

THE PARALLEL MEANING BANK: A FRAMEWORK FOR SEMANTICALLY ANNOTATING MULTIPLE LANGUAGES

Lasha Abzianidze Rik van Noord Chunliu Wang Johan Bos

CLCG, University of Groningen, the Netherlands
{l.abzianidze, r.i.k.van.noord, chunliu.wang,
johan.bos}@rug.nl

Abstract

This paper gives a general description of the ideas behind the Parallel Meaning Bank, a framework with the aim to provide an easy way to annotate compositional semantics for texts written in languages other than English. The annotation procedure is semi-automatic, and comprises seven layers of linguistic information: segmentation, symbolisation, semantic tagging, word sense disambiguation, syntactic structure, thematic role labelling, and co-reference. New languages can be added to the meaning bank as long as the documents are based on translations from English, but also introduce new interesting challenges on the linguistics assumptions underlying the Parallel Meaning Bank.

Keywords and phrases: parallel corpus, semantic annotation, meaning banking, compositional semantics, formal semantics

1 Introduction

The Parallel Meaning Bank (PMB) is a semantically annotated parallel corpus for English, Dutch, German, Italian, Chinese, and Japanese. The key idea behind the PMB is based on the assumption that translations—at least to a large extent—preserve the meaning between the source and target language. Making use of translated texts, annotation for one language can be re-used for the translations, resulting in an economical annotation platform. One of the core ideas is that the human annotations can help improve existing language technology (based on supervised machine learning) in the areas of machine translation, automatic question answering and advanced information retrieval.

The PMB can be viewed as a multilingual version of the Groningen Meaning Bank, GMB [7, 14], an annotation platform designed for the meaning of English texts. Like the GMB, the PMB contains the raw

	x_1	x_2	e_1	t_1
EN	Alfred Nobel invented dynamite in 1866.			
DE	Alfred Nobel erfand 1866 das Dynamit.			
IT	Alfred Nobel inventò la dinamite nel 1866.			
NL	Alfred Nobel vond in 1866 het dynamiet uit.			
	$\text{male.n.02}(x_1)$ $\text{Name}(x_1, \text{alfred_nobel})$ $\text{invent.v.01}(e_1)$ $\text{Time}(e_1, t_1)$ $\text{Result}(e_1, x_2)$ $\text{Agent}(e_1, x_1)$ $\text{time.n.08}(t_1)$ $\text{YearOfCentury}(t_1, 1866)$ $t_1 < \text{now}$ $\text{dynamite.n.01}(x_2)$			

Figure 1: 03/0766 PMB document has four meaning-preserving translations. As a result, each translation is annotated with the same meaning representation.

texts and various layers of linguistic annotation, ultimately resulting in a formal meaning representation based on Discourse Representation Theory (DRT) [25]. The annotations are automatically generated by a pipeline of state-of-the-art natural language processing (NLP) tools and then manually corrected by annotators. Semantic annotation is hard, even for trained linguists. To give an idea what a meaning representation in the PMB looks like, consider Figure 1. These representations are called Discourse Representation Structures (DRSs) in DRT.

This paper gives a general overview of the PMB and describes several aspects of it in more details. First, we describe the seven annotation layers that are used to automatically obtain formal meaning representations (Section 2 and Section 3). Then, we sketch how the semantic annotation can be projected from one language to another (Section 4). This is followed by an overview of applications of the released PMB data (Section 5). Finally, we show how new documents in new languages are added to the PMB and how language technology tools are bootstrapped for new languages (Section 6).

2 The Seven Annotation Layers

There are two main approaches on semantic annotation. The first approach is to go directly from the source text to the target meaning representations, without any layer of analysis in between. An example of this method is the corpus constructed for Abstract Meaning Representations [5]. The second approach, adopted in the PMB, is to view annotation as a sequence of layers of analysis, where each layer builds on the previous layer by adding a piece of (semantic) information to it. In the PMB, seven layers of annotation are distinguished:

1. Tokenisation: detecting sentence boundaries and word tokens;
2. Symbolisation: assigning a non-logical symbol to a word (or multi-word) token. This layer unifies lemmatization and normalization.
3. Word sense disambiguation: assigning concepts to symbols, based on the WordNet [23] sense inventory;
4. Co-reference resolution: marking antecedents for anaphoric expressions;
5. Thematic role labelling: annotate relations between entities using VerbNet roles [11] and comparison operators (e.g., temporal and spatial orders);
6. Syntactic analysis: providing lexical categories for each token and building a syntactic structure for the sentence, based on Combinatory Categorical Grammar [36];
7. Semantic tagging: assigning a semantic type to a word token [3].

