A Tour of Explicit Multilingual Semantics: Word Sense Disambiguation, Semantic Role Labeling and Semantic Parsing

Roberto Navigli, Edoardo Barba, Simone Conia
Sapienza NLP Group
Sapienza University of Rome
first.lastname@uniroma1.it

Rexhina Blloshmi*
Amazon Alexa AI
Berlin, Germany
blloshmi@amazon.de

Abstract

The recent advent of modern pretrained language models has sparked a revolution in Natural Language Processing (NLP), especially in multilingual and cross-lingual applications. Today, such language models have become the *de facto* standard for providing rich input representations to neural systems, achieving unprecedented results in an increasing range of benchmarks. However, questions that often arise are: firstly, whether current language models are, indeed, able to capture explicit, symbolic meaning; secondly, if they are, to what extent; thirdly, and perhaps more importantly, whether current approaches are capable of scaling across languages.

In this cutting-edge tutorial, we will review recent efforts that have aimed at shedding light on meaning in NLP, with a focus on three key open problems in lexical and sentence-level semantics: Word Sense Disambiguation, Semantic Role Labeling, and Semantic Parsing. After a brief introduction, we will spotlight how state-of-the-art models tackle these tasks in multiple languages, showing where they excel and where they fail. We hope that this tutorial will broaden the audience interested in multilingual semantics and inspire researchers to further advance the field.

1 Tutorial Description and Relevance

Over the past few years, the field of Natural Language Processing (NLP) has witnessed tremendous growth, mainly thanks to the increasingly wide availability of modern pretrained language models, such as ELMo (Peters et al., 2018), BERT (Devlin et al., 2019), and BART (Lewis et al., 2020), which have enabled unprecedented results in a broad range of tasks, from Neural Machine Translation to Question Answering, Information Retrieval and Text Summarization, *inter alia*. However, important questions that naturally arise when looking at the recent impressive gains in the field are whether such powerful language models learn to encode *semantic knowledge* and, if they are, to what extent. More importantly, the escalating interest in multilingual NLP demands approaches that are able to identify and transfer semantics across a multitude of languages, especially those for which there is a scarce amount of data available.

In this tutorial, we will review recent studies in lexical and sentence semantics, paying special attention to state-of-the-art approaches and how they tackle multilinguality in three fundamental tasks for Natural Language Understanding (NLU): Word Sense Disambiguation (WSD), Semantic Role Labeling (SRL) and Semantic Parsing (SP). In addition to an introduction to multilingual NLU, for each task we will provide, i) a gentle introduction, ii) an overview of the inventories and resources most commonly adopted, iii) an outline of current approaches with a particular focus on multilinguality and cross-linguality in order to understand their strengths and shortcomings, and also pointing to promising directions for future work. Although there have been previous tutorials on Semantics in NLP, especially on SP (Lopez and Gilroy, 2018; Gardner et al., 2018; Koller et al., 2019), our tutorial will, instead, focus on the challenges of multilinguality and cross-linguality and how recent approaches based on pretrained language models tackle them.

Despite the increasing performance of huge language models in NLU tasks, recent studies have demonstrated that the integration of explicit semantics into deep learning techniques is beneficial not only in terms of performances (Levine et al., 2020), but also interpretability (Wiedemann et al., 2019) and cross-lingual transfer (Blloshmi et al., 2020).

2 Tutorial Structure and Contents

The tutorial will be structured in a bottom-up fashion: participants will be introduced to multilingual

*Proposal written while at Sapienza University of Rome.

Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing: Tutorial Abstracts, pages 35–43
November 20-21, 2022.

semantics, first at the lexical level with WSD, and then at the sentence level with SRL and SP, highlighting the most effective approaches to date, but also their weaknesses and future directions to address these.

2.1 Word Sense Disambiguation (WSD)

The tutorial will start with WSD as the lowest level of semantic abstraction. Its objective is to assign the most appropriate sense to a word in context from a finite set of possible choices (Navigli, 2009), which usually come from predefined sense inventories. Although, at a first glance, WSD may seem a simple task to a human, it has proven to be extremely challenging for machines. Indeed, depending on the sense inventory of choice, different linguistic phenomena may make the task difficult to tackle with standard classification techniques. Nonetheless, being able to link raw text to knowledge bases is fundamental in NLP (McCoy et al., 2019; Bender and Koller, 2020), bringing benefits in several fields, such as Machine Translation (Liu et al., 2018; Pu et al., 2018; Campolungo et al., 2022), Information Extraction (Delli Bovi et al., 2015), and Information Retrieval (Blloshmi et al., 2021b). We will start with an introduction to the task, presenting its most common formulation along with the challenges it poses. Then, we will describe state-of-the-art systems, highlighting their core contributions. Finally, we will conclude by presenting open challenges in multilingual WSD.

