

Semantic Frame Parsing for Information Extraction : the CALOR corpus

Gabriel Marzinotto^{1,2}, Jeremy Auguste², Frederic Bechet², Geraldine Damnati¹, Alexis Nasr²

(1) Orange Labs, Lannion, France

(2) Aix Marseille Univ, Université de Toulon, CNRS, LIS, Marseille, France

{gabriel.marzinotto, geraldine.damnati}@orange.com

{jeremy.auguste, frederic.bechet, alexis.nasr}@lis-lab.fr

Abstract

This paper presents a publicly available corpus of French encyclopedic history texts annotated according to the Berkeley FrameNet formalism. The main difference in our approach compared to previous works on semantic parsing with FrameNet is that we are not interested here in *full text parsing* but rather on *partial parsing*. The goal is to select from the FrameNet resources the minimal set of frames that are going to be useful for the applicative framework targeted, in our case Information Extraction from encyclopedic documents. Such an approach leverages the manual annotation of larger corpora than those obtained through full text parsing and therefore opens the door to alternative methods for Frame parsing than those used so far on the FrameNet 1.5 benchmark corpus. The approaches compared in this study rely on an integrated sequence labeling model which jointly optimizes frame identification and semantic role segmentation and identification. The models compared are CRFs and multitasks bi-LSTMs.

Keywords: Frame Semantic Parsing, LSTM, CRF

1. Introduction

Semantic Frame parsing is a Natural Language Understanding task that involves detecting in a sentence an event or a scenario, called *Frame*, as well as all the elements or roles that can be associated to this event in the sentence, called *Frame Elements*. One of the most popular semantic frame model is the Berkeley FrameNet project developed by ICSI Berkeley (Baker et al., 1998). This model is composed of an inventory of Frames with, for each of them, a list of words, called *Lexical Units* (or LU), that can trigger a frame in a sentence. Besides, for each frame, a list of *Frame Elements* (FE), core or optional, is defined. LUs are pairings of a word with a sense; Frame Elements are the components of a frame, represented by sequences of words in a sentence.

Two kinds of parsing can be done with a Semantic Frame model: *full text parsing* where each word in a sentence is analyzed to check if it can trigger a frame; and *partial parsing* where only a subset of frames and LUs is considered, according to their relevance for a given applicative framework. Annotating all the possible LUs and frames in a sentence is a very difficult (and expensive) task for human annotators, therefore there are very few corpora annotated this way. Moreover not many languages have such resources. Most of previous work in semantic frame parsing that has been done with a *full text parsing* approach have used the benchmark corpus FrameNet 1.5. Although the size of this benchmark is relatively large, a lot of frames have a very small number of occurrences in the corpus. This is due to the very large number of frames considered in the semantic model and this makes this corpus particularly challenging for machine learning methods.

On the contrary *partial parsing* can be made on large corpora at a reasonable cost: because the amount of frames and LUs is limited, the annotators can focus only on a few words for each sentence, making the task much easier than full parsing. Corpus obtained this way contain much more examples for each frame, opening the door to more machine learning methods than it is the case with full parsing.

From an applicative point of view, *partial parsing* is also a more realistic option. Although the different senses from the FrameNet model are generic, models trained on the FrameNet 1.5. corpus are not. Assuming that a full text annotation will be available for each new applicative domain is not an option. Moreover a lot of applicative frameworks using semantic models such as frames, like the Information Extraction framework considered in this study, are not interested in full parses but on the contrary only in some specific senses related to the domain targeted.

An example of this kind of annotation scheme is given in figure 1. As we can see, only two words are considered as lexical units in the sentence: *decide*, which triggers the frame *Deciding* and *order* triggering the frame *Request*.

This paper presents a study on the use of Sequence Labeling models such as CRF and LSTM for Semantic Frame parsing. Unlike previous studies developing a multi-step approach involving first the Frame identification task, then the frame element detection and labeling tasks, we propose an approach detecting simultaneously LUs and Frame Elements making Frame identification and argument selection an integrated process. A simple heuristic filter is used in order to maintain coherence in the hypotheses produced. We compare two popular sequence labeling methods: conditional random fields and recurrent deep neural networks using Long Short-term Memory.

