

A Pilot Arabic Propbank

Mona Diab, Aous Mansouri[†], Martha Palmer[†], Olga Babko-Malaya,
Wajdi Zaghouni[‡], Ann Bies[‡], Mohammed Maamouri[‡]

Columbia University,[†]University of Colorado, BAE Systems, [‡]Linguistic Data Consortium,

mdiab@ccls.columbia.edu, {aous.mansouri,martha.palmer}@colorado.edu,
olga.babko-malaya@baesystems.com,{wajdz,bies,maamouri}@ldc.upenn.edu

Abstract

In this paper, we present the details of creating a pilot Arabic proposition bank (Propbank). Propbanks exist for both English and Chinese. However the morphological and syntactic expression of linguistic phenomena in Arabic yield a very different type of process in creating an Arabic propbank. Hence, we highlight those characteristics of Arabic that make creating a propbank for the language a different challenge compared to the creation of an English Propbank. We believe that many of the lessons learned in dealing with Arabic could generalise to other languages that exhibit equally rich morphology and relatively free word order.

1. Introduction

Recent years have witnessed a surge in available automated resources for the Arabic language.¹ The computational linguistics community is just about starting to exploit these resources toward several interesting scientific and engineering goals. The Arabic language is interesting from a computational linguistic perspective as it is significantly different from English hence creating a challenge for existing technology to be easily portable to Arabic. The Arabic language is inherently complex due to its rich morphology and relative free word order. Moreover, with the existence of several interesting varieties, namely, the spoken vernaculars, we are witnessing the emergence of written dialectal Arabic everyday on the web, however they exist with no standard orthographies or forms.

We have seen many successful strides towards functional systems for Arabic enabling technologies, but we are yet to read about large Arabic NLP applications such as Machine Translation and Information Extraction that are on par in their performance on the English language. The problem is not the existence of data, but rather the existence of data annotated with the relevant level of information that is useful for NLP. Shallow approaches to text processing have been garnering a lot of attention recently. Specifically, shallow approaches to semantic processing are making large advances in the direction of efficiently and effectively deriving tacit semantic information from text (Pradhan et al., 2003; Gildea and Palmer, 2002; Pradhan et al., 2004; Gildea and Jurafsky, 2002; Xue and Palmer, 2004; Chen and Rambow, 2003; Carreras and Màrquez, 2005; Moschitti, 2004; Moschitti et al., 2005; Diab et al., 2007). With the advent of larger and faster computers and better machine learning algorithms the importance of large scale data resources is becoming paramount.

The existence of semantically annotated resources in English such as FrameNet (Baker et al., 1998) and PropBank (Kingsbury and Palmer, 2003; Palmer et al.,

2005) corpora have marked a surge in efficient approaches to automatic semantic labeling of the English language. There is a widely held belief in the NLP and computational linguistics communities that identifying and defining the roles of the arguments of predicates in a sentence has a lot of potential for and is a significant step toward improving important applications such as document retrieval, machine translation, question answering and information extraction. Hence, the process by which predicates and their arguments are identified and roles defined in a sentence would lead to a better performance in such applications.

In the English sentence, ‘John likes apples.’, the predicate is ‘likes’ and the first argument, the subject, is ‘John’, and the second argument, the object, is ‘apples’. ‘John’ would be semantically annotated as the *agent* and ‘apples’ would be the *theme*. Depending on the linguistic theory adopted, the labels for these roles vary. For instance, according to FrameNet, ‘John’ is labeled the *liker* while ‘apples’ is labeled *likee*. According to Propbank, ‘John’ is labeled *ARG0* and ‘apples’ is labeled *ARG1*. It is worth noting that independent of the labeling formalism adopted, the labels do not vary in different syntactic constructions, which is why proposition annotation is different from treebank annotation. For instance, if the example above was in the passive voice, ‘Apples are liked by John’, ‘apples’ is still the *theme*, *ARG1* and *likee*, while ‘John’ remains the *agent*, *ARG0* and *liker*. Likewise for the example ‘John opened the door.’ vs. ‘The door opened’, in both these examples ‘the door’ is the *theme* or according to Propbank annotation, it is labeled *ARG1*.

