

Universal Semantic Annotator: the First Unified API for WSD, SRL and Semantic Parsing

Riccardo Orlando, Simone Conia, Stefano Faralli, Roberto Navigli

Sapienza NLP

Sapienza University of Rome

Piazzale Aldo Moro, 5, 00185 Rome, Italy

{orlando,navigli}@diag.uniroma1.it, {conia,faralli}@di.uniroma1.it

Abstract

In this paper, we present the Universal Semantic Annotator (USeA), which offers the first unified API for high-quality automatic annotations of texts in 100 languages through state-of-the-art systems for Word Sense Disambiguation, Semantic Role Labeling and Semantic Parsing. Together, such annotations can be used to provide users with rich and diverse semantic information, help second-language learners, and allow researchers to integrate explicit semantic knowledge into downstream tasks and real-world applications.

Keywords: Multilingualism and Language Technology for All, Word Sense Disambiguation, Semantic Role Labeling, Semantic Parsing, Multilinguality, API

1. Introduction

The grand goal of Natural Language Processing (NLP) is to realise automatic systems that are able to understand language. In order to do so, NLP combines Linguistics with Machine Learning – in particular Deep Learning – to develop increasingly “intelligent” systems that are able to process, understand, and generate natural language, as opposed to artificial language (Otter et al., 2021). Despite the tremendous progress that we have witnessed in recent years, NLP systems are still a long way from truly understanding what they process, i.e., researchers are still striving to come closer to true Natural Language Understanding (NLU). Indeed, current approaches still struggle to obtain human performance on many complex tasks that involve identifying and extracting the meaning conveyed by texts (Navigli, 2018). This is true especially for those tasks that require explicit semantic knowledge that human beings usually acquire by experiencing the real world (Bender and Koller, 2020); for example, associating the correct meaning to a word in context, understanding agentive-patientive relations between sentential constituents, and representing the interplay between different concepts.

Even if we are still far from achieving true NLU, explicit semantic knowledge is already being employed with success in an increasingly growing body of applications spanning multiple areas of AI including NLP with Information Retrieval (Christensen et al., 2010; Blloshmi et al., 2021b), Question Answering (He et al., 2015; Ramakrishnan et al., 2003), Text Summarization (Hardy and Vlachos, 2018; Liao et al., 2018), and Machine Translation (Marcheggiani et al., 2018; Raganato et al., 2019; Song et al., 2019), but also Computer Vision with Visual Semantic Role Labeling (Gupta and Malik, 2015), Situation Recognition (Yatskar et al., 2016), and Video Understanding (Sadhu et al., 2021),

inter alia. One important downside of such progress is that the complexity reached by current techniques has become a significant issue: understanding how a state-of-the-art system works now requires expert knowledge of linguistic theories and Deep Learning methods. This issue is exacerbated when dealing with multiple languages, since learning to transfer knowledge from possibly distant languages leads to further complications. As a consequence, one could argue that accessibility to modern techniques, rather than their performance, is the primary obstacle to the integration of semantics into practical applications.

To overcome this issue and make high-quality semantic knowledge accessible to a broader audience in multilingual applications, we present the Universal Semantic Annotator (USeA), the first unified API for three core tasks in multilingual NLU:

- Word Sense Disambiguation (WSD): the task of selecting the most appropriate sense for a word in context from a predefined sense inventory;
- Semantic Role Labeling (SRL): the task of extracting the predicate-argument structures within a sentence, also known as a form of shallow semantic parsing;
- Semantic Parsing (SP): the task of representing a text using a formal representation, usually a graph of concepts connected by semantic relations.

For each of the above tasks, USeA transparently employs state-of-the-art models that work in 100 languages and that can be accessed through a unified API. This will ease the integration of NLU models into NLP pipelines (also for low-resource languages), enabling them to exploit explicit semantic information to improve their performance.¹

¹USeA is available for non-commercial purposes at

2. Universal Semantic Annotator

USeA – pronounced u · see – is the first unified API to provide high-quality annotations in 100 languages for three key tasks of NLU: Word Sense Disambiguation, Semantic Role Labeling, and Semantic Parsing. This Section describes how USeA tackles each of these tasks using recent high-performance systems.