Resources for WSD. We will first present the standard resources currently in use for WSD, starting with WordNet (Miller et al., 1990), i.e., the most widely used sense inventory, and Open Multilingual Wordnet (Bond, 2011) and BabelNet (Navigli and Ponzetto, 2012; Navigli et al., 2021), two multilingual extensions of WordNet.

Current approaches in WSD. After the initial success of purely data-driven neural models in WSD (Yuan et al., 2016), subsequent approaches started to leverage information coming from knowledge bases in addition to standard training datasets (Huang et al., 2019; Bevilacqua and Navigli, 2020). We will put a special focus on state-of-the-art systems that rely on relational knowledge (Bevilacqua and Navigli, 2020) and sense definitions as additional knowledge (Blevins and Zettlemoyer, 2020; Barba et al., 2021a,b). We will explain how these approaches are data-efficient and why they are important, especially for low-resource languages.

2.2 Semantic Role Labeling (SRL)

While WSD is concerned with lexical-level meaning, SRL (Gildea and Jurafsky, 2000) investigates sentence-level semantics and is usually described informally as the task of automatically answering the question *"Who did What to Whom, Where, When, and How?"* (Màrquez et al., 2008). More precisely, its objective is to extract the predicate-argument structure of a sentence and, therefore, it is considered by some as a form of shallow Semantic Parsing. Over the years, SRL has been proven to be beneficial in several tasks, such as Question Answering (Shen and Lapata, 2007), Machine Translation (Marcheggiani et al., 2018), Video Understanding (Sadhu et al., 2021), and data augmentation (Ross et al., 2022). Following a general introduction to SRL, the tutorial will highlight some key details about the most popular predicate-argument structure inventories for SRL, the salient characteristics of current state-of-the-art systems, and why everything becomes more complex when trying to tackle multilingual and cross-lingual SRL.

Inventories for SRL. The tutorial will overview the main challenges that current predicate-argument structure inventories pose for multilingual and cross-lingual SRL, with particular focus on PropBank-style inventories (Palmer et al., 2005; Xue, 2008; Jindal et al., 2022), FrameNet (Baker et al., 1998) and VerbAtlas (Di Fabio et al., 2019).

Current approaches in SRL. Given its close ties with syntax, over the years one of the main distinctions between proposed approaches is whether they have chosen to rely on syntactic features (He et al., 2019; Marcheggiani and Titov, 2020; Conia and Navigli, 2020), or not (Marcheggiani et al., 2017; Cai et al., 2018). The tutorial will briefly cover the advantages and disadvantages of relying on syntax in multilingual SRL, but also highlight annotation projection techniques for cross-lingual SRL (Akbik et al., 2015; Daza and Frank, 2020), and how recent trends in multi-task learning (Conia et al., 2021) and generation (Blloshmi et al., 2021a; Paolini et al., 2021; Conia et al., 2022) are going beyond traditional approaches, hinting at new directions in SRL.

2.3 Semantic Parsing (SP)

Finally, the tutorial will bring participants to a higher level of semantic abstraction: SP, indeed, may be seen as "the task of mapping natural lan-

guage sentences into *complete formal meaning representations* which a computer can execute" (Kate and Wong, 2010). Here we focus on formalisms that aim at encoding text in an abstract form that captures aspects of meaning – as opposed to executable formalisms for SP – that can be reusable in various scenarios, thus being domain independent. Indeed, SP formalisms have been successfully integrated into numerous downstream applications, such as Machine Translation (Song et al., 2019), Text Summarization (Hardy and Vlachos, 2018), Human-Robot Interaction (Bonial et al., 2020) and Question Answering (Kapanipathi et al., 2021). Nevertheless, research in SP has mainly focused on English, with only a handful of attempts in other languages.