Systems follow the resources. To date most of the systems exist for English. In this paper, we present recent work on producing a pilot Arabic proposition bank (APB) where predicates are identified together with their relevant arguments in running texts. Propbanks exist for other languages such as German and Chinese (Erk and Pad, 2006; Sun and Jurafsky, 2004). To our knowledge this is the first attempt at producing an Arabic propbank. We follow here the successful model by the previous propbank efforts for other languages taking into consideration

¹In this paper, we use Arabic to refer to Modern Standard Arabic (MSA).

the specifics of the Arabic language. The creation of such a resource would not have been feasible without the existence of several crucial resources: the Arabic Treebank (ATB) (Maamouri et al., 2004) and an Arabic morphological analyser, AraMorph (Buckwalter, 2002).

2. Characteristics of the Arabic Language

Arabic is the language spoken and written by over 300 million people in the world. It is one of the official languages of the United Nations. It is also read by a big portion of Muslims in the world since it is the language of the Quran, the Muslim holy book. The Arabic script is used for Farsi, Urdu and Dari besides Arabic. The spoken form of Arabic is quite different from the written form of the language. It is one of the few languages in the world that exhibit diglossia. Diglossia is a phenomenon where several forms of the same language co-exist in a community. For any native speaker of Arabic, there exist at least two forms of the language, the spoken form which is typically a specific dialect versus a formal written form, referred to as modern standard Arabic (MSA). All the educated people of the Arab world understand MSA, it is the shared formal language. In the context of this resource, we only deal with MSA.²

Arabic is a Semitic language. It is known for its templatic morphology where words are made up of roots and affixes. Clitics agglutinate to words. For instance, the surface word *wbHsnAthm*³ *و بحسناتهم* ‘and by their virtues[fem.]’, can be split into the conjunction *w* ‘and’, preposition *b* ‘by’, the stem *HsnAt* ‘virtues [fem.]’, and possessive pronoun *hm* ‘their’. Arabic is different from English both morphologically and syntactically. Hence, Arabic is a challenging language to existing NLP technology that are too tailored to the nuances of the English language.

From the morphological standpoint, Arabic exhibits rich morphology. Similar to English, Arabic verbs are marked explicitly for tense, voice and person, however in addition, Arabic marks verbs with mood (subjunctive, indicative and jussive) information. For nominals (nouns, adjectives, proper names, and pronouns), Arabic marks case (accusative, genitive and nominative), number, gender and definiteness features. Depending on the genre of the text at hand, not all of those features are explicitly marked on naturally occurring text.

Arabic writing is known for being underspecified for short vowels. Some of the case, mood and voice features are marked only using short vowels. Hence, if the genre of the text were religious such as the case for the Quran or the Bible, or pedagogically oriented books such as those for children or second language Arabic learning, the text is typically fully specified for all the short vowels to enhance readability and disambiguation.

From the syntactic standpoint, Arabic, different from English, is considered a pro-drop language, where the subject of a verb may be implicitly encoded in the verb morphology. Hence, we observe sentences such as *اكل البرتقال* *Akl AlbrtqAl* ‘ate-[he] the-oranges’, where the verb *Akl* encodes that the subject is a 3rd person masculine singular. This sentence is exactly equivalent to *هو اكل البرتقال* *hw Akl AlbrtqAl* ‘he ate the-oranges’. It is worth noting that in the Arabic Treebank (ATB), 30% of all sentences are pro-dropped for subject.

Also Arabic is different from English in that it exhibits a larger degree of free word order. For example, Arabic allows for subject-verb-object (SVO) and verb-subject-object (VSO) argument orders, as well as, OSV and OVS. In the ATB, we observe an equal distribution of both VSO and SVO orders each equally 35% of the time. An example of an SVO sentence is *الرجال اكلوا البرتقال* *AlrjAl AklwA AlbrtqAl* ‘the-men ate-them the-oranges’, this is contrasted with *اكل الرجال البرتقال* *Akl AlrjAl AlbrtqAl* ‘ate the-men the-oranges’.

Arabic exhibits more complex noun phrases than English mainly to express possession. These constructions are known as *idafa* constructions. One example of these complex structures an indefinite noun is followed by a definite noun. For example, *رجل البيت* *rjl Albyt* ‘man the-house’ meaning ‘man of the house’. Accordingly, MSA does not have a special prepositional use to express possession in a manner similar to English.