2.1. Word Sense Disambiguation

Task overview. The task of WSD consists in associating a word in context with its most appropriate sense chosen from a predetermined sense inventory (Navigli, 2009; Bevilacqua et al., 2021b). Discerning the meaning of a word in context is often considered a fundamental step in enabling machine understanding of text (Navigli, 2018): indeed, a word can convey different meanings depending on the context it appears in (Camacho-Collados and Pilehvar, 2018). Over the years, there has been steady progress in this area, so much so that recent approaches (Bevilacqua and Navigli, 2020; Conia and Navigli, 2021; Barba et al., 2021a; Barba et al., 2021b) have achieved results that have come close to or even surpassed the estimated inter-annotator agreement on gold standard benchmarks for English WSD (Raganato et al., 2017), even though recent studies have shown that there is still much work to be done (Maru et al., 2022), especially in multilingual and cross-lingual WSD (Pasini, 2020; Pasini et al., 2021).

Methodology. Unfortunately, ready-to-use tools for automatic WSD have not been able to keep up with the pace of progress of the research community: currently available prepackaged systems are still based purely on graph-based heuristics, such as Babelfy (Moro et al., 2014) and SyntagRank (Scozzafava et al., 2020), or pre-neural techniques, such as SVM-based systems (Papandrea et al., 2017). USeA, instead, introduces an end-to-end WSD system that is based on recently proposed state-of-the-art approaches (Conia and Navigli, 2021; Orlando et al., 2021) and, differently from the prepackaged systems described above, is built on top of a Transformer-based (Vaswani et al., 2017) language model. Nevertheless, employing a multilingual model architecture is just the first step towards multilingual WSD; crucially, thanks to BabelNet 5² (Navigli et al., 2021), a popular multilingual encyclopedic dictionary and semantic network, our system – and therefore USeA – is able to perform WSD in 100 languages.

Implementation. The neural architecture of our WSD system is built on top of the “base” version of XLM-RoBERTa (Conneau et al., 2020), a Transformer-based multilingual language model pretrained in an unsupervised fashion on massive amounts of unlabeled text in multiple languages.

<https://github.com/SapienzaNLP/usea>.

²BabelNet is freely available to use and download for research purposes at <https://babelnet.org>.

Given a word in context, our WSD model, i) builds a contextualized representation of the word using the hidden states coming from the last four layers of XLM-RoBERTa, ii) applies a non-linear transformation to obtain sense-specific word representations, and, iii) computes the output score distribution over the set of candidate senses for the input word, where the candidate set is defined by BabelNet 5.

2.2. Semantic Role Labeling

Task overview. SRL is often defined informally as the task of automatically answering the question “*Who did What, to Whom, Where, When, and How?*” (Márquez et al., 2008). Therefore, its output, which describes the predicate-argument relations that can be inferred from a given sentence, is commonly regarded as a form of shallow semantic parse. Similarly to WSD, the area of SRL has benefited from numerous advances that have resulted in robust end-to-end systems that are also able to perform the task with strong performance in multiple languages (He et al., 2019; Conia and Navigli, 2020).

Methodology. For USeA, we develop and encapsulate an SRL model that falls within the broad category of end-to-end systems, tackling the whole SRL pipeline – predicate identification, predicate sense disambiguation, argument identification and argument classification – in a single forward pass. Differently from other prepackaged SRL systems, such as InVeRo and AllenNLP’s SRL demo, USeA is based on a multilingual neural model (Conia et al., 2021a; Conia et al., 2021b) which is able to perform state-of-the-art SRL not only across-languages, but also using different and heterogeneous linguistic inventories, namely, the English PropBank (Palmer et al., 2005), the Chinese PropBank (Xue, 2008), AnCora (Taulé et al., 2008), and VerbAtlas (Di Fabio et al., 2019), *inter alia*.

