

Meaning Representations for Natural Languages: Design, Models and Applications

* Julia Bonn, † Jeffrey Flanigan, ▷ Jan Hajič,

‡ Ishan Jindal, ◁ Yunyao Li, ◇ Nianwen Xue

*julia.bonn@colorado.edu, †jmflanig@ucsc.edu, ▷hajic@ufal.mff.cuni.cz,

‡ishan.jindal@ibm.com, ◁yunyaoli@gmail.com, ◇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 propose a full-day, cutting-edge tutorial for all stakeholders in the AI community, including NLP researchers, domain-specific practitioners, and students.

1. Introduction

This tutorial aims to introduce the NLP community to an emerging research area that has the potential to create linguistic resources and build computational models that provide critical components for interpretable and controllable NLP systems. While large language models have shown remarkable ability to generate fluent and mostly coherent text, the blackbox nature of these models makes it difficult to know where to tweak these models to fix errors or at least anticipate errors if they cannot easily be fixed. For instance, LLMs are known to hallucinate and generate factually incorrect answers when prompted as there is no mechanism in these models to constrain them to only provide factually correct answers. Addressing this issue requires that first of all the models have access to a body of verifiable facts, and then when generating answers to prompts or queries, do not alter them materially to make the answers factually incorrect. Interpretability and controllability in NLP systems are critical in high-stake application scenarios such as the health domain, where AI systems are used as medical assistants.

In the past few decades, there has been a steady accumulation of semantically annotated resources that are increasingly richer in representation. As these resources become available, steady progress has been made in developing computational models that can automatically parse unstructured text into these semantic representations with increasing accuracy. These models have reached a level of accuracy that makes them useful in practical applications. For example, these models have been used in information extraction, where entities and relations are extracted from unstructured text. It is now conceivable that these models can be used to extract verifiable facts at scale

to build controllable and interpretable systems that can produce factual correct answers. These rich semantic representations are also needed in human-robot interaction (HRI) systems to facilitate on-the-fly grounding so that the robot can establish connections with its surroundings and interact with them in a meaningful way. These meaning representations are easily translated into logical representations to support logical reasoning that LLMs often struggle with, or they can be used to develop NLP systems for low-resource languages where there is insufficient data to train LLMs, but the richness in semantic representation can to some extent make up for the lack of quantity. This tutorial will provide an overview of these semantic representations, the computational models that are trained on them, as well as the practical applications built with these representations. We will also delve into future directions for this line of research and examine how these meaning representations might be used to build interpretable and controllable applications, used in human-robot interaction scenarios, and low-resource settings.

2. Target audience

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

In this tutorial, attendees will

- Develop fluency in core concepts of common meaning representations, state-of-the-art mod-

els 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.

3. Outline

3.1. Background

In this tutorial, we primarily discuss one thread of meaning representations that encompasses the Proposition Bank (PropBank) (Palmer et al., 2005), Abstract Meaning Representations (AMR) (Banarescu et al., 2013) as well as Uniform Meaning Representations (UMR) (Gysel et al., 2021), a recent extension to AMR, but will situate our discussion with a comparison with related meaning representations. We will discuss the representations themselves, as well as 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.

The proposed tutorial is organized as follows:

I. Introduction (15 minutes). This section provides a high-level overview of the evolution of common meaning representation, discussing key concepts, unique challenges, and examples of applications.

II. Common Meaning Representations (150 minutes) This section provides an in-depth review of three common meaning representation – PropBank, Abstract Meaning Representation, and Uniform Meaning Representation. It also provides a brief overview of other common meaning representations and a comparison between these meaning representations. Concretely, we will organize this section as follows:

- **PropBank**

- An intuitive introduction of Propbank-style semantic roles
- Defining predicate-specific semantic roles in frame files
- Semantic roles for complicated predicates
- Relation of propbank-style semantic roles to FrameNet and VerbNet semantic roles

- **Abstract Meaning Representation (AMR)**

This section discusses different aspects of AMR, and covers how AMR represents word senses, semantic roles, named entity types, date entity types, and relations.

- Format and basics
- Some details and design decisions

- Multi-sentence AMRs
- Relation to other formalisms

- **Uniform Meaning Representation (UMR)**

This section overviews Uniform Meaning Representations, and discusses how UMR builds on AMR and extends it to cross-lingual settings.

- Sentence-level representations of UMR: aspect, person, number, and quantification scope
- Document-level representations: temporal and modal dependencies, coreference
- Cross-lingual applicability of UMR.
- UMR-Writer: tool for annotating UMRs

- **Other Related Meaning Representations**

This section provides a brief overview of other common meaning representations such as MRS, Tectogrammatical Representation used in the Prague Dependency Treebanks (PDT), etc.

- Discourse Representation Structures (annotations in Groening Meaning Bank and Parallel Meaning Bank)
- Minimal Recursion Semantics
- Universal Conceptual Cognitive Annotation
- Prague Semantic Dependencies (Tectogrammatical annotation of syntax and semantics in the PDT-style treebanks)

- **Comparison of Meaning Representations**

This section presents a qualitative comparison of the three meaning representations on their commonalities and differences.