2. The CALOR-Frame corpus

We introduce in this paper the corpus CALOR, which is a collection of documents in French language that were hand annotated in frame semantics. This 1.3M words corpus contains documents from 4 different sources: Wikipedia's Archeology portal (WA, 201 documents), Wikipedia's World War 1 portal (WGM, 335 documents), Wikidia's portals of Prehistory and Antiquity (VKH, 183 documents) and ClioTexte's¹ resources about World War One (WW1)

¹<https://clio-texte.clionautes.org/>

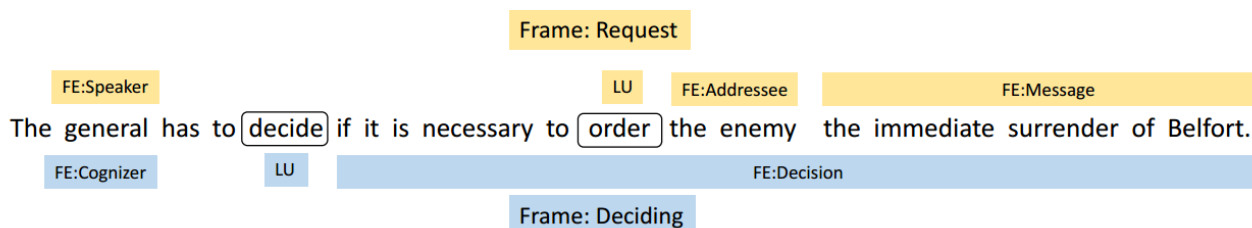


Figure 1: Example of Frame multilabel annotation

(CTGM, 16 documents). These sources were chosen to guarantee both writing style and domain diversity. By having documents from Vikidia (an encyclopedia addressed to children from 8-13 years old) and Wikipedia, presenting subjects of ancient history and archeology we can compare the influence of the writing style on the complexity of the task. The same analysis is possible on the WW1 documents, as Cliotexte (a collection of historical documents such as letters, essays, speeches from WW1) and Wikipedia share a common domain with a completely different writing style. This document selection allows to study the importance of the nature of the training data on the performance of the system on a test set on the same subject. Moreover, having data from two different portals of Wikipedia allows to study the domain dependency problem. For example: evaluate if a model for a frame F , trained on data from the archeology domain can successfully be applied on data from the WW1.

In contrast to full text parsing corpus, the frame semantic annotations of CALOR are limited to a small subset of frames from FrameNet (Baker et al., 1998). As described in the introduction, the goal of this *partial parsing* process is to obtain, at a relatively low cost, a large corpus annotated with frames corresponding to a given applicative context. In our case this applicative context is Information Extraction (IE) from encyclopedic texts, mainly historical texts.

To this purpose we extracted from our 1.3M word text corpus the top 100 most frequent verbs, then we kept those that were more likely to correspond to an action or a situation that would be relevant in our IE context. For example, verbs such as *discover* or *build* are very relevant for exploring archaeological documents. We looked for the corresponding frames of the selected verbs in the Berkeley FrameNet lexicon and kept a set of 53 different frames. If a verb could trigger several frames, we only kept those which were relevant in our corpus.

By adding noun triggers to the list of selected verb triggers we obtain a list of 145 Lexical Units (LU), with 30,950 occurrences in the training corpus. Selecting the most frequent verbs and nouns from the 1.3M words CALOR corpus as frame triggers is a guarantee that the average number of occurrences per frame is high. Therefore, even if the list of frames annotated in the CALOR corpus is small compared to the Framenet corpus, we have a large variety of occurrences for each of them, allowing us to build robust parsers for encyclopedic texts, which is the goal of the CALOR corpus. The list of Frames in CALOR is provided in Table 1.

2.1. The annotation process

Once the corpus, the lexical units and the frame set were chosen, we developed an iterative process for the manual annotation of the CALOR corpus. Preliminarily to this annotation process, the documents were automatically processed by the Macaon (Nasr et al., 2011) tool suite (sentence segmentation, tokenization, POS Tagging, lemmatization, and dependency parsing). Every word within the documents which lemma belongs to the set of selected LUs generates an example to be annotated. It is possible that one sentence generate several examples to annotate if it contains several LUs. We obtained a set of 30,950 examples to annotate, corresponding to all the LUs occurrences in the CALOR corpus. Three annotators were hired for this project. Their goal was to process these 30,950 examples: decide for each of them if its corresponding LU triggers or not one of the 53 frames selected, and finally annotate, if a frame was triggered, all its Frame Elements (FE) occurring in the sentence.

In order to reduce the manual annotation time and perform quality control on the corpus produced, we designed an iterative process based on three principles:

- an automatic pre-annotation scheme based on the frame parser that will be presented in section 4.1.;
- a batch selection process that selects from the unlabeled corpus a set of examples to annotate that corresponds to the same LU in very similar syntactic and lexical contexts.
- an automatic quality control estimator that regularly retrains the frame parser and evaluate its performance thanks to a *k-fold* experiments on the part of the corpus already manually annotated.

