases are included in this model.
- * **BestVP** This model designates the best performing VP model observed so far.
- * **VPCombi (BestVP + VPConnection)** The best performing verb phrase model, together with the verb connection features are employed at this stage.
- * **VPCombi + Collocation features** This model comprises the combination of VP features, as well as the features characterising the context of the connective.
- * **VPCombi + Cooccurrence features** This model is trained with the VPCombi model features combined with the cooccurrence features of the corresponding connective.
- * **VPCombi + Structural feature** The VPCombi model together with the structural feature form the present model.
- * **VPCombi + FDG-based feature** This model comprises the VPCombi model features and the FDG-based feature denoting the functional dependency holding between the two clauses.
- * **VPCombi + Best combination** The present model embeds the features of the VPCombi model, as well as the best combination of features chosen from the four feature classes: collocation, cooccurrence, structural and FDG-based.
- * **All** This model is trained with all feature classes described in Section 4.3.

CLASSIFIER	CONNECTIVE				
	AS	AS LONG AS SO LONG AS	SINCE	WHEN	WHILE WHILST
Baseline	67.38%	73.21%	85.00%	86.86%	52.77%
MainVP (Tense Signature only)	74.19%	64.28%	96.66%	84.74%	58.33%
MainVP (All features)	76.70%	64.28%	96.66%	84.32%	47.22%
SubVP (Tense Signature only)	70.25%	78.57%	90.00%	90.67%	75.00%
SubVP (All features)	74.55%	80.35%	96.66%	87.28%	75.00%
BothVP = MainVP + SubVP	81.72%	75.00%	95.00%	91.94%	72.22%
BestVP = MAX(MainVP, SubVP, BothVP)	81.72%	80.35%	96.66%	91.94%	75.00%
VPCombi = BestVP + VPConnection	81.72%	82.14%	95.00%	92.37%	76.38%
VPCombi + Collocation features	86.02%	67.85%	95.00%	89.40%	65.27%
VPCombi + Cooccurrence features	81.72%	82.14%	96.66%	92.79%	81.94%
VPCombi + Structural feature	81.00%	69.64%	96.66%	90.25%	83.33%
VPCombi + FDG-based feature	83.87%	76.78%	95.00%	90.67%	79.16%
VPCombi + Best combination	88.17%	82.14%	98.33%	92.79%	84.72%
All features	86.37%	71.42%	98.33%	91.10%	73.61%

Table 1. Accuracy of various classifiers in discovering temporal usages of ambiguous connectives

All models' accuracy when classifying each connective use as temporal or not is revealed by Table 1. Figures in bold indicate the best performing model per connective.

The best model for *as* (88.17%) includes the grammatical features of the two verb phrases, the verb phrase connection features, the collocation and functional dependency features. The collocation features proved to be useful only in the case of *as*, due to many cases where the connective was preceded by another *as* followed by an adjective or an adverb, signalling non-temporal usage.

In the case of *as/so long as*, the best model (82.14%) comprises the features characterising the subordinate clause VP and the VPConnection.

Since is best dealt with by the VP features of the main clause, combined with VP-Connection, structural and cooccurrence features (98.33%). The verb phrase of the main clause proves to be very important in the classification of *since*, because a temporal *since*-clause generally requires the Present or Past Perfective in the matrix clause.

The best classifier for *when* (92.79%) combines the features corresponding to both verb phrases, VPConnection and cooccurrence.

In the case of *while/whilst*, the best performing model (84.72%) includes the subordinate clause's VP, the VPConnection, the structural and the FDG-based features.

An examination of errors revealed two main causes. On the one hand, there are cases when the syntactic parser fails in identifying verbs, thus leading to erroneous values being attached to the features attached to the verb phrases of the two clauses.

On the other hand, due to the fact that the classifiers do not rely on a semantic analysis of the clauses connected by a certain connective, two syntactically similar pairs of main-subordinate clauses will lead to the same class being assigned to the connective lying between them. This lack of semantic information leads to many classification errors, as exemplified below:

- (20) *As she held her speech, he thought about what they had spoken before.*
(temporal *as*-clause, correctly classified as temporal)
- (21) *As we expected, my uncle recovered fast.*
(non-temporal *as*-clause, but incorrectly classified as temporal)

Bearing in mind that this research is mainly aimed to be included in a temporal-sensitive Question Answering system, and that a previous investigation of temporal questions has revealed the need to identify what questions should be decomposed in order to have more chances of being correctly answered, we have also performed an experiment to distinguish temporal clauses that serve as time anchor (22) from temporal clauses that, apart from referring to time, carry the meaning of habitual condition (23).

- (22) *Who was the ruler of Egypt **when** the World War II started?*
(first submit to a QA system *When did World War II start?*, then substitute in the original question the temporal clause with the answer, *1939*, and finally resubmit the question *Who was the ruler of Egypt in 1939?*)
- (23) *Where can I find a keychain that beeps or chirps **when** I clap my hands?*
(in this case no decomposition is necessary, *when* being synonym to *whenever* or *at any time that*)

The experiment we have performed employed the data annotated for *when*, more precisely the temporal usages of *when*. Each temporal usage was labelled with one of the classes: TIME_ANCHOR or HABITUAL. Afterwards, using the same features as in our previous experiments, we have evaluated several classifiers for distinguishing between the two temporal usages of *when*. The best performing classifier was found to be a combination of VPConnection and cooccurrence features, with an accuracy of 95.12%, the baseline being 91.21% (evaluated on a set of 205 annotated examples).

