Meaning Representations for Natural Languages: Design, Models and Applications

[†]Jeffrey Flanigan, [⋄]Tim O'Gorman, [‡]Ishan Jindal, [‡]Yunyao Li, Martha Palmer, [♢]Nianwen Xue
[†]jmflanig@ucsc.edu, [⋄]timjogorman@gmail.com, [‡]ishan.jindal@ibm.com,
[‡]yunyaoli@us.ibm.com, palmer@colorado.edu, [♢]xuen@brandeis.edu

Abstract

This tutorial reviews the design of common meaning representations, SoTA models for predicting meaning representations, and the applications of meaning representations in a wide range of downstream NLP tasks and real-world applications. Reporting by a diverse team of NLP researchers from academia and industry with extensive experience in designing, building and using meaning representations, our tutorial has three components: (1) an introduction to common meaning representations, including basic concepts and design challenges; (2) a review of SoTA methods on building models for meaning representations; and (3) an overview of applications of meaning representations in downstream NLP tasks and real-world applications. We will also present qualitative comparisons of common meaning representations and a quantitative study on how their differences impact model performance. Finally, we will share best practices in choosing the right meaning representation for downstream tasks.

1 Background

In this tutorial, we primarily discuss one thread of meaning representations encompassing the Proposition Bank (PropBank) (Palmer et al., 2005), Abstract Meaning Representations (AMR) as well as Uniform Meaning Representations (UMR), a recent extension to AMR. We will discuss the representations themselves, discuss the latest semantic role labeling (SRL) and AMR parsing techniques using these representations, and overview applications of these meaning representations to practical natural language applications.

These approaches all share the use of the predicate-specific semantic roles defined in the Proposition Bank (PropBank) (Palmer et al., 2005). In such an approach, the particular sense of "afford" in "The public was afforded a preview of the show", is sense-tagged as "afford.02" in PropBank, and it requires three semantic roles, *Arg0* the provider, *Arg1* the thing that is provided, and *Arg2* the recipient of *Arg1*. We will seek to provide attendees with good intuitions about the behavior and advantages of how such predicate-specific roles work across these different meaning representations. We will also contextualize how such an approach to semantics compares to other approaches such as FrameNet(Baker et al., 1998).

AMR can be viewed as an extension of PropBank to handle wide-coverage sentence representation. Whereas PropBank is annotated on a predicate-by-predicate basis and predicates are can be viewed as independent, Abstract Meaning Representation (AMR) (Banarescu et al., 2013) adopts PropBank-style semantic roles but also connects the different predicates in a sentence in a graph. Such an AMR graph seeks to represent the meaning of sentences as a single-rooted directed acyclic graph, where the nodes are labeled with entity or predicate types, and edges are labeled with semantic roles (e.g., Arg0, Arg1) or general semantic relations (e.g., *time*, *location*).

AMR captures the essential predicate-argument structure of a sentence that is applicable to a variety of applications as well as to languages such as Chinese. Extensions to AMR attempt to increase coverage beyond the sentence, to add additional semantic phenomena, and to increase cross-linguistic applicability(Gysel et al., 2021). We discuss these extensions with a focus on the new Uniform Meaning Representation(UMR) approach, which extends AMR to add coverage of *Aspect*, *Scope*, *Person* and *Number* annotation to the sentence level representation, adds a document-level representation that captures temporal and modal dependencies as well as coreference relations that can go beyond sentence boundaries, and which defines conventions for AMR-style annotation of languages without existing PropBank lexicons. The discussion of UMR will provide attendees with an understanding of

Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: Tutorial Abstracts, pages 1 - 8

which semantic phenomena are out of scope for AMR and how projects like UMR address them.

In this tutorial we will provide an in-depth discussion of these meaning representations. When doing so, we will also discuss how they are similar to or different from other meaning representations such as semantic dependencies (Oepen et al., 2015), Minimal Recursion Semantics (MRS) (Copestake et al., 2005), Discourse Representation Theory (DRT) (Kamp and Reyle, 2013; Bos et al., 2017) , and UCCA (Abend and Rappoport, 2013).