The iterative process based on these principles can be implemented as follows:

1. *Frame pre-annotation parsing* : an automatic frame parsing process is applied to each example to annotate. It predicts the frame label and the possible FEs for the LU contained in the example. All these automatic annotations are manually checked and eventually corrected by our annotators. At each iteration the frame parser (see section 4.1.) is trained on the subset of the CALOR corpus that is already annotated. For the first iteration, since there is no data to train the parser, all the LUs are labeled with a "no frame" label. Each iteration brings more data to train the frame parser.

Accomplishment	Activity-start	Age	Appointing	Arrest
Arriving	Assistance	Attack	Awareness	Becoming
Becoming-aware	Buildings	Change-of-leadership	Choosing	Colonization
Coming-to-believe	Coming-up-with	Conduct	Contacting	Creating
Death	Deciding	Departing	Dimension	Education-teaching
Existence	Expressing-publicly	Finish-competition	Giving	Hiding-objects
Hostile-encounter	Hunting	Inclusion	Ingestion	Installing
Killing	Leadership	Locating	Losing	Making-arrangements
Motion	Objective-influence	Origin	Participation	Request
Scrutiny	Seeking	Sending	Shoot-projectiles	Statement
Subjective-influence	Using	Verification		

Table 1: List of Frames annotated in the CALOR corpus

2. *Batch selection process*: the examples that have not yet been manually processed are sorted according to their LUs, then, using a similarity measure taking into account the lexical and syntactic context in which the LUs occur. They are then grouped into batches that will be sent to the annotators for manual validation. the goal here is to reduce the cognitive load of the annotators by splitting the corpus to annotate into small batches sharing very similar properties, likely to be annotated the same way.
3. *Manual correction* : a GUI allows annotators to work on the batches of examples produced in the previous step. Annotation is done on text only, no syntactic annotation is provided to the annotators. The frame pre-annotations produced by the frame parser are displayed, annotators can correct them and add what is missing.
4. *Model training and quality control validation* : at each iteration, the frame parsing models are re-trained on the corpus of manually processed examples. A *k-fold* evaluation is also performed to monitor the evolution of the parsing performance of the model when more validated data is added to the training corpus. If the frame parsing performance improves, it is a good indication that the added data is coherent with the annotations already processed in the previous iterations. This can be seen as a quality control measure of the annotation process on the whole dataset, in addition to inter-annotators agreement measures than can also be estimated on small subsets of the corpus.

2.2. Corpus Statistics

Table 2 presents the distribution of the CALOR corpus among its different sources. We observe that the two first rows of this table, corresponding to Wikipedia, represent most of the corpus. After the annotation process, 30,950 LU have been annotated, leading to 26,725 LU associated to a frame and 4,225 (13%) labeled as *OTHER*. In table 2 the columns **# Sentences**, **# Words**, **# Frames**, **# Other** and **# FE** display the number of sentences, words, frames, LUs, and Frame Elements. **% Sentences with Frame** displays the percentage of sentences with at least one frame and **Lexicon** corresponds to the size of the vocabulary of each document source. The CALOR corpus contains 57,688 FE annotations which averages to 2.2 FE per Frame occurrence.

Figure 2 displays the distribution of the number of annotated examples for each Frame in the CALOR corpus.

We observe that the corpus has a large number of examples per frame: half of the frames in CALOR have more than 400 annotated examples and the 10 most frequent frames have more than 900 examples. The most common Frames are *Attack* (triggers: *attaquer*, *attaque*, *offensive*, *bombardement*, *contre-attaque*), *Leadership* (triggers: *commander*, *diriger*, *commandement*), *Activity Start* (triggers: *commencer*, *débuter*, *commencement*, *début*), *Locating* (triggers: *retrouver*, *trouver*, *localisation*) and *Building* (triggers: *construire*, *fabriquer*, *élever*, *construction*, *fabrication*).

2.3. Comparison with other corpora

The comparison between the CALOR-Frame corpus and other corpora with semantic frame annotations is given in table 3. The column **Documents** shows the main sources of documents annotated, **# Sent** counts the number of sentences in each corpus; **% Sent. w/Frame** shows the percentage of sentences that have a Frame annotation; **Word Lexicon** displays the size of the lexicon of each corpus; **Frame lexicon**, **LU lexicon** and **FE lexicon** correspond to the number of frames, LUs and FEs considered in the annotation model; finally, **# Frame occurrences** shows the number of Frames annotated in the whole corpus.

As we can see, the CALOR-Frame corpus is the only corpus that is not oriented towards journalism, news or current events, it is also the corpus with the largest lexicon size. When we compare it to the existing dataset for Semantic Frame Parsing in English (SemEval07) and French (AS-FALDA) (Candito et al., 2014) we observe that the CALOR corpus has the smallest frame lexicon, but the biggest number of annotations.