Formalisms for SP. Over the years, various different formalisms have been proposed to encode semantic structures. We will first overview the most popular formalisms, such as Elementary Dependency Structures (Oepen and Lønning, 2006, EDS), Prague Tectogrammatical Graphs (Hajič et al., 2012, PTG), Universal Conceptual Cognitive Annotation (Abend and Rappoport, 2013, UCCA), Universal Decompositional Semantics (White et al., 2016, UDS), with a main focus on Abstract Meaning Representation (Banarescu et al., 2013) and BabelNet Meaning Representation, its fully-semantic extension (Martínez Lorenzo et al., 2022).

Current approaches in SP. SP is receiving ever growing attention that has led to numerous approaches of different flavors. Indeed, the advantages and disadvantages of parser types are variable across different formalisms. We will focus on two categories of approaches: *graph-based* ones (Zhang et al., 2019; Cai and Lam, 2020), that consist of transducing natural utterances into graphs, and *sequence-to-sequence* ones, that produce linearized graph structures for a given input text (Ge et al., 2019; Bevilacqua et al., 2021a). Due to the recent development of encoder-decoder pretrained architectures, sequence-to-sequence approaches to SP are emerging as the best-performing methods, not only in English (Bevilacqua et al., 2021a), but also in other languages (Procopio et al., 2021b).

3 Type, Prerequisites and Audience

This is a **cutting-edge tutorial**. State-of-the-art approaches for three key areas of multilingual lexical and sentence semantics will be presented, and some of them will be discussed in detail. We expect **80-120 attendees** from different fields as the barriers to entry will be low:

- **Math prerequisites:** Linear algebra, e.g., matrix operations, linear/non-linear functions.
- **Machine Learning prerequisites:** General concepts of classification, e.g., token classification, sequence labeling, sequence-to-sequence.
- **NLP prerequisites:** High-level notions about pretrained language models.

4 Reading List

Recommended work to read before the tutorial:

- Bevilacqua et al. (2021b): a survey on recent trends in **WSD**;
- Blevins and Zettlemoyer (2020) and Barba et al. (2021a): two recent **WSD** systems that take advantage of sense definitions;
- Màrquez et al. (2008) and Hajic et al. (2009): an introduction to **SRL** and the largest gold benchmark for multilingual **SRL**;
- He et al. (2019) and Conia et al. (2021): two recent approaches to multilingual **SRL**, a syntax-aware and a syntax-agnostic one;
- Koller et al. (2019) and Oepen et al. (2020): tutorial on recent work and shared task on **SP**;
- Banarescu et al. (2013) and Bevilacqua et al. (2021a): the introduction to the AMR formalism for **SP** and a state-of-the-art system for AMR parsing and generation;

5 Tutorial Outline (3h)

Part 0: Introduction (10 minutes). Introduction, motivation, goals, how the tutorial is organized.

Part 1: WSD (40 minutes)

- Introduction to WSD, formulation, examples;
- Sense inventories for WSD: WordNet, Open Multilingual WordNet and BabelNet;
- Current approaches in multilingual WSD: purely data-driven vs. knowledge-enhanced supervision; going beyond sense inventories.

QA & Break (10 minutes)

Part 2: SRL (40 minutes)

- Introduction to SRL, formulation, examples;
- Predicate-argument structure inventories: the case of multilingual and cross-lingual SRL;
- Current approaches in multilingual and cross-lingual SRL: syntax-aware vs syntax-agnostic systems, annotation projection techniques, and novel directions.

QA & Break (10 minutes)

Part 3: SP (40 minutes)

- Introduction to SP, formulation, examples;
- Main formalisms for SP;
- Current approaches in cross-lingual SP: annotation projection, data augmentation via translation, generation.

QA & Break (10 minutes)

Part 4: Conclusion (20 minutes). Where to go from here, general considerations, a look to the future of explicit lexical and sentence semantics.

6 Pedagogical Material

Part 1 (WSD), Part 2 (SRL) and Part 3 (SP) will include brief hands-on sessions. These will be supported by interactive demos and Jupyter/iPython/Colab notebooks to invite participants to play with high-performance pretrained systems for WSD, SRL and SP. All material (slides, notebooks, pretrained models) will be freely available online to let discussions continue beyond the tutorial and for teaching purposes.