These annotation layers are demonstrated in Figure 2. The annotation layers provide all information needed to provide a compositional semantic analysis for a sentence (for additional details about the PMB annotation layers see [2]). This is done by using the lambda calculus, and adopting Discourse Representation Theory as semantic formalism, implemented by the semantic parser Boxer [13]. In a final step, the semantic analysis of single sentences are combined into one meaning representation covering the entire text.

3 Annotation Pipeline

Manually creating the seven annotation layers for a large amount of documents is not a feasible task. For this reason, we use an annotation pipeline to automatically segment raw documents, label tokens with token-based annotations, and produce the final meaning representation. The pipeline consists of a sequence of NLP tools each serving for a specific annotation layer. The pipeline of English-specific tools is highlighted with a green background in Figure 3. Below, we briefly describe each NLP tool:¹

- Elephant [21] is used for tokenisation. The tool performs sentence boundary and word token detection as a single labelling task: each

¹Currently, the pipeline lacks specialized NLP tools for word sense disambiguation and co-reference resolution. Therefore, these layers are manually annotated for now.

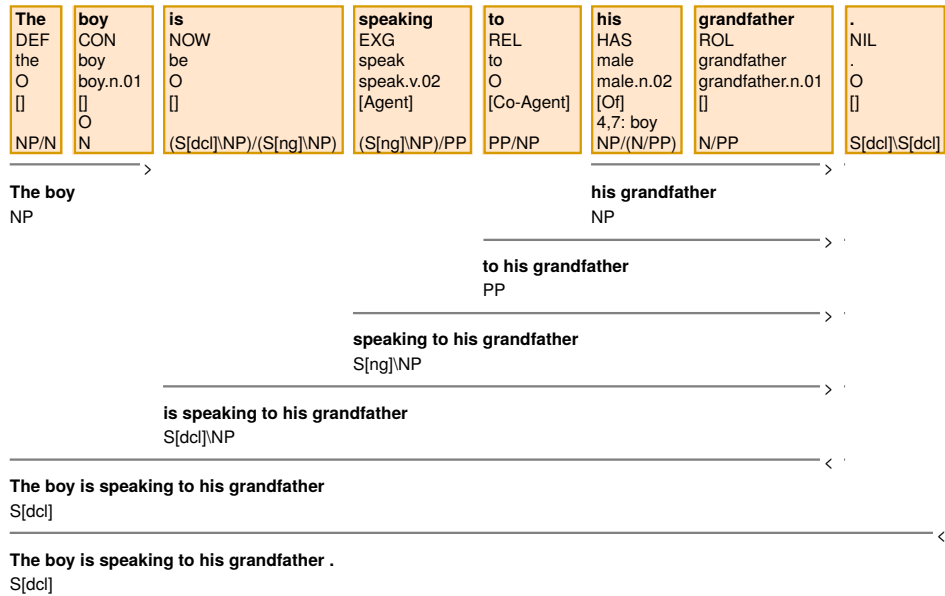


Figure 2: All the seven annotation layers of the English translation of 46/2924 PMB document. The order of layers starting from top: tokenisation, semantic tagging, symbolisation, word sense disambiguation, thematic role labelling, co-reference resolution, and syntactic analysis.

character is labelled with one of the four labels depending on being sentence beginning, token beginning, inside token, and outside token;

- Semantic tagging is carried using the tri-gram based TnT tagger [15];
- The lemmatisation part of symbolisation is done with the help of the lemmatizer Morpha [33]. Currently, we use instance-based learning for the normalisation part. In particular, for every existing combination of lemma and semantic tag in the PMB, the most frequent symbol is memorized which is later reused to tag a token with the corresponding pair of semantic tag and lemma. For example, to get a symbol for a token *eight*, first its lemma **eight** and semantic tags QUC is obtained and then the instance-based learning will assign 8 as a symbol to it
- Obtaining syntactic analysis consists of assigning lexical categories to tokens and constructing a derivation tree over these categories. The both subtasks are performed using EasyCCG [28], a CCG-based parser that requires only tokenised input and pre-trained word embeddings.