3. Frame parsing as a sequence labeling task

As mentioned in the introduction, semantic frame parsing is a structured prediction task where a word can belong to several structures. For example, in figure 1, the word *general* can belong to both the frame *Request* as the frame element *Speaker* and the frame *Deciding* as the *Cognizer*.

In order to consider the Frame parsing task as a *word sequence labeling* task, we need to flatten the frame structures by labeling each word with both semantic and information structure. In this study we have decided to use the simple *B, I, O* encoding for word segments where each word label starts with a *B* if it starts a segment, with an *I* if the word is inside a segment and *O* if it doesn't belong to any segment. Links between segments are represented by word indices.

Document Source	# Sentences	# Words	# Frames	# FE	Lexicon	% Sentence with Frame
WGM (Wikipedia WW1)	30994	686355	14227	32708	42635	34.2%
WA (Wikipedia Archeology)	27023	540653	9943	19892	41418	28.0%
CTGM (Clitexte WW1)	3523	67736	938	1842	10844	21.4%
VKH (Vikidia Prehistory & Antiquity)	5841	85034	1617	3246	11649	21.9%
All	67381	1379778	26725	57688	72127	30.0%

Table 2: Description of the CALOR corpus

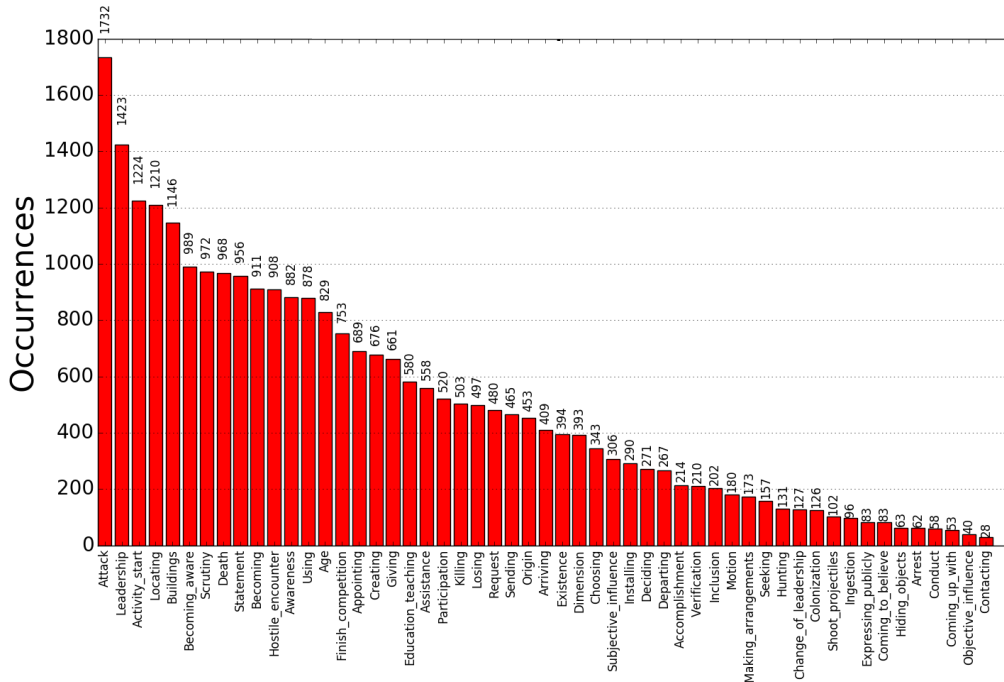


Figure 2: Distribution of the frame occurrences in the CALOR-Frame corpus

An example of such a representation for the sentence of Figure 1 is given in table 4. As we can see column 3 corresponds to the frame *Request* and column 4 to *Deciding*. The LUs triggering the frames are *order* at index 11 for the frame *Request* and *decide* at index 5 for *Deciding*. To each frame element is attached the index of the frame it belongs to through a link to the LU that triggered it.

Using such a representation for sequence labeling models like Conditional Random Fields (CRF) or Recurrent Neural Networks (RNN) models is challenging for two reasons:

1. Multi-labels: each word can receive more than one label according to the number of frames occurring in a sentence. Each label contains the frame as well as the frame element identifiers, therefore there are too many labels to consider building complex labels combining all of them.
2. Linking: an explicit link to the LU triggering a frame is added to each word label of its FEs, as can be seen in table 4. This information is necessary as several expressions of the same frame can occur in a sentence triggered by several LUs. The absolute value of these links is not meaningful and cannot be predicted the same way as the semantic labels are.

In this study we compare two different strategies in order to deal with the multi-label issue, one based on CRF with a multi-model approach (each LU has its own prediction model) and one based on a bi-LSTM model following a multi-task approach. They are described in the next section.