7 Presenters

Roberto Navigli is a Full Professor in the Department of Computer, Control and Management Engineering (DIAG) of Sapienza University of Rome, from which he also obtained his Ph.D. in Computer Science in 2007. At Sapienza he has taught courses for 4 Master's programmes (CS, CS Engineering, AI & Robotics and Data Science), including NLP. He has been a keynote speaker at more than 30 **conferences** and **workshops**, including IJCNLP, IJCAI-ECAI (early career spotlight), AMLD, SwissText+KONVENS, CLNLP, RANLP, TALN, eLex. In 2014, he co-presented a (pre-neural) tutorial on "Multilingual WSD and Entity Linking" at COLING. In 2016, he co-presented a tutorial on "Semantic Representations of Word Senses and Concepts" at ACL. He has worked and published with around **200 researchers** from all over the world in more than 200 papers in the area of NLP with a particular focus on **Natural Language Understanding** and **multilinguality**, attracting 18,000+ citations.

Rexhina Blloshmi is a Machine Learning Scientist at Amazon Alexa AI in Berlin. Her PhD focused on Semantic Parsing. She contributed in this field with several publications in AI and NLP conferences (3 EMNLP, 1 IJCAI and 3 AAAI), mainly on English and Cross-Lingual Abstract Meaning Representation and Semantic Role Labeling, but also on novel formalisms such as BabelNet Meaning Representation.

Edoardo Barba is a third-year PhD Student in NLP at Sapienza University of Rome. His research is mostly focused on Word Sense Disambiguation. He contributed to several articles regarding both state-of-the-art and data efficient systems for WSD (Barba et al., 2021a) as well as Data Augmentation techniques for Multilingual WSD (Barba et al., 2020; Procopio et al., 2021a). Teaching Assistant in 2020 and 2021 for the NLP course at Sapienza (taught in English).

Simone Conia is a third-year PhD Student in NLP at Sapienza University of Rome. His research revolves around multilingual and cross-lingual semantics, with numerous papers on WSD and SRL published at *ACL and other top-tier conferences. Simone is recipient of an **Outstanding Paper Award at NAACL-2021** for his work on cross-lingual SRL. Teaching Assistant in 2020 and 2021 for the NLP course at Sapienza (taught in English).

8 Ethics & Diversity Statement

We do not foresee any major ethical issue for the topics covered in this tutorial. We acknowledge that pretrained language models may show biases towards some stereotypes, cultures, ethnic and/or social groups: perpetrating such biases is not in our intentions. We will cover a variety of languages, including Arabic, Chinese, English, French, German, Italian, Spanish: we hope that our effort can promote new studies aimed at making lexical and sentence semantics increasingly more inclusive of lower-resource languages.

References

Omri Abend and Ari Rappoport. 2013. Universal Conceptual Cognitive Annotation (UCCA). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 228–238, Sofia, Bulgaria. Association for Computational Linguistics.

Alan Akbik, Laura Chiticariu, Marina Danilevsky, Yunyao Li, Shivakumar Vaithyanathan, and Huaiyu Zhu. 2015. Generating high quality proposition Banks for multilingual semantic role labeling. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 397–407, Beijing, China. Association for Computational Linguistics.

Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. The Berkeley FrameNet project. In *36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics, Volume 1*, pages 86–90, Montreal, Quebec, Canada. Association for Computational Linguistics.

Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 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.

Edoardo Barba, Tommaso Pasini, and Roberto Navigli. 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. Association for Computational Linguistics.

Edoardo Barba, Luigi Procopio, Niccolò Campolungo, Tommaso Pasini, and Roberto Navigli. 2020. Mulan: Multilingual label propagation for word sense disambiguation. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 3837–3844. International Joint Conferences on Artificial Intelligence Organization. Main track.

Edoardo Barba, Luigi Procopio, and Roberto Navigli. 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.

Emily M. Bender and Alexander Koller. 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.

Michele Bevilacqua, Rexhina Blloshmi, and Roberto Navigli. 2021a. One SPRING to Rule Them Both: Symmetric AMR Semantic Parsing and Generation without a Complex Pipeline. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(14):12564–12573.

Michele Bevilacqua and Roberto Navigli. 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.

Michele Bevilacqua, Tommaso Pasini, Alessandro Raganato, and Roberto Navigli. 2021b. Recent trends in word sense disambiguation: A survey. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, pages 4330–4338. ijcai.org.