The increasing availability of meaning representation datasets such as PropBank as well as significant advances in modeling techniques have led to increased interest and progress in computational models for meaning representation parsers. In this tutorial, we will discuss models for SRL and AMR tasks. We will start with the traditional SRL models that rely heavily on syntactic feature templates (Xue and Palmer, 2004; Pradhan et al., 2005; Zhao et al., 2009; Akbik and Li, 2016), go on to advanced neural SRL models (He et al., 2017, 2018), and include more recent work (Marcheggiani and Titov, 2020; Fei et al., 2021a,b). For AMR parsing, we will cover early approaches and SoTA methods for graph-based methods (Flanigan et al., 2014; Foland and Martin, 2017; Lyu and Titov, 2018; Cai and Lam, 2019; Zhang et al., 2019b; Zhou et al., 2020), transition-based methods (Wang et al., 2015; Wang and Xue, 2017; Ballesteros and Al-Onaizan, 2017; Fernandez Astudillo et al., 2020; Zhou et al., 2021), grammar-based methods(Peng et al., 2015; Artzi et al., 2015; Chen et al., 2018) sequence-to-sequence methods(Konstas et al., 2017; Xu et al., 2020), and other methods (Pust et al., 2015; Welch et al., 2018; Lindemann et al., 2020; Cai and Lam, 2020; Lee et al., 2020; Lam et al., 2021). We will discuss whole-document AMR parsing (Anikina et al., 2020; Fu et al., 2021).

There is a wide range of NLP tasks that leverage meaning representations as an effective way to infuse knowledge into their models for better performance and interpretability. For instance, SRL has been widely used to build better models for information extraction, such as open information extraction (Christensen et al., 2010; Solawetz and Larson, 2021) and event extraction (Zhang et al., 2020a, 2021), opinion mining (Marasović and Frank, 2018; Zhang et al., 2019a), machine translation (Bastings et al., 2017), natural language inference (Zhang et al., 2020b), and reading comprehension (Guo et al., 2020). Similarly, AMR has been adopted for a variety of downstream NLP tasks such as information extraction (Pan et al., 2015; Garg et al., 2016; Rao et al., 2017), summarization (Liu et al., 2015; Liao et al., 2018), machine translation (Song et al., 2019; Nguyen et al., 2021), question answering (Sachan and Xing, 2016; Mitra and Baral, 2016; Kapanipathi et al., 2021), and dialog (Bonial et al., 2020; Bai et al., 2021). With the increasing availability of high-quality meaning representation parsers, we also see increasing adoption of meaning representation in wide-range of real-world applications, from an enterprise-grade contract understanding system (Agarwal et al., 2021) to customizable targeted sentiment analysis.

2 Tutorial type

We are proposing a 6-hour cutting edge tutorial to cover in depth on the design, modeling, and application of meaning representations.

3 Outline of the tutorial

The proposed tutorial is organized as follows:
I. Introduction (15 minutes). We will provide a high-level overview and evolution of common meaning representation, discussing key concepts, unique challenges and examples of applications.
II. Common Meaning Representations (150 minutes) In this section, we will provide an in-depth review of three common meaning representation – PropBank and FrameNet that have been widely used to train Semantic Role Labeling systems, Abstract Meaning Representation, a sentence-level meaning representation that inherits PropBank-style semantic roles, and Uniform Meaning Representation, a cross-lingual document-level meaning representation that to a large extent inherits the sentence-level representation of AMR. We also provide a brief overview of other common meaning representations as a brief background. We will also discuss the unique challenges around designing meaning representation. Concretely, we will organize this section as follows:

- **PropBank** We start out our discussion with PropBank-style semantic roles and their theoretical underpinnings. In particular, we will discuss the proto-roles of Dowty (Dowty, 1991). We will go over the process of developing the frame files, and how the frame files are used to annotate each predicate instances in the corpus. We will discuss how to annotate compli-

cated predicates such as phrasal verbs and light verb constructions, and end with a brief discussion of how PropBank-style semantic roles are related to FrameNet (Baker et al., 1998) and VerbNet (Schuler, 2005).