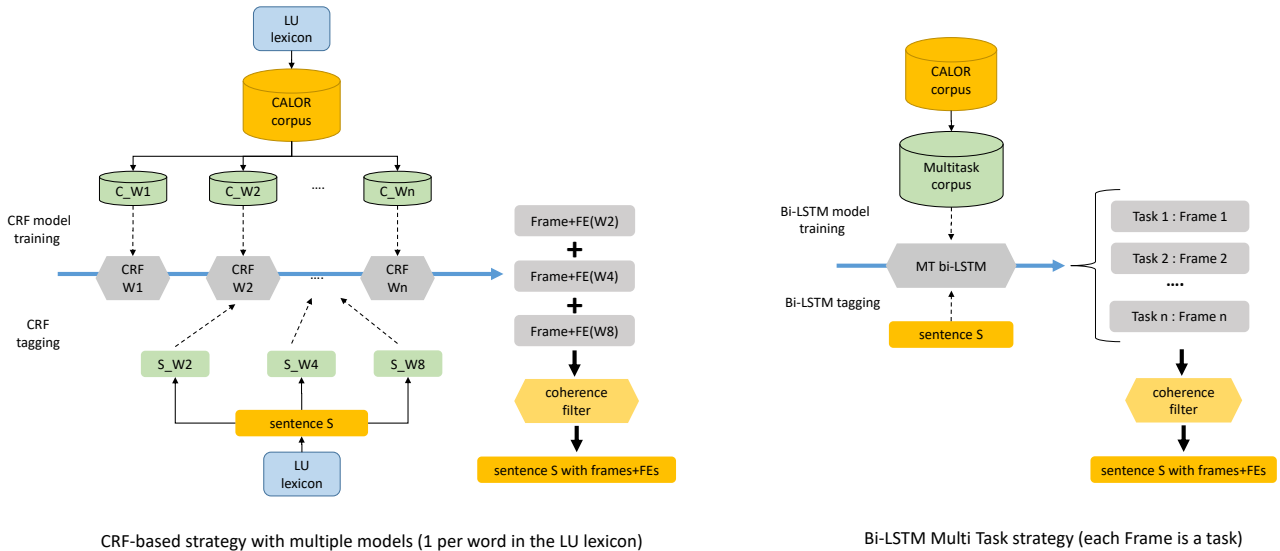
4. Sequence labeling models

4.1. Multi-model CRF approach

CRF-based approaches have been used in many NLP tasks involving sequence labeling such as POS tagging, chunking or named entity recognition (McCallum and Li, 2003). In order to apply CRF to frame parsing, as described in section 3., we need to address the multi-label issue. Since we want to perform frame disambiguation and semantic role detection in one step, and because each word in a sentence cannot trigger more than one frame, we chose a multi-model approach where a CRF-model is trained for each word belonging to the LU lexicon. This approach is described in Figure 3. At training time, the corpus is split according to the LU lexicon: to each word W_i belonging to this lexicon is attached a sub-corpus containing all the sentences C_{W_i} where W_i occurs. For each sentence $s \in C_{W_i}$, W_i can trigger a frame F among all the possible frames for this word in the LU lexicon, or nothing.

Corpus Name	Annotation	Document types	# Sent	% Sent w/ Frame	Word Lexicon	Frame lexicon	LU lexicon	FE lexicon	# Frame occurrences
SemEval07	English Frames	Journals	5946	85.1%	14150	720	3.197	754	23814
ASFALDA	French Frames	Journals (Le Monde)	21634	60.7%	33955	121	782	140	16167
CALOR	French Frames/SRL	Wikipedia Wikidia ClioTexte	67283	33.5%	72127	53	145	148	26725

Table 3: Semantic Frame corpus comparison



CRF-based strategy with multiple models (1 per word in the LU lexicon)

Bi-LSTM Multi Task strategy (each Frame is a task)

Figure 3: Two different strategies (CRF and bi-LSTM multitask) for Frame parsing

1	The	B:Req:Speaker:11	B:Dec:Cogn:5
2	general	I:Req:Speaker:11	I:Dec:Cogn:5
3	has	o	o
4	to	o	o
5	decide	o	LU:Deciding
6	if	o	B:Dec:Decis:5
7	it	o	I:Dec:Decis:5
8	is	o	I:Dec:Decis:5
9	necessary	o	I:Dec:Decis:5
10	to	o	I:Dec:Decis:5
11	order	LU:Request	I:Dec:Decis:5
12	the	B:Req:Address:11	I:Dec:Decis:5
13	enemy	I:Req:Address:11	I:Dec:Decis:5
14	the	B:Req:Message:11	I:Dec:Decis:5
15	immediate	I:Req:Message:11	I:Dec:Decis:5
16	surrender	I:Req:Message:11	I:Dec:Decis:5
17	of	I:Req:Message:11	I:Dec:Decis:5
18	Belfort	I:Req:Message:11	I:Dec:Decis:5

Table 4: Example of corpus with B,I,O format

For example, the sentence shown in table 4 will be duplicated into two sub-corpora, C_{order} with column 3 and C_{decide} with column 4. A CRF model is trained on each C_{w_i} sub-corpus.