Figure 3: The PMB pipeline: a sequence of NLP tools that processes raw texts and outputs formal meaning representations. The hand icon indicates functionality of overwriting parts of the system outputs with manual annotations.

- Thematic role labelling is done with a tagger based on Conditional Random Fields [26]. The tagger employs semantic tags, symbols and CCG lexical categories as features to predict thematic roles.

The output of each tool can be manually corrected by human annotators.² In this way, we use a human-in-the-loop approach to obtain gold standard annotation layers and the final meaning representations. We also apply bootstrapping with the gold standard annotation layers to retrain and further improve the quality of the NLP tools. This aims at reducing human annotation efforts while still retaining high quality system outputs.

4 Annotation Projection

The previous two sections gave a rough overview of what is required to provide a compositional analysis for the meaning of a text for one language. For historical reasons, this language is English, because of the tools developed earlier in the Groningen Meaning Bank [6]. Instead of starting from scratch and implementing a pipeline for other languages, we follow a different approach in the PMB. This approach is called *annotation projection*, and requires that the English text has an adequate translation in the language of your choice. The first languages that we added in the PMB were languages close to English, such as other Germanic languages (Dutch and German) and Italian, a Romance language.

²The PMB documents can be manually annotated with the PMB explorer, an online annotation environment, available at: <https://pmb.let.rug.nl/explorer>. Anybody can register and annotate the documents.

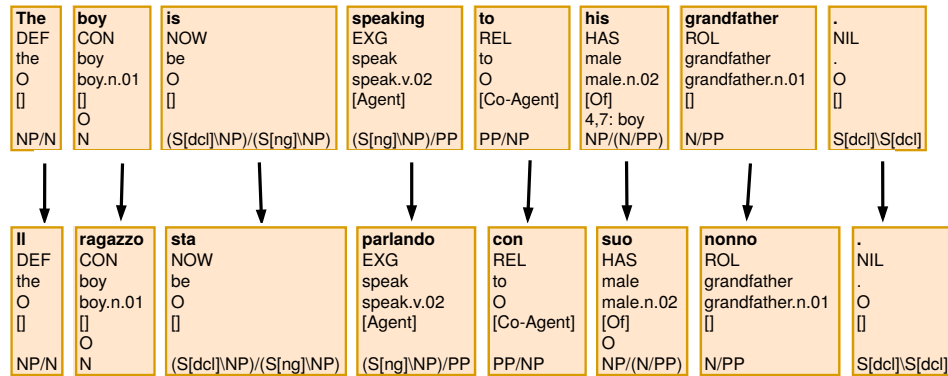


Figure 4: An example of a complete annotation projection: all the seven annotation layers are projected from English to Italian.

The idea of semantic projection is extremely simple, but its implementation is surprisingly challenging even for closely-related languages. The assumption that a translation doesn't change much of the meaning, is the driving force in this approach. But for reasons of scalability, we are not just interested in the final meaning representation, but also in the compositional analysis supporting this final meaning representation. This makes projection more challenging.

In the PMB, annotation projection is implemented using word alignment between English and the target language.³ The alignments provide clues how to transfer the layers of annotation from English to the other languages [19]. For cases where the syntactic structure of the target language is similar to that of the source language (English), this is often straightforward. Figure 4 shows one of such cases where a literal translation leads to a perfect word alignment and therefore to a complete annotation projection. This leads to the very same meaning representation for the Italian translation that the English translation had.

But translations are not always in a perfect word-to-word and order-preserving correspondence as in the previous example. Even closely-related language show different behaviour with respect to a word order, multi-word expressions, definiteness, use of articles, and noun-noun compounds. So automatic projection requires the help of human annotators to provide corrections.

In the PMB, we go beyond the mere annotation projection as it is brittle for wide-coverage translations. To do so, using the same NLP tools, we (re-)train semantic tagging and syntactic parsing models for non-English

³We employ GIZA++ [34] to automatically induce word alignments.