Terra Blevins and Luke Zettlemoyer. 2020. Moving down the long tail of word sense disambiguation with gloss informed bi-encoders. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1006–1017, Online. Association for Computational Linguistics.

Rexhina Blloshmi, Simone Conia, Rocco Tripodi, and Roberto Navigli. 2021a. Generating senses and roles: An end-to-end model for dependency- and span-based semantic role labeling. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 3786–3793. International Joint Conferences on Artificial Intelligence Organization. Main Track.

Rexhina Blloshmi, Tommaso Pasini, Niccolò Campolungo, Somnath Banerjee, Roberto Navigli, and Gabriella Pasi. 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*, pages 1030–1041, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Rexhina Blloshmi, Rocco Tripodi, and Roberto Navigli. 2020. XL-AMR: Enabling cross-lingual AMR parsing with transfer learning techniques. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2487–2500, Online. Association for Computational Linguistics.

Francis Bond. 2011. A survey of wordnets and their licenses.

Claire Bonial, Lucia Donatelli, Mitchell Abrams, Stephanie M. Lukin, Stephen Tratz, Matthew Marge,

Ron Artstein, David Traum, and Clare Voss. 2020. Dialogue-AMR: Abstract Meaning Representation for dialogue. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 684–695, Marseille, France. European Language Resources Association.

Deng Cai and Wai Lam. 2020. AMR parsing via graph-sequence iterative inference. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1290–1301, Online. Association for Computational Linguistics.

Jiaxun Cai, Shexia He, Zuchao Li, and Hai Zhao. 2018. A full end-to-end semantic role labeler, syntactic-agnostic over syntactic-aware? In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2753–2765, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

Niccolò Campolungo, Federico Martelli, Francesco Saina, and Roberto Navigli. 2022. DiBiMT: A novel benchmark for measuring Word Sense Disambiguation biases in Machine Translation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4331–4352, Dublin, Ireland. Association for Computational Linguistics.

Simone Conia, Andrea Bacciu, and Roberto Navigli. 2021. 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.

Simone Conia, Edoardo Barba, Alessandro Scirè, and Roberto Navigli. 2022. Semantic Role Labeling meets definition modeling: Using natural language to describe predicate-argument structures. In *Findings of the Association for Computational Linguistics: EMNLP 2022*. Association for Computational Linguistics.

Simone Conia and Roberto Navigli. 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.

Angel Daza and Anette Frank. 2020. X-SRL: A parallel cross-lingual semantic role labeling dataset. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3904–3914, Online. Association for Computational Linguistics.

Claudio Delli Bovi, Luca Telesca, and Roberto Navigli. 2015. Large-scale information extraction from textual definitions through deep syntactic and semantic analysis. *Transactions of the Association for Computational Linguistics*, 3:529–543.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. 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)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Andrea Di Fabio, Simone Conia, and Roberto Navigli. 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.

Matt Gardner, Pradeep Dasigi, Srinivasan Iyer, Alane Suhr, and Luke Zettlemoyer. 2018. Neural semantic parsing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts*, pages 17–18, Melbourne, Australia. Association for Computational Linguistics.

DongLai Ge, Junhui Li, Muhua Zhu, and Shoushan Li. 2019. Modeling source syntax and semantics for neural amr parsing. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 4975–4981. International Joint Conferences on Artificial Intelligence Organization.

Daniel Gildea and Daniel Jurafsky. 2000. Automatic labeling of semantic roles. In *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics*, pages 512–520, Hong Kong. Association for Computational Linguistics.

Jan Hajic, Massimiliano Ciaramita, Richard Johansson, Daisuke Kawahara, Maria Antònia Martí, Lluís Màrquez, Adam Meyers, Joakim Nivre, Sebastian Padó, Jan Stepánek, Pavel Stranák, Mihai Surdeanu, Nianwen Xue, and Yi Zhang. 2009. The conll-2009 shared task: Syntactic and semantic dependencies in multiple languages. In *Proceedings of the Thirteenth Conference on Computational Natural Language Learning: Shared Task, CoNLL 2009, Boulder, Colorado, USA, June 4, 2009*, pages 1–18. ACL.