Implementation. Similarly to the WSD module, our SRL model is also based on XLM-RoBERTa. Thus, given an input sentence the system, i) builds a sequence of contextualized word representations using the hidden states coming from the last four layers of the underlying language model, ii) identifies the predicates in the sentence, iii) disambiguates each identified predicate according to each supported linguistic inventory, and iv) for each disambiguated predicate, identifies its arguments and assigns a semantic role to every predicate-argument pair.

2.3. Semantic Parsing

Task overview. Semantic Parsing can be seen as “the task of mapping natural language sentences into *complete formal meaning representations*” (Kate and Wong, 2010). Here, we focus on Semantic Parsing with Abstract Meaning Representation (Banarescu et al., 2013) – often called AMR parsing – which is concerned with encoding the meaning of a sentence in an abstract, domain-independent, form, as opposed to ex-

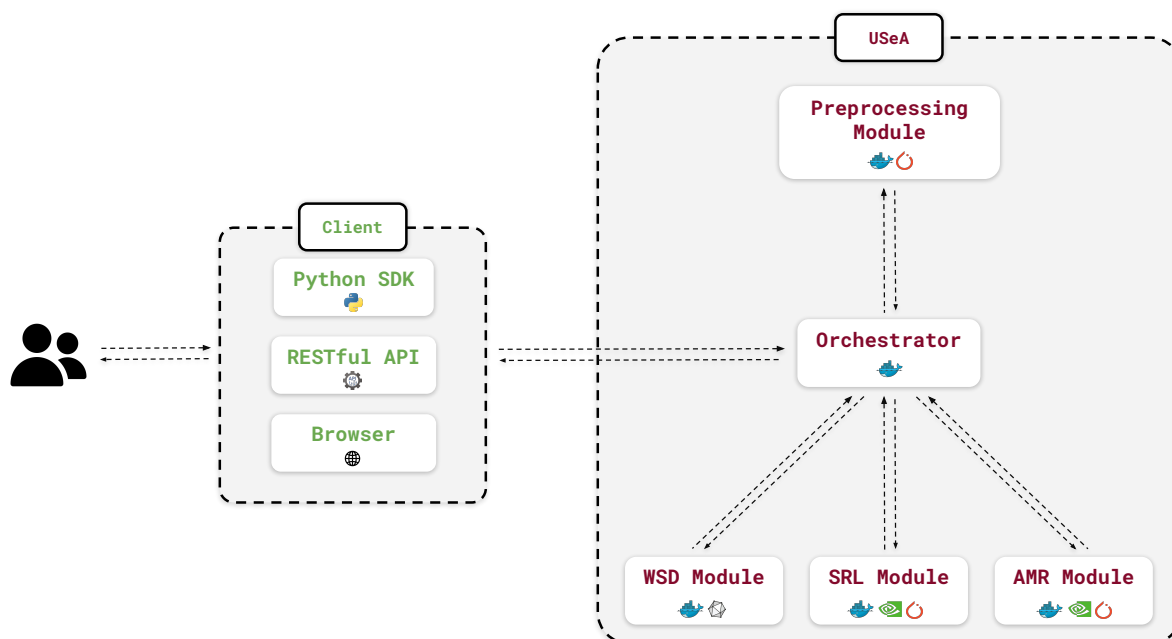


Figure 1: An overview of how USeA works: i) a user simply sends a raw-text document to the USeA server and receives the corresponding semantic information (word senses, semantic role labels, semantic parses). Inside the server, the Orchestrator Module takes care of processing the input text using other task-specific modules (preprocessing, WSD, SRL and AMR Parsing).

ecutable semantic parsing, e.g., text-to-SQL. In AMR, the semantics of a sentence is represented as a rooted directed acyclic graph, where the nodes represent concepts and the edges represent their semantic relations (Banarescu et al., 2013). The wellspring of information that AMR graphs provide has resulted in promising improvements in a variety of downstream tasks, therefore, it is key to have AMR parsers that produce structures with high quality and in as many languages as possible.