At decoding time, when processing a sentence S , the same process is applied: first S is duplicated for each word w_i of S belonging to the LU lexicon. Then the CRF model corresponding to each w_i is applied and the different predictions made by the CRF models are merged.

This approach has the advantage of keeping the number of

possible labels to predict for each CRF relatively small, limited to the frames that can be triggered by the word considered. Therefore the ambiguity is limited and CRFs can be trained efficiently even with a large number of features. However the drawback is that the training data is split across words in the LU lexicon, therefore similarities among LU are not exploited. This situation is acceptable if enough training examples are provided for each LUs, which is the case for the CALOR corpus.

4.2. Multi-task LSTM approach

Deep Neural Networks (DNN) with word embedding is the state of the art approach for semantic frame parsing (Hermann et al., 2014). More recently *recurrent neural networks* (RNN) with Long Short Memory (LSTM) cells have been applied to several semantic tagging tasks such as *slot filling* (Mesnil et al., 2015) or even frame parsing (Hakkani-Tür et al., 2016; Tafforeau et al., 2016) for Spoken Language Understanding.

Following these previous works, we propose in this study a single-layered bidirectional LSTM sequence to sequence architecture to perform frame tagging. To deal with the multi-label issue we could train a biLSTM model per LU, using the same approach as for the CRF, however, the number of examples per LU is reduced and neural networks do not perform well on small datasets. We would face the same problem if instead we decided to train one biLSTM per

frame. The remaining possibility is to build a single biLSTM to predict all possible frames in order to automatically learn a feature representation meaningful to all frames. We chose a multi-task approach (biLSTM-MT) similar to the one proposed by (Tafforeau et al., 2016), which models each frame as an isolated task. In this model two LSTM models, one forward and one backward are concatenated and shared among all tasks. Then a task-specific fully-connected output layer is added for each task.

In this work we consider each frame of our FrameNet model as a different task. This approach is described in figure 3. At decoding time each sentence is processed by the network and a distribution probability on the labels of each task for each word is produced. We only keep the labels above a certain threshold.

4.3. Coherence filter

Once the sequence tagging process is performed, each word is labeled with a frame element and position labels or a null label (O) as presented in table 4. Because the labels given at the word level might not be coherent at the frame level, we apply a *coherence filter* to the output of the tagging process. This filter is in charge of removing incoherences (FEs not starting with a B label; FEs without a frame) and linking the FEs to the LU that triggered the corresponding frame.

This filter implement is a very simple strategy: in a given sentence, if a word W of index i is labeled as a LU trigger for frame F , we link all the FEs detected in the sentence with the same frame label F to the LU w_i . At the end of this process, all FEs that have not been linked to a LU are removed.

4.4. Feature selection

The feature sets used for the CRF and the biLSTM approaches differ. For the CRF model, each training sample contains only one trigger word which is clearly identified, therefore we can use this information in order to add global constraints on the feature set of each word to process. On the opposite the multi-task biLSTM models can process several triggers in the same sentence, therefore the feature set cannot be biased toward a specific trigger and only local features are considered.

For the CRF models we consider 3 features: word lemma, part-of-speech (POS) and the syntax dependency path between the word to process and the potential frame trigger in the sentence. This dependency path is built through the concatenation of the syntactic functions between the word and the trigger. In the general case, a trigger is not necessarily at the root of the syntactic tree, for this reason, the dependency paths are composed of both links from child to parent (ascending links) and from parent to child (descending links) we make distinction of both types of links in the way we encode the dependency path.

For the biLSTM-MT we consider 4 features: word embeddings (Glove embeddings of 200 dimensions trained on French Wikipedia), POS, syntactic function (without a link) and a boolean indicator of whether the word belongs to the lexicon LU or not. All the features are encoded as trainable embeddings, we allow the network to adapt them during the training of the frame parsing task.

For both systems, in order to extract lemmas, POS and syntactic dependency trees we processed the frame annotated corpus using MACAON (Nasr et al., 2010) trained with a set of POS and dependencies similar to the one proposed in the Paris French TreeBank (Abeillé et al., 2003; Abeillé and Barrier, 2004).