Jan Hajič, Eva Hajičová, Jarmila Panevová, Petr Sgall, Ondřej Bojar, Silvie Cinková, Eva Fučíková, Marie Mikulová, Petr Pajas, Jan Popelka, Jiří Semecký, Jana Šindlerová, Jan Štěpánek, Josef Toman, Zdeňka Urešová, and Zdeněk Žabokrtský. 2012. Announcing Prague Czech-English Dependency Treebank 2.0. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 3153–3160, Istanbul, Turkey. European Language Resources Association (ELRA).

Hardy Hardy and Andreas Vlachos. 2018. Guided neural language generation for abstractive summariza-

tion 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.

Shexia He, Zuchao Li, and Hai Zhao. 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.

Luyao Huang, Chi Sun, Xipeng Qiu, and Xuanjing Huang. 2019. GlossBERT: BERT for word sense disambiguation with gloss knowledge. 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 3509–3514, Hong Kong, China. Association for Computational Linguistics.

Ishan Jindal, Alexandre Rademaker, Michał Ulewicz, Ha Linh, Huyen Nguyen, Khoi-Nguyen Tran, Huaiyu Zhu, and Yunyao Li. 2022. Universal proposition bank 2.0. In *Proceedings of the Language Resources and Evaluation Conference*, pages 1700–1711, Marseille, France. European Language Resources Association.

Pavan Kapanipathi, Ibrahim Abdelaziz, Srinivas Ravishankar, Salim Roukos, Alexander Gray, Ramón Fernandez Astudillo, Maria Chang, Cristina Cornelio, Saswati Dana, Achille Fokoue, Dinesh Garg, Alfio Gliozzo, Sairam Gurajada, Hima Karanam, Naweed Khan, Dinesh Khandelwal, Young-Suk Lee, Yunyao Li, Francois Luus, Ndivhuwo Makondo, Nandana Mihindukulasooriya, Tahira Naseem, Sumit Neelam, Lucian Popa, Revanth Gangi Reddy, Ryan Riegel, Gaetano Rossiello, Udit Sharma, G P Shrivatsa Bhargav, and Mo Yu. 2021. Leveraging Abstract Meaning Representation for knowledge base question answering. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3884–3894, Online. Association for Computational Linguistics.

Rohit J. Kate and Yuk Wah Wong. 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.

Alexander Koller, Stephan Oepen, and Weiwei Sun. 2019. Graph-based meaning representations: Design and processing. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts*, pages 6–11, Florence, Italy. Association for Computational Linguistics.

Yoav Levine, Barak Lenz, Or Dagan, Ori Ram, Dan Padnos, Or Sharir, Shai Shalev-Shwartz, Amnon Shashua, and Yoav Shoham. 2020. SenseBERT: Driving some sense into BERT. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4656–4667, Online. Association for Computational Linguistics.

Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 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.

Frederick Liu, Han Lu, and Graham Neubig. 2018. Handling homographs in neural machine translation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1336–1345, New Orleans, Louisiana. Association for Computational Linguistics.

Adam Lopez and Sorcha Gilroy. 2018. Graph formalisms for meaning representations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: Tutorial Abstracts*, Melbourne, Australia. Association for Computational Linguistics.

Diego Marcheggiani, Jasmijn Bastings, and Ivan Titov. 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.

Diego Marcheggiani, Anton Frolov, and Ivan Titov. 2017. A simple and accurate syntax-agnostic neural model for dependency-based semantic role labeling. In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 411–420, Vancouver, Canada. Association for Computational Linguistics.

Diego Marcheggiani and Ivan Titov. 2020. Graph convolutions over constituent trees for syntax-aware semantic role labeling. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3915–3928, Online. Association for Computational Linguistics.

Lluís Màrquez, Xavier Carreras, Kenneth C. Litkowski, and Suzanne Stevenson. 2008. Special issue introduction: Semantic role labeling: An introduction to the special issue. *Computational Linguistics*, 34(2):145–159.

Abelardo Carlos Martínez Lorenzo, Marco Maru, and Roberto Navigli. 2022. Fully-Semantic Parsing and

Generation: the BabelNet Meaning Representation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1727–1741, Dublin, Ireland. Association for Computational Linguistics.

Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics.

George A. Miller, Richard Beckwith, Christiane D. Fellbaum, Derek Gross, and Katherine J. Miller. 1990. Introduction to wordnet: An on-line lexical database. *International Journal of Lexicography*, 3:235–244.