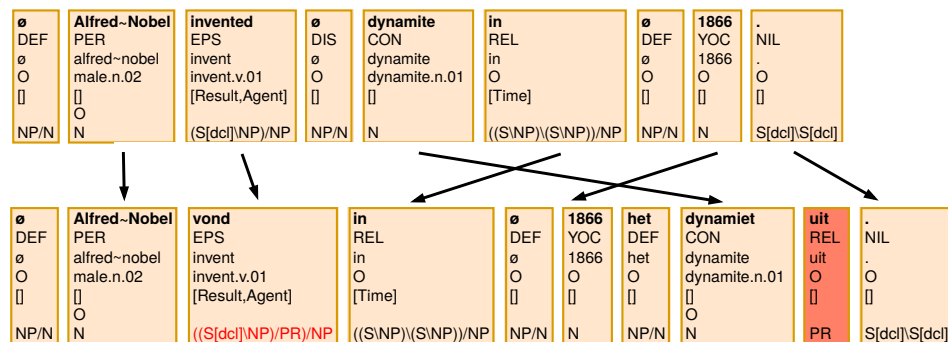


Figure 5: An imperfect annotation projection is compensated by the language-specific syntactic parsing model.

languages. Initially the training data consisted of translations with perfect annotation projections. Gradually the training data increased as a result of reprocessing the rest of the translations with new models and correcting manually where necessary. For example, the annotation projection in Figure 5 fails for the syntactic analysis layer due to the difference in a word order of the Dutch translation. But with the help of the in-house trained Dutch model of the parser, it is possible to automatically recover a correct syntactic analysis of the Dutch translation, which eventually leads to the same meaning representation (see Figure 1).⁴

Figure 3 shows the PMB pipeline of NLP tools that simultaneously processes documents in five languages. While currently only symbols and thematic roles are projected for the Dutch, German, and Italian translations, the Japanese translations also get semantic tags projected from the English translations. In the near future, we plan to retrain Japanese-specific model for the semantic tagging.

Currently we are investigating what consequences semantic annotation projection has on languages that behave significantly different from English from a linguistic perspective. Here we think of languages such as Chinese and Japanese, and perhaps also Kartvelian languages such as Georgian [35]. These languages add pressure on the principles of the PMB, in particular on the extent one can adopt a single framework for each layer in the semantic analysis pipeline. To give a first example, the semantic tags might be subject to extension of the tagset for new languages that show phenomena that cannot be captured with the existing semantic categories. To give a second example, we assume CCG as the theory of syntactic structure suitable for all languages. CCG starts with a base of atomic categories, which work

⁴To verify whether projected annotations yield the same meaning representation as of English, we perform fine-grained matching of meaning representations [40].

well for Germanic languages, but other languages could be hard to adopt in the parameters provided for English. In future work we need to take a closer look at such a wider perspective. As a final example, in Chinese, there are less syntactic constraints for verbs, but there is widespread use of pro-drop, and a larger distribution of ambiguous constructions, such as the relative clause and verbal coordination. In addition, the inherent ambiguities caused by both verbal coordination and relative clauses of Chinese make semantic parsing more difficult than syntactic parsing [48].

5 Applications

The PMB annotations are released periodically, free of charge.⁵ It includes gold standard data, which is fully manually corrected, as well as silver (partially manually corrected) and bronze (with no manual corrections) data. The releases so far contain documents for English, German, Italian and Dutch, but for future releases we plan to include Chinese and Japanese. An overview of the releases is shown in Table 1.

Table 1: Number of released documents per language for the five current PMB releases.

Release	Quality	EN	DE	IT	NL
PMB-1.0.0	Gold	2,049	641	387	394
PMB-2.0.0	Gold	3,925	1,048	568	527
	Silver	66,693	611	266	192
PMB-2.1.0	Gold	4,555	1,175	635	586
	Silver	71,308	688	306	207
PMB-2.2.0	Gold	5,929	1,419	724	633
	Silver	67,965	4,235	2,515	1,051
	Bronze	120,622	102,998	61,504	20,554
PMB-3.0.0	Gold	8,403	1,979	1,062	1,012
	Silver	97,598	5,250	2,772	1,301
	Bronze	146,371	121,111	64,305	21,550

One of the goals of the PMB releases is to aid DRS parsing, a task in which a model has to automatically produce a DRS from raw text. These produced DRSs can then potentially be of benefit in other language related tasks, such as machine translation or question answering. Early approaches used rule-based system for only small fragments of English [24,43], though

⁵<https://pmb.let.rug.nl/data.php>

wide-coverage semantic parsers that use supervised machine learning were also developed, mainly on the GMB data [12, 27, 13, 30, 31].