3. Pilot Arabic Propbank

3.1. Definition

In a propbank, a proposition is annotated with the predicate and its arguments are identified. The sentence *مشروع الامم المتحدة فرض مهلة نهائية ل إتاحة الفرصة امام قبرص* *m\$rwE AlAmm AlmtHdp frD mhlp nhAyp l AtAHp AlfrSp AmAm qbrS*. meaning ‘The United Nations’ project imposed a final grace period as an opportunity for Cyprus’, is annotated as follows: $[m\$rwE AlAmm AlmtHdp]_{ARG0} [frD]_{PREDICATE} [mhlp nhA]_{yp}]_{ARG1} [l AtAHp AlfrSp AmAm qbrS]_{ARGM-PRP}$.

In the propbank annotation framework, the semantic annotations are based on an underlying syntactic structure. In the case of the Arabic propbank, we relied on the ATB constituency parses for the underlying syntax, hence in this example: ARG0 corresponds to the subject of the sentence which is *m\$rwE AlAmm AlmtHdp* ‘United Nations project’; ARG1 corresponds to the object position, namely, *mhlp nhAyp* ‘final grace period’. The predicate has an ARGM-PRP (purpose argument) in *l AtAHp AlfrSp AmAm qbrS* ‘as an opportunity for Cyprus’.

3.2. Design

The design of the Arabic propbank is very similar in terms of steps to the design steps taken for previous languages. The general procedure is based on creating framesets for

²In this paper we use MSA and Arabic interchangeably.

³We use the Buckwalter transliteration scheme to show romanized Arabic (Buckwalter, 2002).

verbs and then using them as annotation guidelines for the annotators. The framesets identify the predicate and its possible arguments. For example for the verb *AqAm AqAm* has the following three framesets and some associated examples:

- وتم اكتشاف اول منزل اقامت فيه عائلة باخ
w tm Akt\$Af Awl mnzl AqAmt fyh EA}lp bAx
 ‘The first house that Bach’s family **lived in** had been discovered’

PREDICATE: *AqAmt* ‘to live’
ARG0 is ‘experiencer’: *E}lp bAx* ‘Bach’s family’
ARG1 is ‘location’: *fyh*⁴ ‘in it’

- و كانت مجموعة اسلامية مسلحة اقامت حاجزا وهميا على الطريق الدولي رقم واحد
w kAnt mjmwEp AslAmyp mslHp AqAmt HAjzA whmyA ELY AlTryq Aldwly rqm wAHd
 ‘An armed Islamic group had **set up** a false barrier on international route no.1.’

PREDICATE: *AqAmt* ‘to set up’
ARG0 is ‘organizer’: *mjmwEp AslAmyp mslHp* ‘an armed Islamic group’
ARG1 is ‘event’: *HjzA whmyA* ‘a false barrier’
ARGM-LOC: *ELY AlTryq Aldwly rqm wAHd* ‘on international route no. 1’

- و قال: يجب ان نقيم دعوى ونوكل محاميا
w qAl: ”yjb An nqym dEwY wnwkl mHAMyAF”
 ‘He said: “We need to file a complaint and hire an attorney.”’

PREDICATE: *nqym*⁵ ‘to file’,
ARG0 is ‘agent’: NONE * (pro-dropped marked with an inflection on the verb)
ARG1 is ‘theme’: *dEwY* ‘claim’

Accordingly, the creation of the Arabic propbank is an iterative process once the set of predicates to be annotated are chosen. The iteration occurs between the frame creation phase and the annotation phase. The steps taken are described as follows:

1. Once a predicate (in this case a verb) is chosen, framers look at an average sample size of 50 instances per predicate found in the corpus in order to get an idea of its syntactic behavior. Framers utilize other tools to help examine the predicate at hand thoroughly - especially for low frequency ones - by consulting dictionaries and additional corpora (including the internet).

2. A frame is created based on the semantic-syntactic interaction of the verb. Arguments are chosen based on what the verb requires in order to complete its meaning. If the verb has more than one sense, it is divided into more than one frame depending on how it relates to its arguments. The arguments are chosen based not only on what is deemed semantically necessary, but also on frequency of usage.