- **Abstract Meaning Representation (AMR)** We next discuss different aspects of AMR, and cover how AMR represents word senses, semantic roles, named entity types, date entity types, and relations.
- **Uniform Meaning Representation (UMR)** Finally we will discuss Uniform Meaning Representations, and discuss how UMR builds on AMR. We will also discuss the cross-lingual aspect of UMR.
- **Other Related Meaning Representations** We will provide a brief overview on other common meaning representations such as MRS, etc.
- **Comparison of Meaning Representations** We will then present a qualitative comparison of the three meaning representations on their commonalities and differences.
- **Building Meaning Representation Datasets** Finally, we will close this section with discussions on the general approaches, challenges, and emerging trend in building datasets for meaning representations.

III. Modeling Meaning Representation (100 minutes) We will next discuss computational models for SRL and AMR parsing, from early approaches to current end-to-end SoTA methods. We will discuss gaps and challenges in building and evaluating such models. We will also share a quantitative comparison study based on SoTA models and demonstrates how the differences of the meaning representations lead to differences in model performance on various examples.

IV. Applying Meaning Representation (75 minutes) We will share applications of the meaning representations for a wide range of tasks from information extraction to question answering. We will discuss how the differences in these meaning representations discussed earlier impact the choice of which one(s) to use for which downstream tasks.

V. Open Questions and Future work (15 minutes) We will conclude the tutorial by raising several open research questions in this space (e.g., creating datasets for training and evaluation at scale) and ways we as a community might work forward on these issues.

4 Breadth of the tutorial

This tutorial will have three components. The first component (45%) will introduce core concepts related to meaning representations, common meaning representations and key challenges in designing (including scaling to different languages) and developing those meaning representations. The second component (30%) will review the state-of-the-art models for two common meaning representations: SRL and AMR. It will also provide a quantitative comparison study of how the differences in meaning representations impact model performance. Finally, the last component (25%) will show how real-world applications as well as research projects leverage meaning representations for better performance and more transparency and how to decide which meaning representation to use based on downstream tasks.

5 Diversity of the team

This tutorial is to be given a team of researchers from five different institutions across academia and industry, both junior instructors (including 1 assistant professor, and 2 junior industry researcher) and researchers with extensive experience in academic and corporate research settings. The team includes creators, modelers, and users of common meaning representations. The team also has a good gender balance (two female and four male instructors).

6 Target audience and objectives

This tutorial welcomes all stakeholders in the NLP community, including NLP researchers, domain-specific practitioners, and students. In this tutorial, attendees will

- Develop fluency in core concepts of common meaning representations, state-of-the-art models for producing these meaning representations, and potential use cases.
- Gain insights into the practical benefits and challenges around leveraging meaning representations for downstream applications.
- Discuss and reflect on open questions related to meaning representations.

7 Prerequisites

As stated before, our tutorial presumes no prior knowledge on the core concepts of meaning representation. However, a basic understanding of NLP,

machine learning (especially, deep learning) concepts may be helpful. We intend to introduce the necessary concepts related to meaning representation during the introductory section of the tutorial.

8 Reading list

We aim to make the tutorial self-contained, but it will be helpful if the attendees can get some basic understanding of this field by going through the following reading list: PropBank: (Palmer et al., 2005), AMR: (Banarescu et al., 2013), UMR: (Gysel et al., 2021), SRL models: (Pradhan et al., 2005; He et al., 2017), and AMR models: (Flanigan et al., 2014; Lyu and Titov, 2018; Xu et al., 2020).

9 Audience size estimation

We are proposing a cutting edge tutorial on meaning representation. No similar tutorial has been given in ACL/EMNLP/NAACL/COLING in the past five years. Since meaning representation is an important topic in NLP, we expect that this tutorial will be popular with 50 - 100 attendees.