- Alignment to text / compositionality
- Logical and executable forms
- Lexicon and ontology differences
- Task-specific representations
- Discourse-level representations

- **Building Meaning Representation Datasets**

This section discusses the general approaches, challenges, and emerging trend in building data sets for meaning representations.

III. Modeling Meaning Representation (100 minutes) This section discusses computational models for SRL and AMR parsing, from early approaches to current end-to-end SoTA methods.

- Semantic role labeling
- AMR parsing
- AMR generation

IV. Applying Meaning Representation (75 minutes) This section shares applications of the meaning representations for a wide range of tasks from information extraction to question answering. This section also discusses how the differences in these meaning representations impact the choice of which one(s) to use for which downstream tasks.

- Applications of Meaning Representations
- Case Studies

V. Open Questions and Future Directions (15 minutes) The final section concludes the tutorial by raising open research questions about the representation, modeling, and application of meaning representations in NLP and how they could complement LLMs.

4. Diversity considerations

Representing languages of the world. We devote considerable time to discuss the meaning representation for low-resource languages, which tend to have distinct linguistic properties that have previously received little attention. This contributes to greater fairness in the field.

Diversity of the team. This tutorial is to be given by a team of researchers from six different institutions across academia and industry, both junior instructors (including 1 assistant professor, 1 advanced PhD student, and 1 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).

5. Reading list

[LAW'2013] 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

[NAACL'2022] Li Zhang, Ishan Jindal, and Yunyao Li. “Label definitions improve semantic role labeling.” In Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pp. 5613-5620. 2022.

[LREC'2022] Ishan Jindal, Alexandre Rademaker, Michał Ulewicz, Ha Linh, Huyen Nguyen, Khoi-Nguyen Tran, Huaiyu Zhu, and Yunyao Li. “Universal proposition bank 2.0.” In Proceedings of

the Thirteenth Language Resources and Evaluation Conference, pp. 1700-1711. 2022.

[KI'2021] Jens E. L. Van Gysel, Meagan Vigus, Jayeol Chun, Kenneth Lai, Sarah Moeller, Jiarui Yao, Tim O’Gorman, Andrew Cowell, William Croft, Chu-Ren Huang, Jan Hajič, James H. Martin, Stephan Oepen, Martha Palmer, James Pustejovsky, Rosa Vallejos, and Nianwen Xue. “Designing a uniform meaning representation for natural language processing.” *KI-Künstliche Intelligenz* 35, no. 3-4 (2021): 343-360.

[NAACL'2018] 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.

[NAACL'2016] Flanigan, Jeffrey, Chris Dyer, Noah A. Smith, and Jaime G. Carbonell. “Generation from abstract meaning representation using tree transducers.” In Proceedings of the 2016 conference of the north american chapter of the association for computational linguistics: Human language technologies, pp. 731-739. 2016.

[NAACL'2015] Wang, Chuan, Nianwen Xue, and Sameer Pradhan. “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, pp. 366-375. 2015.

[ACL'2014] Flanigan, Jeffrey, Sam Thomson, Jaime G. Carbonell, Chris Dyer, and Noah A. Smith. “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), pp. 1426-1436. 2014.

6. Presenters

Julia Bonn is an advanced Ph.D. student in Linguistics and Cognitive Science at the University of Colorado, Boulder. During her last 14 years as a Senior Research Assistant at CLEAR, she has been a long-term contributor to PropBank and the PropBank Roleset Lexicon, Verbnet, AMR, and UMR. She is also the developer of SpatialAMR, an extension to AMR annotation for fine-grained, multimodal annotation of spatially rich corpora. Her research interests center on bringing multimodality and pragmatics into cross-lingual meaning representations, and development of lexical resources for these applications with a special focus on how such resources can be designed to better serve polysynthetic languages.

Jan Hajič is the director of the large research infrastructure for Language Resources, Digital Hu-

manities and Arts LINDAT/CLARIAH-CZ, which is part of the EU's CLARIN, DARIAH and EHRI networks. He is also the vice-director of the Institute of Formal and Applied Linguistics at Charles University, Prague, Czech Republic. His interests span the morphology and part-of-speech tagging of inflective languages, machine translation, deep language understanding, and the application of statistical machine learning in NLP. His work experience includes both industrial research (IBM Research Yorktown Heights, NY, USA, in 1991-1993) and academia (Charles University in Prague, Czech Republic and Johns Hopkins University, Baltimore, MD, USA, 1999-2000, adjunct position at University of Colorado, USA, 2017-2025). He has published more than 200 conference and journal papers, a book and book chapters, encyclopedia and handbook entries. He regularly teaches both regular courses as well as tutorials and lectures at various international training schools. He has been the PI or Co-PI of numerous international as well as large national grants and projects (EU and NSF). He is the chair of the Executive Board of META-NET, European research network in language technology, and is a member of several other international boards and committees.