3. The annotators tag each instance of the verb and its arguments based on the frames that have been created. Adjuncts have a preset number of tags: e.g. LOC for location, TMP for temporal, etc. If the annotators find instances within a verb that do not fit any of the created frames, a framer is consulted and a discussion ensues around whether or not another frame needs to be created to fit the specific instance.

In designing such a resource, several relevant constraints had to be considered.

Given the rich morphology of Arabic as well as its templatic morphology, there exist several possibilities for the level of granularity for the propbank entry representation. The question arises whether we should use a lemma or the root given that traditional Arabic dictionaries list words using the roots. In our current formulation, a lemma is the underlying derivation of root and pattern including the vocalic structure. While the root comprises only the radicals which are typically three letters.⁶ Given the Arabic writing system, deriving both the lemma and the root from the surface form of the word is not straight forward. Since the writing system is underspecified for the vowels, even after stripping out the inflections, the underlying lemma is not deterministic. Yet, there exist automatic systems that can predict the underlying lemma with high accuracy (Habash and Rambow, 2007). While we do not know of highly accurate systems for root prediction, one can also envision an automated approach to the task.⁷ However, conceptually, listing predicates by root allows for too high a level of abstraction as several lemmas could belong to the same root. Lemmas on the other hand encode exactly the right level of morphological and semantic information that is relevant for propositions. A templatic pattern in Arabic dictates the number of arguments for a verb, hence it is at the syntactic-semantic interface. Accordingly, we decided to use lemmas for indexing our predicates in the propbank repository.⁸

Syntactic annotations assumed in the ATB are abided by.⁹ Subtle decisions made by ATB syntacticians with

⁶Most of the Arabic verbs are trilateral.

⁷We do recognize how difficult automatically identifying the underlying root is especially when one of the weak consonants *w*, *A*, *y* is involved.

⁸We do mark the roots on the frameset entries such that at a future point in time we may link framesets pertaining to the same root.

⁹However, in the process, inconsistencies in the syntactic annotations for the ATB were reported to the ATB syntactic annotators.

⁴Traced back to *Awl mnzl* ‘first house’

⁵This is a light verb construction.

respect to unexplicit information in the surface language expression such as dealing with traces in the case of pro-dropped arguments are exploited in the APB annotation. Moreover, since Arabic script is underspecified by nature for short vowels, the morphological decisions adopted by the treebank morphological annotators are abided by for the APB annotations. Hence, identifying passivization and case endings on single nominals were given in the ATB and used in the APB annotations. For example, in a sentence such as *ftH AlbAb bAlmftAH* `فتح الباب بالفتاح`, where *ftH* ‘open’ is the predicate, depending on the assumed lemma for the predicate, the meaning of the sentence may change leading to a change in the argument structure. If *ftH* is assumed to be an active verb, then the underlying form is *fataHa*, then the sentence means ‘the key opened the door’, where the predicate is *fataHa* ‘open’, ARG0 is the pro-dropped subject typically marked in the ATB as ‘-NONE-*’, ARG1 is *AlbAb* ‘the-door’, and ARG2 is *bAlmftAH*. However, if the ATB morphological annotators deemed the predicate passive, *futiHa*, then the sentence meaning changes leading to a difference in number of arguments. Hence, the meaning changes to ‘the door was opened by the key’, where *futiHa* is the predicate, *AlbAb* is ARG1 and *bAlmftAH* is ARG2. It is worth noting that in Arabic agents are not allowed to actualize in the passive construction, therefore a passive predicate can not have an ARG0 in Arabic.

The current pilot APB is based on 200k words from the ATB3 version 2. It comprises 560 predicates annotated by at least one annotator. The chosen verbs occur at least 12 times in the corpus covering 80% of the ATB data. The predicates are fully specified for diacritization as lemmas. The lemmas are devoid of clitics such as object pronouns and inflections such as subject markers whether indicating a pro-drop or as an agreement marker. Hence, no two derivationally variant verbs are conflated.