10 Open Access

We agree to allow the publication of our slides and video recording of our tutorial in the ACL Anthology.

11 Technique Equipment

To give this tutorial, we need to have internet access and a projector or large screen. No special requirements needed.

12 Preferred Venue

Due to travel restrictions of our instructors, we prefer NAACL and ACL over the other venues.

13 Ethics Statement

Infusing meaning representations into NLP models are shown to be effective in injecting knowledge into such models. As such, meaning representations allow deep understanding of languages and identify more nuanced instances of ethics concerns (e.g. biases). Furthermore, meaning representations allow the building of fully interpretable yet effective models. We hope that this tutorial helps the audience to develop a deeper appreciation for such topics and equips them with powerful tools to mitigate recent concerns that have arisen with NLP models with regard to explainability and bias.

14 Author biographies

Martha Palmer is the Helen & Hubert Croft Professor of Engineering in the Computer Science Department, and Arts & Sciences Professor of Distinction for Linguistics, at the University of Colorado, with over 300 peer-reviewed publications. Her research is focused on capturing elements of the meanings of words that can comprise automatic representations of complex sentences and documents in many languages. She is a co-Director of CLEAR, an ACL Fellow, and an AAAI Fellow.

Nianwen Xue is a Professor in the Computer Science Department and the Language & Linguistics Program at Brandeis University. His core research interests include developing linguistic corpora annotated with syntactic, semantic, and discourse structures, as well as machine learning approaches to syntactic, semantic and discourse parsing. He is an action editor for Computational Linguistics.

Ishan Jindal is a Research Staff Member with IBM Research - Almaden. His research interest lies at the intersection of machine learning and NLP, primarily in semantic parsing and model analysis for enterprise use cases. He regularly publishes papers at ML and NLP conferences.

Jeffrey Flanigan is an Assistant Professor in the Computer Science and Engineering Department at University of California Santa Cruz. He research interests are in semantic parsing and generation, with a focus on AMR, and using semantic representations in downstream applications such as summarization and machine translation. Previously he has given a tutorial in AMR at NAACL 2015.

Tim O'Gorman is a Senior Research Scientist at Thorn. He was involved in AMR 2.0 and 3.0 annotations, the Multi-sentence AMR corpus, and updates to PropBank. He co-organized the CoNLL'19 and '20 Meaning Representation Parsing shared task. His interests are in the extensions of meaning representations to cross-sentence phenomena.

Yunyao Li is a Distinguished Research Staff Member and Senior Research Manager with IBM Research - Almaden. Her expertise is at the intersection of NLP, databases, HCI, and information retrieval. Her work has resulted in 80+ peer-reviewed publications and transferred into 20+ commercial products. She regularly gives talks and tutorials, such as Explainability for NLP (AACL'20, KDD'21), and Deep Learning on Graphs for NLP (NAACL'21, KDD'21, IJCAI'21). She is an ACM Distinguished Member.

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.

Arvind Agarwal, Laura Chiticariu, Poornima Chozhiyath Raman, Marina Danilevsky, Diman Ghazi, Ankush Gupta, Shanmukha C. Guttula, Yannis Katsis, Rajasekar Krishnamurthy, Yunyao Li, Shubham Mudgal, Vitobha Munigala, Nicholas Phan, Dhaval Sonawane, Sneha Srinivasan, Sudarshan R. Thitte, Mitesh Vasa, Ramiya Venkatachalam, Vinitha Yaski, and Huaiyu Zhu. 2021. Development of an enterprise-grade contract understanding system. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Industry Papers, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 222–229. Association for Computational Linguistics.

Alan Akbik and Yunyao Li. 2016. K-SRL: Instance-based learning for semantic role labeling. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 599–608, Osaka, Japan. The COLING 2016 Organizing Committee.