Methodology. While early work in AMR parsing was based heavily on graph neural networks, recent studies have proposed employing neural-based sequence-to-sequence models (Bevilacqua et al., 2021a; Biloshmi et al., 2021a). Therefore, we follow this line of research and propose a sequence-to-sequence model for text-to-AMR parsing with the focus on supporting 100 languages. Given an input sentence, our sequence-to-sequence model is trained to generate a linearized version of its corresponding AMR graph representation.

Implementation. As mentioned above, our AMR Parsing system is based on a sequence-to-sequence model, namely, SPRING (Biloshmi et al., 2021a), which we extended to support multiple languages. SPRING is a sequence-to-sequence Transformer-based model that operates as a *parser* by “translating” an input sentence into a linearized AMR graph. Unfortunately, SPRING is an English-only model, so, in order to achieve the goal of supporting AMR Parsing in 100 languages, we modify its original architecture by re-

placing the underlying pretrained language model from BART (Lewis et al., 2020) to mT5 (Xue et al., 2021).

3. USeA’s Pipeline and Infrastructure

To facilitate the integration of lexical and sentence-level semantic knowledge into real world applications, USeA offers a full end-to-end multilingual pipeline for WSD, SRL and AMR Parsing. The pipeline of USeA is organized into five self-contained modules – data preprocessing, WSD, SRL, AMR Parsing, and orchestration – that are transparent to the end user, as shown in Figure 1. In this Section, we provide more details about our pipeline and its infrastructure, focusing on the orchestration and preprocessing modules.

Orchestration. The orchestrator module is at the core of USeA’s pipeline, as it represents the entry point for the semantic API, i.e., each user request is first received by the orchestrator module, which then takes care of submitting other task-specific requests to the other modules in the infrastructure. Being an end-to-end system, the end user is only required to send raw text to our service. However, under the hood, the orchestrator module first sends the input text to the data preprocessing module in order to obtain the information required by the WSD, SRL and AMR Parsing modules. The annotations obtained from the three semantic modules are then combined and sent back to the user.

Data Preprocessing. One of the main advantages of our system is that it is fully self-contained. Indeed,

	English datasets						Multilingual datasets		
	SE2	SE3	SE07	SE13	SE15	ALL	SE13	SE15	XL-WSD
Moro et al. (2014)	67.0	63.5	51.6	66.4	70.3	65.5	65.6	–	52.9
Papandrea et al. (2017)	73.8	70.8	64.2	67.2	71.5	–	–	–	–
Scozzafava et al. (2020)	71.6	72.0	59.3	72.2	75.8	71.7	73.2	66.2	57.7
USeA _{WSD}	77.8	76.0	72.1	77.7	81.5	77.5	76.8	73.0	66.2

Table 1: English WSD results in F_1 scores on Senseval-2 (SE2), Senseval-3 (SE3), SemEval-2007 (SE07), SemEval-2013 (SE13), SemEval-2015 (SE15), and the concatenation of the datasets (ALL). We also include results on multilingual WSD in SemEval-2013 (DE, ES, FR, IT), SemEval-2015 (IT, ES), and XL-WSD (average over 17 languages, English excluded).

	Catalan	Czech	German	English	Spanish	Chinese
AllenNLP’s SRL demo	–	–	–	86.5	–	–
InVeRo	–	–	–	86.2	–	–
USeA _{SRL}	83.3	85.9	87.0	86.8	81.8	84.9

Table 2: Comparison between USeA and other recent automatic tools for SRL. F_1 scores on argument labeling with pre-identified predicates on the CoNLL-2012 English test set and the CoNLL-2009 test sets converted from dependency-based to span-based.