The main advantage of the PMB is that it contains gold standard data for evaluating the parsers. This is in contrast to the GMB, which contains partially manually corrected evaluation sets that are not guaranteed to be gold standard. This allowed for the organization of a shared task on English DRS parsing in PMB format [4]. Five systems participated in this shared task, which all used neural networks in some capacity. Three systems used sequence-to-sequence models based on the first PMB-based DRS parser [41], which was extended by including linguistic features [42, 39] and by swapping the bi-LSTM encoder/decoder for a transformer model [29], which was the winning system. The two other systems consisted of a transition-based parser that relied on stack-LSTMs [20] and a neural graph parsing system that converted the DRSs to a more general graph format before parsing [22]. The latter is also the first system that produced results for German, Italian and Dutch DRS parsing.

There are also other applications of the PMB data. For one, semantic tagging can be useful as either an auxiliary task to improve a main task [10, 9, 1], or as a general dataset for evaluating neural architectures [8, 32, 16, 18]. Moreover, PMB data has been used in research on natural language inference [45] and machine translation [17].

6 A Look at the Future: Extending the PMB

The PMB can be extended in terms of introducing new documents or new translations. Translations may belong to languages that are new or already covered in the PMB. In case a translation belongs to a new language, its integration in the PMB requires more work as the new language needs to be processed by the PMB pipeline. In this section, we describe the procedure and conditions for extending the PMB.

The simplest extension procedure is when adding translations to PMB documents in one of the PMB (non-English) languages, let's say L_{PMB} . In this case, no new documents are created, and there is no need to develop new NLP tools as the PMB pipeline can already process texts in L_{PMB} . If the PMB uses the projection method for L_{PMB} , then it is necessary to align the new L_{PMB} -translations to the existing English translations. For the best results, the alignment is usually done on all PMB English- L_{PMB} bitexts. This might affect the alignments of old PMB documents and annotations of the projected layers, consequently. Since the alignment is carried out on more bitexts than before, the assumption is that the quality of alignments improves. Whether the change influences alignments negatively, this can

be verified for the translations already having a gold standard annotation for the projected layers. The difference for the projected layers will show up as conflicts with the gold standard.

Adding a new parallel corpus to the PMB involves adding completely new documents. Taking the architecture of the PMB into account, one of the languages of the new corpus must be English. Let's first consider the scenario when all the languages of the corpus are covered by the PMB. All new documents (consisting of translations) get new `part/doc` identifiers and are uniformly distributed over all the 100 parts of the PMB. If the newly added documents belong to a text genre new to the PMB, some NLP tools in the pipeline might require further adaptation. For example, if the documents belong to the social media domain, one might need to correct the tokenization or semantic tagging of slang words and retrain the corresponding tools on the corrected annotations. Additionally, the procedures of inducing new alignments and verifying the changes caused by them are also applicable in this scenario.

The case where newly added parallel corpus contains translations not belonging to the PMB languages is the most laborious. New languages require their own annotation pipelines. Here, we describe our first experiences from adding Japanese [46] and Chinese, using translations from Tatoeba.⁶

To enable the annotation projection from English to Japanese, it is necessary to extract word alignments from the bitext, which itself presupposes tokenisation of the Japanese translations. Since we strive to use the same NLP tools with language-specific models for each annotation layer, we trained a Japanese model of the Elephant tokenizer.⁷ After extracting the word alignments, token-based annotations were projected for one-to-one word alignments. Since English and Japanese are languages with radically different typologies, the annotation projection for the syntactic analysis failed for almost all Japanese translations. As syntactic analyses play a key role for obtaining meaning representations in the PMB because they contribute to defining lexical semantics and guiding composition of phrasal semantics, a quick integration required a Japanese CCG parser in the PMB pipeline. Fortunately, there exists a Japanese CCG parser, `depCCG` [47]. We trained a new Japanese model for `EasyCCG` on the output of `depCCG`. We opted for training a new model to keep the PMB pipeline lean rather than integrating an additional tool in it. In the near future, we plan to train a Japanese model for the semantic tagging in order to eliminate "holes" in the semantic tagging layer caused by the annotation projection.