Tatiana Anikina, Alexander Koller, and Michael Roth. 2020. Predicting coreference in Abstract Meaning Representations. In *Proceedings of the Third Workshop on Computational Models of Reference, Anaphora and Coreference*, pages 33–38, Barcelona, Spain (online). Association for Computational Linguistics.

Yoav Artzi, Kenton Lee, and Luke Zettlemoyer. 2015. Broad-coverage CCG semantic parsing with AMR. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1699–1710, Lisbon, Portugal. Association for Computational Linguistics.

Xuefeng Bai, Yulong Chen, Linfeng Song, and Yue Zhang. 2021. Semantic representation for dialogue modeling. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4430–4445, Online. 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.

Miguel Ballesteros and Yaser Al-Onaizan. 2017. AMR parsing using stack-LSTMs. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1269–1275, Copenhagen, Denmark. 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.

Jasmijn Bastings, Ivan Titov, W. Aziz, Diego Marcheggiani, and Khalil Sima'an. 2017. Graph convolutional encoders for syntax-aware neural machine translation. In *EMNLP*.

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 Twelfth Language Resources and Evaluation Conference*, pages 684–695, Marseille, France. European Language Resources Association.

Johan Bos, Valerio Basile, Kilian Evang, Noortje Venhuizen, and Johannes Bjerva. 2017. The groningen meaning bank. In Nancy Ide and James Pustejovsky, editors, *Handbook of Linguistic Annotation*, volume 2, pages 463–496. Springer.

Deng Cai and Wai Lam. 2019. Core semantic first: A top-down approach for amr parsing. *arXiv preprint arXiv:1909.04303*.

Deng Cai and Wai Lam. 2020. Amr parsing via graph-sequence iterative inference. *arXiv preprint arXiv:2004.05572*.

Yufei Chen, Weiwei Sun, and Xiaojun Wan. 2018. Accurate SHRG-based semantic parsing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 408–418, Melbourne, Australia. Association for Computational Linguistics.

Janara Christensen, Mausam, Stephen Soderland, and Oren Etzioni. 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.

Ann Copestake, Dan Flickinger, Carl Pollard, and Ivan A Sag. 2005. Minimal recursion semantics: An introduction. *Research on language and computation*, 3(2):281–332.

David Dowty. 1991. Thematic proto-roles and argument selection. *language*, 67(3):547–619.

Hao Fei, Fei Li, Bobo Li, and Donghong Ji. 2021a. Encoder-decoder based unified semantic role labeling with label-aware syntax. In *Proc. AAAI Conf. Artif. Intell.*, pages 1479–1488.

Hao Fei, Meishan Zhang, Bobo Li, and Donghong Ji. 2021b. End-to-end semantic role labeling with neural transition-based model. In *Proc. AAAI Conf. Artif. Intell.*, pages 566–575.

Ramón Fernandez Astudillo, Miguel Ballesteros, Tahira Naseem, Austin Blodgett, and Radu Florian. 2020. Transition-based parsing with stack-transformers. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1001–1007, Online. Association for Computational Linguistics.

Jeffrey Flanigan, Sam Thomson, Jaime Carbonell, Chris Dyer, and Noah A. Smith. 2014. A discriminative graph-based parser for the Abstract Meaning Representation. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1426–1436, Baltimore, Maryland. Association for Computational Linguistics.

William Foland and James H. Martin. 2017. Abstract Meaning Representation parsing using LSTM recurrent neural networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 463–472, Vancouver, Canada. Association for Computational Linguistics.

Qiankun Fu, Linfeng Song, Wenyu Du, and Yue Zhang. 2021. End-to-end AMR coreference resolution. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4204–4214, Online. Association for Computational Linguistics.

Sahil Garg, Aram Galstyan, Ulf Hermjakob, and Daniel Marcu. 2016. Extracting biomolecular interactions using semantic parsing of biomedical text. In *Thirtieth AAAI Conference on Artificial Intelligence*.

Shaoru Guo, Yong Guan, Ru Li, Xiaoli Li, and Hongye Tan. 2020. Incorporating syntax and frame semantics in neural network for machine reading comprehension. In *COLING*.