	AMR 3.0		AMR 2.0				
	EN		EN	DE	ES	IT	ZH
Lyu et al. (2021)	75.8		76.8	–	–	–	–
Zhou et al. (2021)	81.2		82.8	–	–	–	–
SPRING (Bevilacqua et al., 2021a)	83.0		84.5	–	–	–	–
Procopio et al. (2021)	80.0		81.7	54.8	60.4	63.6	47.8
USeA _{AMR-Parsing}	80.9		81.3	58.8	61.2	60.1	45.3

Table 3: SMATCH scores obtained by USeA compared with recent literature on AMR 3.0 (English) and AMR 2.0 (English and Multilingual).

the linguistic information needed by the other modules, e.g., part-of-speech tagging for WSD, is obtained through a data preprocessing module so that users do not have to worry about setting up complex pipelines themselves. In general, our preprocessing module takes care of producing all that is usually needed by other NLP systems, e.g., language identification, document splitting, tokenization, lemmatization, and part-of-speech tagging. In order to support as many languages as possible while keeping low hardware requirements, our preprocessing module is built around Trankit (Nguyen et al., 2021) and supports 100 languages with a single preprocessing model.

4. Experiments and Results

In this Section, we provide an overview of the results achieved by each USeA model on its corresponding task. While there is no prepackaged tool that offers automatic annotations for all three of the tasks we take into account, we report how USeA fares in comparison with the literature of each task.

Results in WSD. We compare our system with other prepackaged tools for automatic WSD on a set of

gold standard benchmarks for English (Raganato et al., 2017) and multilingual (Pasini et al., 2021) all-words WSD, for a total 17 languages. The results reported in Table 1 show that USeA outperforms its competitors by a wide margin, especially in multilingual WSD (+8.5% in F_1 Score on XL-WSD).

Results in SRL. We report the performance of our SRL system on two gold standard benchmarks for SRL, CoNLL-2009³ (Hajič et al., 2009) and CoNLL-2012 (Pradhan et al., 2012), covering six languages, namely, Catalan, Chinese, Czech, English, German and Spanish. We highlight that, not only does USeA provide state-of-the-art annotations for English SRL (+0.3% and +0.5% in F_1 score compared to InVeRo and AllenNLP’s SRL demo, respectively), but it is also the first ready-to-use tool to offer automatic annotations in other languages, as shown in Table 2.

³The CoNLL-2009 dataset was originally intended for dependency-based SRL. We convert dependency-based annotations to span-based annotations using the gold syntactic trees provided in the dataset.

Results in AMR Parsing. Finally, we evaluate the performance of our AMR Parsing system on the gold standard dataset of AMR 2.0⁴ and its extension, AMR 3.0⁵, which, to the best of our knowledge, are the largest datasets with gold AMR graphs. Furthermore, we also evaluate our system in a cross-lingual setting using *Abstract Meaning Representation 2.0 - Four Translations* (Damonte and Cohen, 2020), a corpus that contains the translations into Chinese (ZH), German (DE), Italian (IT) and Spanish (ES) of the sentences in the test set of AMR 2.0. Even though the AMR system we developed for USeA makes use of a multilingual language model (mT5) to support 100 languages, it is still competitive with the systems recently proposed by Bevilacqua et al. (2021a), which takes advantage of an English-only language model (Lewis et al., 2020, BART), and Procopio et al. (2021), which employs a multilingual language model that, however, only supports 25 languages (Liu et al., 2020, mBART), as shown in Table 3.

5. Conclusion

In this paper, we introduced Universal Semantic Annotator (USeA), the first unified set of APIs for automatically annotating text with explicit semantic knowledge in 100 languages. With USeA, our main objective is to re-engineer and improve state-of-the-art systems that are currently available only to a narrow group of users and, instead, make them more accessible to a broader audience, including researchers who may be interested in taking advantage of explicit semantics in their research areas and language learners who may like to ease their study with automatic tools.

6. Acknowledgments

The authors gratefully acknowledge the support of the ERC Consolidator Grant MOUSSE No. 726487 and the European Language Grid project No. 825627 (Universal Semantic Annotator, USeA) under the European Union’s Horizon 2020 research and innovation programme.