We are currently adding (Mandarin) Chinese translations for the PMB

⁶<https://tatoeba.org>

⁷The training data was obtained by processing the Japanese translations in the PMB with the `UDPipe 1.2.0` [37] and the model `japanese-gsd-ud-2.3-181115`.

documents. While doing so, we are taking a route similar to the one we took for Japanese. To train the Chinese model for Elephant, we used the output from jieba.⁸ The EasyCCG model was trained on the CCG derivation trees which were obtained from the Chinese Treebank [44] following [38].

The current undertakings of adding more languages to the framework doesn't mean that all problems are solved. The entire PMB enterprise emits a formal flavour of universality of language analysis. This is reflected in the practical use of our language technology pipeline, with the aim of using the same NLP tools but employing the language-specific models as the only variable element. We have reached a high level of generalization, but there are also many refinements that seek improvement, in particular on the ontological, categorial, and contextual level. The only way to make progress in this area of computational semantics is by considering other languages and getting your hands dirty!

References

1. ABDOU, M., KULMIZEV, A., RAVISHANKAR, V., ABZIANIDZE, L., AND BOS, J. What can we learn from semantic tagging? In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing* (2018), pp. 4881–4889.
2. ABZIANIDZE, L., BJERVA, J., EVANG, K., HAAGSMA, H., VAN NOORD, R., LUDMANN, P., NGUYEN, D.-D., AND BOS, J. The Parallel Meaning Bank: Towards a multilingual corpus of translations annotated with compositional meaning representations. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers* (Valencia, Spain, April 2017), Association for Computational Linguistics, pp. 242–247.
3. ABZIANIDZE, L., AND BOS, J. Towards universal semantic tagging. In *Proceedings of the 12th International Conference on Computational Semantics (IWCS 2017) – Short Papers* (Montpellier, France, September 2017), Association for Computational Linguistics.
4. ABZIANIDZE, L., VAN NOORD, R., HAAGSMA, H., AND BOS, J. The first shared task on discourse representation structure parsing. In *Proceedings of the IWCS Shared Task on Semantic Parsing* (Gothenburg, Sweden, May 2019), Association for Computational Linguistics.

⁸<https://github.com/fxsjy/jieba>

5. BANARESCU, L., BONIAL, C., CAI, S., GEORGESCU, M., GRIFFITT, K., HERMJAKOB, U., KNIGHT, K., KOEHN, P., PALMER, M., AND SCHNEIDER, N. Abstract Meaning Representation for sembanking. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse* (Sofia, Bulgaria, 2013), pp. 178–186.
6. BASILE, V., BOS, J., EVANG, K., AND VENHUIZEN, N. Developing a large semantically annotated corpus. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC 2012)* (Istanbul, Turkey, 2012), pp. 3196–3200.
7. BASILE, V., BOS, J., EVANG, K., AND VENHUIZEN, N. A platform for collaborative semantic annotation. In *Proceedings of the Demonstrations at the 13th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2012)* (Avignon, France, 2012), pp. 92–96.
8. BELINKOV, Y., MÀRQUEZ, L., SAJJAD, H., DURRANI, N., DALVI, F., AND GLASS, J. Evaluating layers of representation in neural machine translation on part-of-speech and semantic tagging tasks. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)* (2017), pp. 1–10.
9. BJERVA, J. Will my auxiliary tagging task help? estimating auxiliary tasks effectivity in multi-task learning. In *Proceedings of the 21st Nordic Conference on Computational Linguistics* (2017), pp. 216–220.
10. BJERVA, J., PLANK, B., AND BOS, J. Semantic tagging with deep residual networks. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers* (Osaka, Japan, 2016), pp. 3531–3541.
11. BONIAL, C., CORVEY, W. J., PALMER, M., PETUKHOVA, V., AND BUNT, H. A hierarchical unification of LIRICS and VerbNet semantic roles. In *Proceedings of the 5th IEEE International Conference on Semantic Computing (ICSC 2011)* (2011), pp. 483–489.
12. BOS, J. Wide-Coverage Semantic Analysis with Boxer. In *Semantics in Text Processing. STEP 2008 Conference Proceedings*, J. Bos and R. Delmonte, Eds., vol. 1 of *Research in Computational Semantics*. College Publications, 2008, pp. 277–286.
13. BOS, J. Open-domain semantic parsing with Boxer. In *Proceedings of the 20th Nordic Conference of Computational Linguistics (NODALIDA 2015)* (2015), B. Megyesi, Ed., pp. 301–304.