Jens E. L. Van Gysel, Meagan Vigus, Jayeol Chun, Kenneth Lai, Sarah Moeller, Jiarui Yao, Timothy J. O'Gorman, Andrew Cowell, W. Bruce Croft, Chu Ren Huang, Jan Hajic, James H. Martin, Stephan Oepen, Martha Palmer, James Pustejovsky, Rosa Vallejos, and Nianwen Xue. 2021. Designing a uniform meaning representation for natural language processing. *Künstliche Intelligenz*, pages 1–18.

Luheng He, Kenton Lee, Mike Lewis, and Luke Zettlemoyer. 2017. Deep semantic role labeling: What works and what's next. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 473–483, Vancouver, Canada. Association for Computational Linguistics.

Shexia He, Zuchao Li, Hai Zhao, and Hongxiao Bai. 2018. Syntax for semantic role labeling, to be, or not to be. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2061–2071, Melbourne, Australia. Association for Computational Linguistics.

Hans Kamp and Uwe Reyle. 2013. *From discourse to logic: Introduction to modeltheoretic semantics of natural language, formal logic and discourse representation theory*, volume 42. Springer Science & Business Media.

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.

Ioannis Konstas, Srinivasan Iyer, Mark Yatskar, Yejin Choi, and Luke Zettlemoyer. 2017. Neural AMR: Sequence-to-sequence models for parsing and generation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 146–157, Vancouver, Canada. Association for Computational Linguistics.

Hoang Thanh Lam, Gabriele Picco, Yufang Hou, Young-Suk Lee, Lam M. Nguyen, Dzung T. Phan, Vanessa López, and Ramón Fernandez Astudillo. 2021. Ensembling graph predictions for AMR parsing. *CoRR*, abs/2110.09131.

Young-Suk Lee, Ramon Fernandez Astudillo, Tahira Naseem, Revanth Gangi Reddy, Radu Florian, and Salim Roukos. 2020. Pushing the limits of amr parsing with self-learning. *arXiv preprint arXiv:2010.10673*.

Kexin Liao, Logan Lebanoff, and Fei Liu. 2018. Abstract meaning representation for multi-document summarization. *arXiv preprint arXiv:1806.05655*.

Matthias Lindemann, Jonas Groschwitz, and Alexander Koller. 2020. Fast semantic parsing with well-typedness guarantees. In *Proceedings of the 2020*

Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 3929–3951, Online. Association for Computational Linguistics.

Fei Liu, Jeffrey Flanigan, Sam Thomson, Norman Sadeh, and Noah A. Smith. 2015. Toward abstractive summarization using semantic representations. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1077–1086, Denver, Colorado. Association for Computational Linguistics.

Chunchuan Lyu and Ivan Titov. 2018. AMR parsing as graph prediction with latent alignment. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 397–407, Melbourne, Australia. Association for Computational Linguistics.

Ana Marasović and Anette Frank. 2018. Srl4orl: Improving opinion role labeling using multi-task learning with semantic role labeling. In *NAACL*.

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.

Arindam Mitra and Chitta Baral. 2016. Addressing a question answering challenge by combining statistical methods with inductive rule learning and reasoning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30.

Long HB Nguyen, Viet H Pham, and Dien Dinh. 2021. Improving neural machine translation with amr semantic graphs. *Mathematical Problems in Engineering*, 2021.

Stephan Oepen, Marco Kuhlmann, Yusuke Miyao, Daniel Zeman, Silvie Cinková, Dan Flickinger, Jan Hajic, and Zdenka Uresova. 2015. Semeval 2015 task 18: Broad-coverage semantic dependency parsing. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 915–926.

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

Xiaoman Pan, Taylor Cassidy, Ulf Hermjakob, Heng Ji, and Kevin Knight. 2015. Unsupervised entity linking with Abstract Meaning Representation. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1130–1139, Denver, Colorado. Association for Computational Linguistics.