Jeffrey Flanigan is an Assistant Professor in the Department of Computer Science and Engineering at the University of California Santa Cruz. His research includes semantic parsing and generation, question answering, and the use of semantic representations in downstream applications such as summarization and machine translation. Previously he has given a tutorial in AMR at NAACL 2015, and a tutorial on Meaning Representations at EMNLP 2022. He served as a senior area chair for CoNLL in 2022.

Ishan Jindal is a Staff Research Scientist with IBM Research - Almaden. He got his PhD degree in Electrical Engineering from Wayne State University, Michigan. His research interest lies at the intersection of Machine Learning (Deep Learning) and Natural Language Processing (NLP), with a particular focus on multilingual shallow semantic parsing and model analysis for enterprise use cases and their applications in various NLP downstream applications. His work has been published at top-tier conferences, including ICASSP, EMNLP, NAACL, ICDM, ISIT, Big Data, and LREC. He has served as an area chair PC member in many conferences (e.g., ACL, EMNLP, NAACL, EACL, and AAAI) and journals (e.g., TNNLS and TACL).

Yunyao Li is the Director of Machine Learning, Adobe Experience Platform. She was the Head of Machine Learning at the Apple Knowledge Platform and a Distinguished Research Staff Member and Senior Research Manager with IBM Research.

She is particularly known for her work in scalable NLP, enterprise search, and database usability. She was an IBM Master Inventor. Her technical contributions have been recognized by prestigious awards on a regular basis, such as IBM Corporate Technical Award (2022), IBM Outstanding Research Achievement Awards (2021, 2020, 2019), ISWC Best Demo Award (2020), and YWCA's Tribute to Women Award (2019), among others. She is a member of inaugural New Voices Program of the American National Academies and represented US young scientists at World Laureates Forum Young Scientists Forum in 2019. Regularly organizes conferences, workshops, and panels at top AI conferences and served on prestigious program committees, editorial board and review panels. She is an ACM Distinguished Member and an elected member of the North American Chapter of the Association for Computational Linguistics (NAACL) Executive Board (2023-2024).

Nianwen Xue is a Professor and chair 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, and currently serves on the editorial boards of Language Resources and Evaluation (LRE). He also served as the editor-in-chief of the ACM Transactions on Asian and Low-Resource Language Information Processing (TALLIP) from 2016 to 2019, and has frequently served as area chairs for ACL, EMNLP, and COLING. He is the program co-chair of the 2024 Joint International Conference on Computational Linguistics, Language Resources, and Evaluation.

7. 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 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.

8. Bibliographical References

Omri Abend and Ari Rappoport. 2013. Universal conceptual cognitive annotation (ucca). In *Proceedings of the 51st Annual Meeting of the As-*

- sociation 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.
- 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.
- 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*.
- Johan Bos, Valerio Basile, Kilian Evang, Noortje Venhuizen, and Johannes Bjerva. 2017a. The groningen meaning bank. In Nancy Ide and James Pustejovsky, editors, *Handbook of Linguistic Annotation*, volume 2, pages 463–496. Springer.
- Johan Bos, Valerio Basile, Kilian Evang, Noortje J Venhuizen, and Johannes Bjerva. 2017b. The groningen meaning bank. In *Handbook of linguistic annotation*, 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*.
- 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.
- Veena G, Deepa Gupta, Akshay Anil, and Akhil M S. 2019. An ontology driven question answering system for legal documents. *2019 2nd International Conference on Intelligent Computing, Instrumentation and Control Technologies (ICICICT)*, 1:947–951.
- 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.
- 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.
- Hoang Thanh Lam, Gabriele Picco, Yufang Hou, Young-Suk Lee, Lam M. Nguyen, Dzong 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*.
- Ling Liu, Ishan Jindal, and Yunyao Li. 2022. Is semantic-aware bert more linguistically aware? a case study on natural language inference. In *Annual Conference of the North American Chapter of the 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*.
- 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(1).
- 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.
- Karin Kipper Schuler. 2005. *VerbNet: A broad-coverage, comprehensive verb lexicon*. University of Pennsylvania.
- Peng Shi and Jimmy Lin. 2019. Simple bert models for relation extraction and semantic role labeling. *arXiv preprint arXiv:1904.05255*.
- Jacob Solawetz and Stefan Larson. 2021. Lsoie: A large-scale dataset for supervised open information extraction. In *EACL*.
- 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*.
- Li Zhang, Ishan Jindal, and Yunyao Li. 2022. Label definitions improve semantic role labeling. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5613–5620.
- Meishan Zhang, Peilin Liang, and Guohong Fu. 2019. Enhancing opinion role labeling with semantic-aware word representations from semantic role labeling. In *NAACL*.
- 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.
- Jiawei Zhou, Tahira Naseem, Ramón Fernandez Astudillo, and Radu Florian. 2021. Amr parsing with action-pointer transformer. *arXiv preprint arXiv:2104.14674*.