14. BOS, J., BASILE, V., EVANG, K., VENHUIZEN, N., AND BJERVA, J. The Groningen Meaning Bank. In *Handbook of Linguistic Annotation*, N. Ide and J. Pustejovsky, Eds. Springer Netherlands, 2017.
15. BRANTS, T. Tnt: A statistical part-of-speech tagger. In *Proceedings of the Sixth Conference on Applied Natural Language Processing* (Stroudsburg, PA, USA, 2000), ANLC '00, Association for Computational Linguistics, pp. 224–231.
16. DALVI, F., SAJJAD, H., DURRANI, N., AND BELINKOV, Y. Exploiting redundancy in pre-trained language models for efficient transfer learning. *arXiv preprint arXiv:2004.04010* (2020).
17. DURRANI, N., DALVI, F., SAJJAD, H., BELINKOV, Y., AND NAKOV, P. One size does not fit all: Comparing NMT representations of different granularities. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)* (Minneapolis, Minnesota, June 2019), Association for Computational Linguistics, pp. 1504–1516.
18. EK, A., BERNARDY, J.-P., AND LAPPIN, S. Language modeling with syntactic and semantic representation for sentence acceptability predictions. In *NEAL Proceedings of the 22nd Nordic Conference on Computational Linguistics (NoDaLiDa), September 30-October 2, Turku, Finland* (2019), no. 167, Linköping University Electronic Press, pp. 76–85.
19. EVANG, K. *Cross-lingual Semantic Parsing with Categorical Grammars*. PhD thesis, University of Groningen, 2016.
20. EVANG, K. Transition-based DRS parsing using stack-LSTMs. In *Proceedings of the IWCS 2019 Shared Task on Semantic Parsing* (2019).
21. EVANG, K., BASILE, V., CHRUPALA, G., AND BOS, J. Elephant: Sequence labeling for word and sentence segmentation. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (Seattle, Washington, USA, 2013), pp. 1422–1426.
22. FANCELLU, F., GILROY, S., LOPEZ, A., AND LAPATA, M. Semantic graph parsing with recurrent neural network DAG grammars. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)* (Hong Kong, China, Nov. 2019), Association for Computational Linguistics, pp. 2769–2778.

23. FELLBAUM, C., Ed. *WordNet. An Electronic Lexical Database*. The MIT Press, Cambridge, Ma., USA, 1998.
24. JOHNSON, M., AND KLEIN, E. Discourse, anaphora and parsing. In *11th International Conference on Computational Linguistics. Proceedings of Coling '86* (Bonn, Germany, 1986), pp. 669–675.
25. KAMP, H., AND REYLE, U. *From Discourse to Logic; An Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and DRT*. Kluwer, Dordrecht, 1993.
26. LAFFERTY, J. D., MCCALLUM, A., AND PEREIRA, F. C. N. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the Eighteenth International Conference on Machine Learning* (San Francisco, CA, USA, 2001), ICML '01, Morgan Kaufmann Publishers Inc., p. 282–289.
27. LE, P., AND ZUIDEMA, W. Learning compositional semantics for open domain semantic parsing. In *Proceedings of COLING 2012* (Mumbai, India, 2012), The COLING 2012 Organizing Committee, pp. 1535–1552.
28. LEWIS, M., AND STEEDMAN, M. A* CCG parsing with a supertag-factored model. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (Doha, Qatar, 2014), pp. 990–1000.
29. LIU, J., COHEN, S., AND LAPATA, M. Discourse representation structure parsing with recurrent neural networks and the transformer model. In *Proceedings of the IWCS 2019 Shared Task on Semantic Parsing* (2019).
30. LIU, J., COHEN, S. B., AND LAPATA, M. Discourse representation structure parsing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (Melbourne, Australia, 2018), vol. 1, pp. 429–439.
31. LIU, J., COHEN, S. B., AND LAPATA, M. Discourse representation parsing for sentences and documents. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (Florence, Italy, July 2019), Association for Computational Linguistics, pp. 6248–6262.
32. LIU, N. F., GARDNER, M., BELINKOV, Y., PETERS, M. E., AND SMITH, N. A. Linguistic knowledge and transferability of contextual