Xiaochang Peng, Linfeng Song, and Daniel Gildea. 2015. A synchronous hyperedge replacement grammar based approach for AMR parsing. In *Proceedings of the Nineteenth Conference on Computational Natural Language Learning*, pages 32–41, Beijing, China. Association for Computational Linguistics.

Sameer Pradhan, Wayne Ward, Kadri Hacioglu, James Martin, and Daniel Jurafsky. 2005. Semantic role labeling using different syntactic views. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL'05)*, pages 581–588, Ann Arbor, Michigan. Association for Computational Linguistics.

Michael Pust, Ulf Hermjakob, Kevin Knight, Daniel Marcu, and Jonathan May. 2015. Parsing English into Abstract Meaning Representation using syntax-based machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1143–1154, Lisbon, Portugal. Association for Computational Linguistics.

Sudha Rao, Daniel Marcu, Kevin Knight, and Hal Daumé III. 2017. Biomedical event extraction using Abstract Meaning Representation. In *BioNLP 2017*, pages 126–135, Vancouver, Canada,. Association for Computational Linguistics.

Mrinmaya Sachan and Eric Xing. 2016. Machine comprehension using rich semantic representations. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 486–492, Berlin, Germany. Association for Computational Linguistics.

Karin Kipper Schuler. 2005. *VerbNet: A broad-coverage, comprehensive verb lexicon*. University of Pennsylvania.

Jacob Solawetz and Stefan Larson. 2021. Lsoie: A large-scale dataset for supervised open information extraction. In *EACL*.

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.

Chuan Wang and Nianwen Xue. 2017. Getting the most out of AMR parsing. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1257–1268, Copenhagen, Denmark. Association for Computational Linguistics.

Chuan Wang, Nianwen Xue, and Sameer Pradhan. 2015. A transition-based algorithm for AMR parsing. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 366–375, Denver, Colorado. Association for Computational Linguistics.

Charles Welch, Jonathan K. Kummerfeld, Song Feng, and Rada Mihalcea. 2018. World knowledge for Abstract Meaning Representation parsing. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).

Dongqin Xu, Junhui Li, Muhua Zhu, Min Zhang, and Guodong Zhou. 2020. Improving amr parsing with sequence-to-sequence pre-training. *arXiv preprint arXiv:2010.01771*.

Nianwen Xue and Martha Palmer. 2004. Calibrating features for semantic role labeling. In *EMNLP*, pages 88–94.

Hongming Zhang, Haoyu Wang, and Dan Roth. 2020a. Unsupervised label-aware event trigger and argument classification. *ArXiv*, abs/2012.15243.

Hongming Zhang, Haoyu Wang, and Dan Roth. 2021. Zero-shot label-aware event trigger and argument classification. In *FINDINGS*.

Meishan Zhang, Peilin Liang, and Guohong Fu. 2019a. Enhancing opinion role labeling with semantic-aware word representations from semantic role labeling. In *NAACL*.

Sheng Zhang, Xutai Ma, Kevin Duh, and Benjamin Van Durme. 2019b. 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.

Zhuosheng Zhang, Yuwei Wu, Zhao Hai, Z. Li, Shuailiang Zhang, Xi Zhou, and Xiang Zhou. 2020b. Semantics-aware bert for language understanding. *ArXiv*, abs/1909.02209.

Hai Zhao, Wenliang Chen, and Chunyu Kit. 2009. Semantic dependency parsing of NomBank and PropBank: An efficient integrated approach via a large-scale feature selection. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pages 30–39, Singapore. Association for Computational Linguistics.

Jiawei Zhou, Tahira Naseem, Ramón Fernandez Astudillo, and Radu Florian. 2021. Amr parsing with action-pointer transformer. *arXiv preprint arXiv:2104.14674*.

Qiji Zhou, Yue Zhang, Donghong Ji, and Hao Tang. 2020. AMR parsing with latent structural information. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4306–4319, Online. Association for Computational Linguistics.