This classifier can be employed in setting apart habitual conditions, irrespective of the temporal connective used to connect the habitual sequence of events.

6 Conclusion

NLP applications place increasing demand on the processing of temporal information under any form it may appear in text. Temporal clauses are used to establish temporal relations between events, but also to bring into focus a novel temporal referent whose unique identifiability in the reader's memory is presupposed, thus updating the current reference time.

The present paper proposes a machine learning approach to the identification of temporal clauses, by training a classifier for each temporal connective manifesting semantic ambiguity. There is a variation in performance between different subordinators, with the classifiers for *as* and *while/whilst* at 21%, respectively 32%, above the baseline. The average accuracy across all investigated connectives is 89.23%, significantly above the average baseline of 73.04%. We believe that an increased size of the training set could lead to an improved performance. In the case of all connectives, the most informative features have proved to be those derived from the verb phrases of the main and subordinate clauses.

The approach presented in this paper is robust, domain-independent and highly relevant to future work involving temporally ordering events, producing TimeML compliant data, and finally improving temporal-sensitive QA. Future work will also investigate other machine learning algorithms that might prove more suited to the present task, as well as the correlation between the semantic classes of the verbs occurring in the main and subordinate clauses, and the temporal value of the subordinate clause.

References

1. W. Daelemans, J. Zavrel, K. Slood, and A. Bosch. TiMBL Tilburg Memory Based Learner, version 51. Ilk technical report 04-02, 2004.
2. D. Davidson. *The Logical Form of Action Sentences*. University of Pittsburgh Press, 1967.
3. E. Filatova and E. Hovy. Assigning Time-Stamps to Event-Clauses. In *Proceedings of the 2001 ACL Workshop on Temporal and Spatial Information Processing*, 2001.
4. B. Grosz and C. Sidner. Attentions, Intentions, and the Structure of Discourse. In *Computational Linguistics*. 1986.
5. J. Haiman and S. Thompson. *Clause Combining in Grammar and Discourse*. John Benjamins, 1988.
6. J. Hirschberg and D. Litman. Empirical Studies on the Disambiguation of Cue Phrases. In *Computational Linguistics*. 1993.
7. B. Hutchinson. Automatic Classification of Discourse Markers by their Cooccurrences. In *Proceedings of ESSLLI'03 Workshop on The Meaning and Implementation of Discourse Particles*, 2003.
8. B. Hutchinson. Acquiring the Meaning of Discourse Markers. In *Proceedings of ACL 2004*, 2004.
9. M. Lapata and A. Lascarides. Inferring Sentence-Internal Temporal Relations. In *Proceedings of HLT-NAACL 2004*, 2004.
10. A. Lascarides and J. Oberlander. Temporal Connectives in a Discourse Context. In *Proceedings of the European Chapter of the Association for Computational Linguistics*, 1993.
11. I. Mani and B. Shiffman. Temporally Anchoring and Ordering Events in News. In J. Pustejovsky and R. Gaizauskas, editors, *Time and Event Recognition in Natural Language*. John Benjamins, 2004.
12. I. Mani and G. Wilson. Robust temporal processing of news. In *Proceedings of ACL 2000*, 2000.
13. M. Moens and M. Steedman. Temporal Ontology and Temporal Reference. In *Computational Linguistics*. 1988.
14. Livia Polanyi. *A Formal Model of the Structure of Discourse*. Journal of Pragmatics, 1988.
15. G. Puscasu. A Framework for Temporal Resolution. In *Proceedings of the LREC2004*, 2004.
16. G. Puscasu. A Multilingual Method for Clause Splitting. In *Proceedings of the 7th Annual Colloquium for the UK Special Interest Group for Computational Linguistics*, 2004.
17. J. Pustejovsky, R. Sauri, A. Setzer, R. Gaizauskas, and R. Ingria. TimeML Annotation Guidelines Version 1.0. <http://www.cs.brandeis.edu/jamesp/arda/time/>, 2002.
18. R. Quirk, S. Greenbaum, G. Leech, and J. Svartvik. *A Comprehensive Grammar of the English Language*. Longman, 1985.
19. Reuters Corpus. Volume 1, English language. <http://about.reuters.com/>, 2000.
20. Geoffrey Sampson. *English for the computer: the SUSANNE corpus and analytic scheme*. Oxford University Press, 1995.
21. F. Schilder and C. Habel. From Temporal Expressions to Temporal Information: Semantic Tagging of News Messages. In *Proceedings of the 2001 ACL Workshop on Temporal and Spatial Information Processing*, 2001.
22. A. Setzer. *Temporal Information in Newswire Articles: An Annotation Scheme and Corpus Study*. PhD thesis, University of Sheffield, 2001.
23. A. Setzer and R. Gaizauskas. On the Importance of Annotating Event-Event Temporal Relations in Text. In *Proceedings of the LREC Workshop on Temporal Annotation Standards*, 2002.
24. P. Tapanainen and T. Jaervinen. A non-projective dependency parser. In *Proceedings of the 5th Conference of Applied Natural Language Processing, ACL*, 1997.
25. B. Webber. Tense as Discourse Anaphor. *Computational Linguistics*, 14(2):61 – 73, 1988.