- representations. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)* (Minneapolis, Minnesota, June 2019), Association for Computational Linguistics, pp. 1073–1094.
33. MINNEN, G., CARROLL, J., AND PEARCE, D. Applied morphological processing of English. *Natural Language Engineering* 7, 3 (2001), 207–223.
 34. OCH, F. J., AND NEY, H. A systematic comparison of various statistical alignment models. *Computational Linguistics* 29, 1 (2003), 19–51.
 35. PKHAKADZE, K., CHIKVINIDZE, M., CHICHUA, G., MASKHARASHVILI, A., AND BERIASHVILI, I. An overview of the trial version of the georgian self-developing intellectual corpus necessary for creating georgian text analyzer, speech processing, and automatic translation systems. In *Reports of Enlarged Session of the Seminar of I. Vekua Institute of Applied Mathematics* (2014), vol. 28, I. Vekua Institute of Applied Mathematics.
 36. STEEDMAN, M. *The Syntactic Process*. The MIT Press, Cambridge, Ma., USA, 2001.
 37. STRAKA, M., HAJIČ, J., AND STRAKOVÁ, J. UDPipe: Trainable pipeline for processing CoNLL-u files performing tokenization, morphological analysis, POS tagging and parsing. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)* (Portorož, Slovenia, May 2016), European Language Resources Association (ELRA), pp. 4290–4297.
 38. TSE, D., AND CURRAN, J. R. Chinese CCGbank: extracting CCG derivations from the Penn Chinese treebank. In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)* (Beijing, China, Aug. 2010), Coling 2010 Organizing Committee, pp. 1083–1091.
 39. VAN NOORD, R. Neural Boxer at the IWCS shared task on DRS parsing. In *Proceedings of the IWCS 2019 Shared Task on Semantic Parsing* (2019).
 40. VAN NOORD, R., ABZIANIDZE, L., HAAGSMA, H., AND BOS, J. Evaluating scoped meaning representations. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)* (Miyazaki, Japan, 2018).

41. VAN NOORD, R., ABZIANIDZE, L., TORAL, A., AND BOS, J. Exploring neural methods for parsing discourse representation structures. *Transactions of the Association for Computational Linguistics* 6 (2018), 619–633.
42. VAN NOORD, R., TORAL, A., AND BOS, J. Linguistic information in neural semantic parsing with multiple encoders. In *IWCS 2019-13th International Conference on Computational Semantics-Short papers* (2019).
43. WADA, H., AND ASHER, N. BUILDERS: An implementation of DR theory and LFG. In *11th International Conference on Computational Linguistics. Proceedings of Coling '86* (Bonn, Germany, 1986), pp. 540–545.
44. XUE, N., ZHANG, X., JIANG, Z., PALMER, M., XIA, F., CHIOU, F.-D., AND CHANG, M. Chinese treebank 9.0 ldc2016t13, 2016.
45. YANAKA, H., MINESHIMA, K., BEKKI, D., INUI, K., SEKINE, S., ABZIANIDZE, L., AND BOS, J. Can neural networks understand monotonicity reasoning? In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP* (2019), pp. 31–40.
46. YANAKA, H., MINESHIMA, K., YAMADA, A., YAMAGUCHI, Y., KUBOTA, Y., ABZIANIDZE, L., AND BOS, J. Building a japanese version of parallel meaning bank. In *26th Annual Meeting of the Association for Natural Language Processing* (2020), pp. 1145–1158.
47. YOSHIKAWA, M., NOJI, H., AND MATSUMOTO, Y. A* CCG parsing with a supertag and dependency factored model. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (Vancouver, Canada, July 2017), Association for Computational Linguistics, pp. 277–287.
48. YU, K., MIYAO, Y., MATSUZAKI, T., WANG, X., AND TSUJII, J. Analysis of the difficulties in Chinese deep parsing. In *Proceedings of the 12th International Conference on Parsing Technologies* (Dublin, Ireland, Oct. 2011), Association for Computational Linguistics, pp. 48–57.