The main differences between the multi-model CRF and multi-task biLSTM models are summarized in table 5.

5. Evaluation

5.1. Experimental setup

When a sentence is processed, there are 4 steps or sub-tasks that take place either explicitly or implicitly in the frame parsing process. Even though our approaches perform frame detection and semantic role labeling in one step, we detail the scores for each sub-task because they are relevant indicators of the performance of a parser, and they serve as point of comparison between our models. These 4 sub-task are:

1. *trigger identification* (TI) which decides whether a word in a sentence can trigger a frame or not;
2. *trigger classification* (TC) which assigns the frame label to the trigger word detected to form a LU;
3. *role filler identification* (RI) which detects potential semantic role fillers in the sentence for the frame detected;
4. *role filler classification* (RC) which assigns a label to each role filler detected in order to obtain the *Frame Elements* (FE) of the frame detected.

In this study we consider these 4 tasks as a cascade process: an error in task 1 will lead to several errors in task 4 since no correct FEs will be detected if the frame is not triggered. The CALOR corpus is annotated with a small number of frames compared to the FrameNet 1.5 corpus, therefore the ambiguity for subtasks 1 and 2 is rather small. The most complex task is of course task 4, *role filler classification*, since every detection and classification of LUs and FEs has to be correct. That is why we will pay a special attention to this task to compare the models in the following experiments.

To carry out our experiments we divide the corpus CALOR assigning 80% of the frame occurrences to the train set and 20% to the test set. This split is done in such a way that the frame distribution remains as similar as possible between train and test while considering a document as an atomic unit that cannot be subdivided and should be either in the train or in the test set. This split does not take into account the LU distribution, a LU can therefore appear both in train and test, only in train or only in test.

Similarly to previous work on semantic frame parsing we will use *precision*, *recall* and *f-measure* on the 4 subtasks presented in order to evaluate our approaches. We also compute *precision/recall* curves by using different acceptance thresholds on the frame and semantic role hypotheses output by our models. In this study we set the operating point for comparing our models to the *Equal Error Rate* (EER) between the precision and recall measures.

	CRF-LU	biLSTM-MT
Lexical Feat	Lemma	Glove Word Embeddings
Morpho-Syntactic Feat	POS	POS Embeddings
Syntactic Feat	Dep. Path to Trigger	Syntactic Function Embeddings
Model	145 CRF (1 CRF per LU)	1 biLSTM 54 Taks (1 Task per Frame)
Scaling	Many LU	Many Taks
Labeling Process	1 Application per LU per Phrase	Parse all the Frames in 1 Application
Information Sharing	No Sharing	Implicit Sharing Between Frames

Table 5: Comparative Overview of CRF-LU and biLSTM-MT models

5.2. Overall results

The comparison between the multi-model CRF (CRF-MM) and multi-task biLSTM (biLSTM-MT) performance over the set of 4 sub-tasks presented in the previous section is presented in Table 6. Both systems use their full feature set as presented in table 5 with EER between precision and recall as the chosen operating point. As expected performance for subtasks 1 and 2 are very good, much better than those reported on the FrameNet 1.5 corpus (Hermann et al., 2014). This is due to the *partial parsing* approach used here where only a subset of the FrameNet model is used.

CRF-MM model performs better on the trigger identification and classification tasks (subtasks 1 and 2) while the biLSTM-MT model is better on role identification and classification (subtasks 3 and 4). The reason for this behavior is that it is easier for a CRF model to identify the proper frame for each trigger word, as they can only trigger a few frames, while for the biLSTM-MT all frames are in competition. This reduction of ambiguity for CRF-MM models has a cost: splitting the training corpus according to the LU lexicon and therefore reducing the training data for each model. If this is not an issue for the two low ambiguity subtasks 1 and 2, the situation is different for subtasks 3 and 4. Indeed biLSTM-MT is a single model that learns from all training data, it is able to automatically learn relevant features and capture semantic aspects of text using a shared layer and then use these features to classify tokens in roles of different frames learned one by tasks. This ability to use the whole corpus leads the biLSTM-MT model to outperform CRF-MM for tasks 3 and 4.

This result is confirmed by figure 4 which displays the precision/recall curves of both methods on subtask 4. As we can see, biLSTM-MT is better in terms of maximal F-measure. It is interesting to observe that both models do not show the same precision-recall trade-off. The CRF-MM is able to parse text at a high precision with a low recall, this is not the case of biLSTM-MT, that handles a much larger set of frames and is more prone to precision errors. On the other hand, biLSTM-MT achieves a better recall at a fairly good precision, this happens because it is trained on the full dataset and it is able to learn syntactic patterns from different frames and extrapolate them from one frame to the other.