Roberto Navigli. 2009. Word sense disambiguation: A survey. *ACM computing surveys (CSUR)*, 41(2):1–69.

Roberto Navigli, Michele Bevilacqua, Simone Conia, Dario Montagnini, and Francesco Cecconi. 2021. Ten years of babelnet: A survey. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, pages 4559–4567. ijcai.org.

Roberto Navigli and Simone Paolo Ponzetto. 2012. Babelnet: The automatic construction, evaluation and application of a wide-coverage multilingual semantic network. *Artificial Intelligence*, 193:217–250.

Stephan Oepen, Omri Abend, Lasha Abzianidze, Johan Bos, Jan Hajic, Daniel Hershcovich, Bin Li, Tim O'Gorman, Nianwen Xue, and Daniel Zeman. 2020. MRP 2020: The second shared task on cross-framework and cross-lingual meaning representation parsing. In *Proceedings of the CoNLL 2020 Shared Task: Cross-Framework Meaning Representation Parsing*, pages 1–22, Online. Association for Computational Linguistics.

Stephan Oepen and Jan Tore Lønning. 2006. Discriminant-based MRS banking. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC'06)*, Genoa, Italy. European Language Resources Association (ELRA).

Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. The Proposition Bank: An annotated corpus of semantic roles. *Computational Linguistics*, 31(1):71–106.

Giovanni Paolini, Ben Athiwaratkun, Jason Krone, Jie Ma, Alessandro Achille, Rishita Anubhai, Cícero Nogueira dos Santos, Bing Xiang, and Stefano Soatto. 2021. Structured prediction as translation between augmented natural languages. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.

Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.

Luigi Procopio, Edoardo Barba, Federico Martelli, and Roberto Navigli. 2021a. Multimirror: Neural cross-lingual word alignment for multilingual word sense disambiguation. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 3915–3921. International Joint Conferences on Artificial Intelligence Organization. Main Track.

Luigi Procopio, Rocco Tripodi, and Roberto Navigli. 2021b. 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.

Xiao Pu, Nikolaos Pappas, James Henderson, and Andrei Popescu-Belis. 2018. Integrating Weakly Supervised Word Sense Disambiguation into Neural Machine Translation. *Transactions of the Association for Computational Linguistics*, 6:635–649.

Alexis Ross, Tongshuang Wu, Hao Peng, Matthew E. Peters, and Matt Gardner. 2022. Tailor: Generating and perturbing text with semantic controls. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 3194–3213. Association for Computational Linguistics.

Arka Sadhu, Tanmay Gupta, Mark Yatskar, Ram Nevatia, and Aniruddha Kembhavi. 2021. Visual semantic role labeling for video understanding. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 5589–5600. Computer Vision Foundation / IEEE.

Dan Shen and Mirella Lapata. 2007. Using semantic roles to improve question answering. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 12–21, Prague, Czech Republic. Association for Computational Linguistics.

Linfeng Song, Daniel Gildea, Yue Zhang, Zhiguo Wang, and Jinsong Su. 2019. Semantic neural machine translation using AMR. *Transactions of the Association for Computational Linguistics*, 7:19–31.

Aaron Steven White, Drew Reisinger, Keisuke Sakaguchi, Tim Vieira, Sheng Zhang, Rachel Rudinger,

Kyle Rawlins, and Benjamin Van Durme. 2016. Universal decompositional semantics on Universal Dependencies. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1713–1723, Austin, Texas. Association for Computational Linguistics.

Gregor Wiedemann, Steffen Remus, Avi Chawla, and Chris Biemann. 2019. Does BERT make any sense? interpretable word sense disambiguation with contextualized embeddings. In *Proceedings of the 15th Conference on Natural Language Processing, KONVENS 2019, Erlangen, Germany, October 9-11, 2019*.

Nianwen Xue. 2008. Labeling Chinese predicates with semantic roles. *Computational Linguistics*, 34(2):225–255.

Dayu Yuan, Julian Richardson, Ryan Doherty, Colin Evans, and Eric Altendorf. 2016. Semi-supervised word sense disambiguation with neural models. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 1374–1385, Osaka, Japan. The COLING 2016 Organizing Committee.

Sheng Zhang, Xutai Ma, Kevin Duh, and Benjamin Van Durme. 2019. AMR parsing as sequence-to-graph transduction. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 80–94, Florence, Italy. Association for Computational Linguistics.