7. Bibliographical References

Banarescu, L., Bonial, C., Cai, S., Georgescu, M., Griffitt, K., Hermjakob, U., Knight, K., Koehn, P., Palmer, M., and Schneider, N. (2013). Abstract Meaning Representation for sembanking. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 178–186, Sofia, Bulgaria. Association for Computational Linguistics.

⁴<https://catalog.ldc.upenn.edu/LDC2017T10>

⁵<https://catalog.ldc.upenn.edu/LDC2020T02>

- Barba, E., Pasini, T., and Navigli, R. (2021a). ESC: Redesigning WSD with extractive sense comprehension. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4661–4672, Online, June. Association for Computational Linguistics.
- Barba, E., Procopio, L., and Navigli, R. (2021b). ConSeC: Word Sense Disambiguation as continuous sense comprehension. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1492–1503, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Bender, E. M. and Koller, A. (2020). Climbing towards NLU: On meaning, form, and understanding in the age of data. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5185–5198, Online. Association for Computational Linguistics.
- Bevilacqua, M. and Navigli, R. (2020). Breaking through the 80% glass ceiling: Raising the state of the art in Word Sense Disambiguation by incorporating knowledge graph information. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2854–2864, Online. Association for Computational Linguistics.
- Bevilacqua, M., Blloshmi, R., and Navigli, R. (2021a). One SPRING to rule them both: Symmetric AMR semantic parsing and generation without a complex pipeline. *Proceedings of AAAI*, 35(14):12564–12573.
- Bevilacqua, M., Pasini, T., Raganato, A., and Navigli, R. (2021b). Recent trends in Word Sense Disambiguation: A survey. In *Proceedings of IJCAI-21*, pages 4330–4338.
- Blloshmi, R., Bevilacqua, M., Fabiano, E., Caruso, V., and Navigli, R. (2021a). SPRING Goes Online: End-to-End AMR Parsing and Generation. In *Proceedings of EMNLP*, pages 134–142.
- Blloshmi, R., Pasini, T., Campolungo, N., Banarjee, S., Navigli, R., and Pasi, G. (2021b). IR like a SIR: Sense-enhanced information retrieval for multiple languages. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Punta Cana, Dominican Republic.
- Camacho-Collados, J. and Pilehvar, M. T. (2018). From word to sense embeddings: A survey on vector representations of meaning. *J. Artif. Intell. Res.*, 63:743–788.
- Christensen, J., Mausam, Soderland, S., and Etzioni, O. (2010). Semantic Role Labeling for open information extraction. In *Proceedings of the NAACL HLT 2010 First International Workshop on Formalisms and Methodology for Learning by Reading*, pages 52–60, Los Angeles, California. Association for Computational Linguistics.