6. Conclusion

In this paper we introduced the CALOR corpus as well as a comparison between two different semantic frame parsing

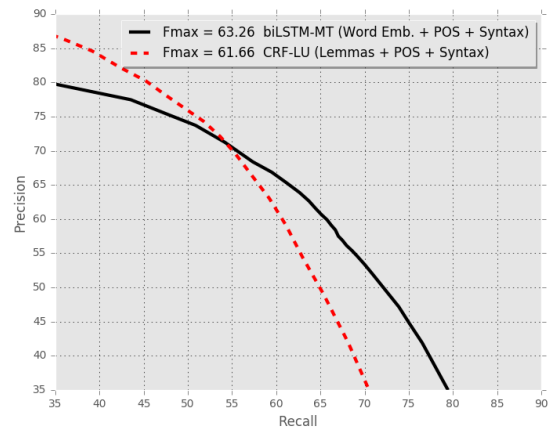


Figure 4: Precision Recall curves for CRF-LU and biLSTM-MT on the full feature set

	Fmeasure CRF- MM	Fmeasure biLSTM- MT
Trigger Ident. (TI)	96.4	95.5
Trigger Classif. (TC)	95.2	92.3
Role Ident. (RI)	69.7	70.3
Role Classif. (RC)	60.6	63.2

Table 6: Comparison of our two models on their full feature set operating at EER

models that consider the task as a sequence labeling task. The main contribution of this work is to propose a new corpus, publicly available, that contain semantic frame annotations that differ from previous corpus such as FrameNet and SemEval. In our case only *partial* annotation is considered, allowing to annotate much larger corpora at a lower cost than full text annotation. Only a small subset of the FrameNet lexicon is used, however the amount of data annotated for each frame is much larger than in other corpora, allowing to develop and test different parsing methods.

In this study two frame parsing models are compared, one based on CRF following a multi-model approach and other based on biLSTM with a multi-task approach. Experiments show that biLSTM-MT model achieves a better recall, while CRF-MM achieves better precision, this is due to the architecture of each model, in CRF-MM we divide frame parsing into small subtasks one per LU, reducing the

number of possible labels in each decision, thus augmenting the precision. On the other hand biLSTM-MT is able to share data across LUs boosting its capacity to deal with complex syntactic patterns and being able to retrieve more frame elements during parsing.

7. Bibliographical References

- Abeillé, A. and Barrier, N. (2004). Enriching a french treebank. In *LREC*.
- Abeillé, A., Clément, L., and Toussnel, F. (2003). Building a treebank for french. *Treebanks*, pages 165–187.
- Baker, C. F., Fillmore, C. J., and Lowe, J. B. (1998). The berkeley framenet project. In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics - Volume 1, ACL '98*, pages 86–90, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Candito, M., Amsili, P., Barque, L., Benamara, F., Chalendar, G., Djemaa, M., Haas, P., Huyghe, R., Mathieu, Y., Muller, P., Sagot, B., and Vieu, L. (2014). Developing a french framenet: Methodology and first results. *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC)*.
- Hakkani-Tür, D., Tur, G., Celikyilmaz, A., Chen, Y.-N., Gao, J., Deng, L., and Wang, Y.-Y. (2016). Multi-domain joint semantic frame parsing using bi-directional rnn-lstm. In *Proceedings of The 17th Annual Meeting of the International Speech Communication Association*.
- Hermann, K. M., Das, D., Weston, J., and Ganchev, K. (2014). Semantic frame identification with distributed word representations. In *ACL (1)*, pages 1448–1458.
- McCallum, A. and Li, W. (2003). Early results for named entity recognition with conditional random fields, feature induction and web-enhanced lexicons. In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003 - Volume 4, CONLL '03*, pages 188–191, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Mesnil, G., Dauphin, Y., Yao, K., Bengio, Y., Deng, L., Hakkani-Tur, D., He, X., Heck, L., Tur, G., Yu, D., et al. (2015). Using recurrent neural networks for slot filling in spoken language understanding. *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)*, 23(3):530–539.
- Nasr, A., Bechet, F., and Rey, J.-F. (2010). Macaon : Une chaîne linguistique pour le traitement de graphes de mots. In *Traitement Automatique des Langues Naturelles - session de démonstrations*, Montreal.
- Nasr, A., Béchet, F., Rey, J.-F., Favre, B., and Le Roux, J. (2011). Macaon: An nlp tool suite for processing word lattices. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: Systems Demonstrations, HLT '11*, pages 86–91, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Tafforeau, J., Bechet, F., Artiere, T., and Favre, B. (2016). Joint syntactic and semantic analysis with a multitask deep learning framework for spoken language understanding. *Interspeech 2016*, pages 3260–3264.