APB defines an overall 24 argument types. The argument types follow in the traditional Propbank annotation style with 5 numbered arguments (ARG0, ARG1, ARG2, ARG3, ARG4) and 19 adjunctive arguments (ARG0-STR, ARG1-PRD, ARG1-STR, ARG2-STR, ARGM, ARGM-ADV, ARGM-BNF, ARGM-CAU, ARGM-CND, ARGM-DIR, ARGM-DIS, ARGM-EXT, ARGM-LOC, ARGM-MNR, ARGM-NEG, ARGM-PRD, ARGM-PRP, ARGM-REC, ARGM-TMP).

4. Challenges to APB Annotation

There are several interesting similarities between the English PB (EPB) and the Arabic Propbank (APB) annotation process. For instance, in the manner in which both handle misplaced arguments, the propbank annotators rely on the empty categories (traces and null elements) to assign the relevant argument role labels. However of more interest to this paper are the places where there are differences between the two languages that lead to variations in the annotations. As shown earlier in Section 2., there are several crucial structural and morphological differences

between English and Arabic. These differences made the annotation process different from the process for English Propbank annotation.

Many of the sentences in the Arabic treebank data are very long with many relative clauses which lead to many issues with scoping of arguments. Several sentences exhibit VP elipsis in conjunctive sentences, therefore the annotators had to assign arguments to an elided verb. This raised the issue of whether the verbs need to be rendered explicit in the treebank and then annotated.

The *idafa* constructions, expressing possession, are significantly present indicating NPs with embedded structure. This leads to complex scoping questions. English does not exhibit similar complex NP structures.

The relative free word order in Arabic combined with the implicit case morphology resulted in some issues in annotation. Even with the adoption of the syntactic decisions made by ATB annotators, the APB annotators had long discussions on whether some arguments are ARG0 or ARG1 since the structure could license both readings.

Small clauses are the cases where verbs like ‘expect’, ‘appear’, and ‘consider’, take clauses as complements. In the EPB, the complements are marked as ARG1. Similarly for equivalent Arabic verbs that take a small clause as a complement, the complement is annotated as ARG1. The problem in Arabic arose with verbs like {*iEotabara* | `اعتبر` ‘consider’, where in one of its framesets it was found that it can take a small clause as its argument rather than a complete clause or a noun phrase. In the active, it isn’t so much a problem except that the syntactic boundaries aren’t intuitive for the Arabic speaker.

In the following example, the dropped pronoun is ARG0 and the clause that starts with the accusative pronominal clitic is ARG1:

```
(S (VP (yaEotabiruw|يعتبروا |‘they-consider’)
  [(NP-SBJ (-NONE- *))ARG0]
  [(S (NP-SBJ (ha|ها |‘it’))
    (ADJP-PRD (gayora|غير |‘not’)
      (mujodiyapK|مجدية |useful)))]ARG1-S]
(PUNC .)
```

This gets further complicated in the passive. Per the ATB, there is always a trace in a passive construction between the actualized subject and the underlying subject that has been dropped due to the passive construction. In small clauses, that trace is the head of the small clause and it appears as if there are two consecutive empty categories which confuses the annotators. Accordingly, this was dealt with such that the annotators only mark the small clause as the argument for the verb.¹⁰

This construction is illustrated in the following passive

¹⁰During post-processing the trace to the underlying subject will be coreferenced.

verb example:

(SBAR (WHNP-3 (Al atiy|التي|‘which’))
(S (VP (AuEotubirat|اعتبرت|‘considered-FEM’)
(NP-SBJ-3 (-NONE- *T*))
[(S (NP-SBJ-3 (-NONE- *))
(NP-PRD (niha|هي|‘final’)))]_{ARG1}]
[(PP-TMP (baEoda|بعد|‘after’)
(NP (AlHarobi|الحرب|‘the-war’)))]_{ARGM-TMP}]
(PUNC .))

Hence, the annotators need to keep in mind the fact that not only is the argument for the verb a clause, its boundaries aren’t necessarily intuitive to the speaker.