- Conia, S. and Navigli, R. (2020). Bridging the gap in multilingual Semantic Role Labeling: A language-agnostic approach. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1396–1410, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Conia, S. and Navigli, R. (2021). Framing Word Sense Disambiguation as a multi-label problem for model-agnostic knowledge integration. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3269–3275, Online. Association for Computational Linguistics.
- Conia, S., Bacciu, A., and Navigli, R. (2021a). Unifying cross-lingual Semantic Role Labeling with heterogeneous linguistic resources. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 338–351, Online. Association for Computational Linguistics.
- Conia, S., Orlando, R., Brignone, F., Cecconi, F., and Navigli, R. (2021b). InVeRo-XL: Making cross-lingual Semantic Role Labeling accessible with Intelligible Verbs and Roles. In *Proceedings of EMNLP*.
- Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., Grave, E., Ott, M., Zettlemoyer, L., and Stoyanov, V. (2020). Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Damonte, M. and Cohen, S. (2020). Abstract Meaning Representation 2.0 - Four Translations. Web Download, Philadelphia: Linguistic Data Consortium.
- Di Fabio, A., Conia, S., and Navigli, R. (2019). VerbAtlas: A novel large-scale verbal semantic resource and its application to Semantic Role Labeling. 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)*, pages 627–637, Hong Kong, China. Association for Computational Linguistics.
- Gupta, S. and Malik, J. (2015). Visual Semantic Role Labeling. *arXiv*, abs/1505.04474.
- Hajič, J., Ciaramita, M., Johansson, R., Kawahara, D., Martí, M. A., Màrquez, L., Meyers, A., Nivre, J., Padó, S., Štěpánek, J., Straňák, P., Surdeanu, M., Xue, N., and Zhang, Y. (2009). The CoNLL-2009 shared task: Syntactic and semantic dependencies in multiple languages. In *Proceedings of the Thirteenth Conference on Computational Natural Language Learning (CoNLL 2009): Shared Task*, pages 1–18, Boulder, Colorado. Association for Computational Linguistics.
- Hardy, H. and Vlachos, A. (2018). Guided neural language generation for abstractive summarization using Abstract Meaning Representation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 768–773, Brussels, Belgium. Association for Computational Linguistics.
- He, L., Lewis, M., and Zettlemoyer, L. (2015). Question-answer driven Semantic Role Labeling: Using natural language to annotate natural language. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 643–653, Lisbon, Portugal. Association for Computational Linguistics.
- He, S., Li, Z., and Zhao, H. (2019). Syntax-aware multilingual Semantic Role Labeling. 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)*, pages 5350–5359, Hong Kong, China. Association for Computational Linguistics.
- Kate, R. J. and Wong, Y. W. (2010). Semantic parsing: The task, the state of the art and the future. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts*, page 6, Uppsala, Sweden. Association for Computational Linguistics.
- Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., and Zettlemoyer, L. (2020). BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Liao, K., Lebanoff, L., and Liu, F. (2018). Abstract Meaning Representation for multi-document summarization. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1178–1190, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Liu, Y., Gu, J., Goyal, N., Li, X., Edunov, S., Ghazvininejad, M., Lewis, M., and Zettlemoyer, L. (2020). Multilingual denoising pre-training for neural machine translation. *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Lyu, C., Cohen, S. B., and Titov, I. (2021). A differentiable relaxation of graph segmentation and alignment for AMR parsing. In *Proceedings of EMNLP*, pages 9075–9091.
- Marcheggiani, D., Bastings, J., and Titov, I. (2018). Exploiting semantics in neural machine translation with graph convolutional networks. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 486–492, New Orleans, Louisiana. Association for Computational Linguistics.