Arabic, in contrast to English, does not allow for agent arguments to actualize in the passive construction. Therefore, a sentence in English such as ‘John was hit by Mary’, if translated to Arabic in the passive form, *Durib jwn bimAry* جون ضرب بماري will be translated as ‘hit John by-Mary [’s hand]’, where *bimAry* is implicitly an instrument. Accordingly, in the English construction ‘by-Mary’ will be annotated as ARG0, where in the Arabic construction it will be annotated by ARG2. In fact, there will never be an ARG0 with a passive in Arabic. The only argument of any subjective nature allowed is an instrumental subject and it is annotated as ARG2.

5. Conclusions and Future Directions

We presented a description of a pilot Arabic propbank creation effort. We describe some of the divergences from English and some of the challenges faced in the annotation process. In the future, we plan to expand the scope of the investigated verbs in light of new treebank annotations. Moreover, we would like to investigate applying the same frames to dependency treebanks such as the Prague Dependency treebank. We would like to add more genres and domains to the corpora annotated. Finally, we plan to investigate methods for automatically bootstrapping the annotation process.

6. Acknowledgements

This research has been funded by DARPA through a subcontract to BBN Inc. to University of Colorado, LDC-UPenn and Columbia University.

7. References

Collin F. Baker, 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.

Tim Buckwalter. 2002. Buckwalter Arabic Morphological Analyzer Version 1.0. Linguistic Data Consortium, University of Pennsylvania, 2002. LDC Catalog No.: LDC2002L49.

Xavier Carreras and Lluís Màrquez. 2005. Introduction to the CoNLL-2005 shared task: Semantic role labeling. In *Proceedings of CoNLL-2005*, Ann Arbor, Michigan.

John Chen 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*, Sapporo, Japan.

Mona Diab, Alessandro Moschitti, and Daniele Pighin. 2007. Cunit: A semantic role labeling system for modern standard arabic. In *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, pages 133–136, Prague, Czech Republic, June. Association for Computational Linguistics.

Katrin Erk and Sebastian Pad. 2006. Shalmaneser - a flexible toolbox for semantic role assignment. In *Proceedings of LREC-06*.

Daniel Gildea and Daniel Jurafsky. 2002. Automatic labeling of semantic roles. *Computational Linguistics*, 28(3):245–288.

Daniel Gildea and Martha Palmer. 2002. The necessity of parsing for predicate argument recognition. In *Proceedings of the 40th Annual Conference of the Association for Computational Linguistics (ACL-02)*, Philadelphia, PA, USA.

Nizar Habash and Owen Rambow. 2007. Arabic diacritization through full morphological tagging. In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Companion Volume, Short Papers*, pages 53–56, Rochester, New York, April. Association for Computational Linguistics.

Paul Kingsbury and Martha Palmer. 2003. Propbank: the next level of treebank. In *Proceedings of Treebanks and Lexical Theories*.

Mohamed Maamouri, Ann Bies, Tim Buckwalter, and Wig dan Mekki. 2004. The penn arabic treebank : Building a large-scale annotated arabic corpus.

Alessandro Moschitti, Ana-Maria Giuglea, Bonaventura Coppola, and Roberto Basili. 2005. Hierarchical semantic role labeling. In *Proceedings of CoNLL-2005*, Ann Arbor, Michigan.

Alessandro Moschitti. 2004. A study on convolution kernels for shallow semantic parsing. In *proceedings of the 42th Conference on Association for Computational Linguistic (ACL-2004)*, Barcelona, Spain.

Martha Palmer, Dan Gildea, and Paul Kingsbury. 2005. The proposition bank: A corpus annotated with semantic roles. In *Computational Linguistics Journal*, number 31:1.

Sameer Pradhan, Kadri Hacioglu, Wayne Ward, James H. Martin, and Daniel Jurafsky. 2003. Semantic role parsing: Adding semantic structure to unstructured text. In *Proceedings of ICDM-2003*, Melbourne, USA.

Sameer S Pradhan, 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. Association for Computational Linguistics.

Honglin Sun and Daniel Jurafsky. 2004. Shallow semantic parsing of chinese. In Daniel Marcu Susan Dumais and Salim Roukos, editors, *HLT-NAACL 2004: Main Proceedings*, pages 249–256, Boston, Massachusetts, USA, May 2 - May 7. Association for Computational Linguistics.

Nianwen Xue 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. Association for Computational Linguistics.