- Màrquez, L., Carreras, X., Litkowski, K. C., and Stevenson, S. (2008). Semantic Role Labeling: An introduction to the special issue. *Computational Linguistics*, 34(2):145–159.
- Maru, M., Conia, S., Bevilacqua, M., and Navigli, R. (2022). Nibbling at the hard core of Word Sense Disambiguation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL 2022)*, Dublin, Ireland, May. Association for Computational Linguistics.
- Moro, A., Raganato, A., and Navigli, R. (2014). Entity linking meets Word Sense Disambiguation: A unified approach. *Transactions of the Association for Computational Linguistics*, 2:231–244.
- Navigli, R., Bevilacqua, M., Conia, S., Montagnini, D., and Ceconi, F. (2021). Ten years of BabelNet: A survey. In *Proceedings of IJCAI-21*, pages 4559–4567.
- Navigli, R. (2009). Word Sense Disambiguation: A survey. *ACM Comput. Surv.*, 41(2).
- Navigli, R. (2018). Natural Language Understanding: Instructions for (present and future) use. In Jérôme Lang, editor, *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 5697–5702. ijcai.org.
- Nguyen, M. V., Lai, V. D., Pourn Ben Veyseh, A., and Nguyen, T. H. (2021). Trankit: A light-weight transformer-based toolkit for multilingual natural language processing. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 80–90, Online. Association for Computational Linguistics.
- Orlando, R., Conia, S., Brignone, F., Ceconi, F., and Navigli, R. (2021). AMuSE-WSD: An all-in-one multilingual system for easy Word Sense Disambiguation. In *Proceedings of EMNLP*.
- Otter, D. W., Medina, J. R., and Kalita, J. K. (2021). A survey of the usages of deep learning for natural language processing. *IEEE Transactions on Neural Networks and Learning Systems*, 32(2):604–624.
- Palmer, M., Gildea, D., and Kingsbury, P. (2005). The Proposition Bank: An annotated corpus of semantic roles. *Computational Linguistics*, 31(1):71–106.
- Papandrea, S., Raganato, A., and Delli Bovi, C. (2017). SupWSD: A flexible toolkit for supervised Word Sense Disambiguation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 103–108, Copenhagen, Denmark. Association for Computational Linguistics.
- Pasini, T., Raganato, A., and Navigli, R. (2021). XL-WSD: An extra-large and cross-lingual evaluation framework for Word Sense Disambiguation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(15):13648–13656, May.
- Pasini, T. (2020). The knowledge acquisition bottleneck problem in multilingual word sense disambiguation. In Christian Bessiere, editor, *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 4936–4942. ijcai.org.
- Pradhan, S., Moschitti, A., Xue, N., Uryupina, O., and Zhang, Y. (2012). CoNLL-2012 shared task: Modeling multilingual unrestricted coreference in OntoNotes. In *Joint Conference on EMNLP and CoNLL - Shared Task*, pages 1–40, Jeju Island, Korea. Association for Computational Linguistics.
- Procopio, L., Tripodi, R., and Navigli, R. (2021). SGL: Speaking the graph languages of semantic parsing via multilingual translation. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 325–337, Online. Association for Computational Linguistics.
- Raganato, A., Camacho-Collados, J., and Navigli, R. (2017). Word Sense Disambiguation: A unified evaluation framework and empirical comparison. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 99–110, Valencia, Spain. Association for Computational Linguistics.
- Raganato, A., Scherrer, Y., and Tiedemann, J. (2019). The MuCoW test suite at WMT 2019: Automatically harvested multilingual contrastive Word Sense Disambiguation test sets for machine translation. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 470–480, Florence, Italy. Association for Computational Linguistics.
- Ramakrishnan, G., Jadhav, A., Joshi, A., Chakrabarti, S., and Bhattacharyya, P. (2003). Question answering via Bayesian inference on lexical relations. In *Proceedings of the ACL 2003 Workshop on Multilingual Summarization and Question Answering*, pages 1–10, Sapporo, Japan. Association for Computational Linguistics.
- Sadhu, A., Gupta, T., Yatskar, M., Nevatia, R., and Kembhavi, A. (2021). Visual Semantic Role Labeling for video understanding. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Sciozzafava, F., Maru, M., Brignone, F., Torrisi, G., and Navigli, R. (2020). Personalized PageRank with syntagmatic information for multilingual Word Sense Disambiguation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 37–46, Online. Association for Computational Linguistics.
- Song, L., Gildea, D., Zhang, Y., Wang, Z., and Su, J. (2019). Semantic neural machine translation using AMR. *Transactions of the Association for Compu-*

- tational Linguistics*, 7:19–31.
- Taulé, M., Martí, M. A., and Recasens, M. (2008). AnCorra: Multilevel annotated corpora for Catalan and Spanish. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*, Marrakech, Morocco. European Language Resources Association (ELRA).
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. (2017). Attention is all you need. In Isabelle Guyon, et al., editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Xue, L., Constant, N., Roberts, A., Kale, M., Al-Rfou, R., Siddhant, A., Barua, A., and Raffel, C. (2021). mT5: A massively multilingual pre-trained text-to-text transformer. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.
- Xue, N. (2008). Labeling Chinese predicates with semantic roles. *Computational Linguistics*, 34(2):225–255.
- Yatskar, M., Zettlemoyer, L. S., and Farhadi, A. (2016). Situation recognition: Visual Semantic Role Labeling for image understanding. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 5534–5542. IEEE Computer Society.
- Zhou, J., Naseem, T., Fernandez Astudillo, R., and Florian, R. (2021). AMR parsing with action-pointer transformer. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5585–5598, Online. Association for Computational Linguistics.