Meaning Representations: Recent Advances in Parsing and Generation

Juri Opitz

October 2020

Abstract

In this course we will visit recent research that – given a *theory of symbolic meaning* – aims at developing functions for generating natural language text from meaning representations ('*Generation*') and generating meaning representations from from natural language text ('*Parsing*').

1 Content of this course

"What's the meaning of text, signs, and symbols?" This question has troubled philosophers for ages. In this course, we aren't likely to find an answer to this question (because we aren't even trying!). Instead, we will pleasantly circumnavigate this question and sail in calmer waters, where we are already given a theory of meaning, and therefore we are absolved from all potential blame for the inevitable loopholes in our theories of meaning (because we will not develop them!). We will also not create data for this task, so that, here too, we are absolved from blame of everything that can go wrong in complex data annotation.

Instead, the task we consider is just to learn a mapping between (compositions of) symbols in one domain (that is phrases composed of words or characters of natural language text) and (compositions of) symbols from another domain, the domain of (abstract) meaning (representations). To learn our 'meaning mappings', we won't design any rules based on our intuitions of how the world works, leaving all this dirty work to neural networks.

More precisely, we visit recent research on neural approaches that provide high performance on meaning representation benchmark data sets. Thereby, we will improve upon our technical knowledge about complex, high-performance machine learning systems. Moreover, by doing this, perhaps, we may, or may not, hope to extend our knowledge about the meaning of signs and symbols, obtaining not only a gain of technical knowledge but possibly also a gain in knowledge about the meaning of signs and symbols.

Does all this excite you (or slightly raises your attention)? Then please join this course, we will have a good time.

2 Requirements to pass the course

1. Participation in our weekly heiconf-meetings. A link will be sent to you.
2. Presentation of a research paper (max. 25 minutes + max. 10 minutes discussion).
3. either i) a term paper or ii) a small implementation project with a technical report of the experimental settings and the results (or iii), a second presentation, subject to availability).

3 Schedule

First meeting (9.11.2020) Welcome and small introduction.

3.1 Subsequent meetings

After the first session, we will try to roughly stick to the following agenda:

Laying the groundwork

- **16.11.2020, Basics I: DRS and AMR meaning representation theories (2 presenters).** We will discuss two prominent theories of text meaning. Discourse representation theory [8] and abstract meaning representation [2].
- **23.11.2020, Basics II. Evaluation and IAA metrics for parsing and generation evaluation (2 presenters).** We discuss the Smatch [5] algorithm that compares AMR graphs, and we visit Bleu [16] that is commonly used for evaluation of all kinds of text generations (including meaning representation to text). In addition, we ask ourselves: what do humans actually have to say about the quality of text generated from AMRs? [13].
- **30.11.2020, Neural sequence-to-sequence models (2 presenters).** We discuss the 'classical' piece of Sutskever et al. [19] and 'the transformer' [21]. This prepares us well for the next session(s).

Assessing recent advances in MR parsing and generation

- **7.12.2020, Neural sequence-to-sequence (seq2seq) for AMR parsing and AMR generation (2 presenters).** We will discuss Konstas et al. [9], who find out that neural seq2seq models work well for AMR parsing *and* AMR generation. Then we see some more 'tricks of the trade' and a character-based AMR parsing seq2seq approach by Van Noord and Bos [20].

- **14.12.2020, MT pre-training greatly boosts seq2seq AMR parsing models (1 presenter)** We discuss the recent work of Xu et al. [23], who show that large-scale MT pre-training provides very useful inductive biases for AMR parsing with seq2seq.
- **11.1.2021, Graph encoders for generation from AMR (2 presenters)**. We discuss the work of Song et al. [18] who use a recurrent graph encoder for better sequence generation from AMR (without the need of silver data). And we discuss the language model-based approach by Mager et al. [12] who show that pre-trained LMs (GPT2 [17]) can be taught to understand AMR language for improved generation.
- **18.1.2021, AMR parsing as graph prediction I (2 presenters).** We will discuss Zhang et al. [24] who encode the sentence with BERT [7] and generate AMR nodes, finally predicting relations between pairs of nodes. And we discuss the work of Cai and Lam [4] who use iterative graph decoding for refined AMR graph generation and a BERT sentence encoding.
- **25.1.2021, Transition-based AMR parsing (2 presenters)** We discuss the work of Wang et al. [22], to get started on transition-based AMR parsing and then discuss the work of Lindemann et al. [10], who aim at greatly increased AMR parsing speed.
- **1.2.2021, Cross-lingual AMR parsing (2 presenters)** We will talk about the 'early' work of Damonte and Cohen [6] and the recent work by Blloshmi et al. [3], who both target a setup that predicts (English) AMR graphs from sentences of all kinds of languages.
- **8.2.2021, Let's move to discourse level! (2 presenters)** We will see how i) we can use structural decoding for DRS parsing [11] and ii) find out which things we can learn from a recent shared task in DRS parsing [1].
- **15.2.2021 Papers of your choice (max. 3 presenters)** If you have a paper that fits well in the framework of this seminar but you find it missing from the proposed schedule and wished it was part of it, this is your chance to present it!
- **22.2.2021** Wrap-up, project and term paper discussion.

4 Possible implementation projects

There are several possible implementation projects. Many of the above papers provide code that can be installed and experimented with and the institute possesses a license for the AMR data. Alternatively, you are free to propose something, and we can discuss about it, whether it will be feasible. Furthermore, if you are interested in the quality assessment of automatically generated AMRs or text generation from AMRs, you could build on some code that the teacher of this seminar wrote [14, 15].

5 Contacting the teacher

You can ask questions directly after the sessions have ended, or by writing an email to `opitz@cl.uni-heidelberg.de` (Please put `[MRPG]` in the email subject).

References

[1] Lasha Abzianidze, Rik van Noord, Hessel Haagsma, and Johan Bos. The first shared task on discourse representation structure parsing. *arXiv preprint arXiv:2005.13399*, 2020.

[2] Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. Abstract meaning representation for sembanking. In *Proceedings of the 7th linguistic annotation workshop and interoperability with discourse*, pages 178–186, 2013.

[3] Rexhina Blloshmi, Rocco Tripodi, and Roberto Navigli. Xl-amr: Enabling cross-lingual amr parsing with transfer learning techniques. In *Proc. of EMNLP*, 2020.

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

[5] Shu Cai and Kevin Knight. Smatch: an evaluation metric for semantic feature structures. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 748–752, 2013.

[6] Marco Damonte and Shay B Cohen. Cross-lingual abstract meaning representation parsing. *arXiv preprint arXiv:1704.04539*, 2017.

[7] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

[8] Hans Kamp, Josef Van Genabith, and Uwe Reyle. Discourse representation theory. In *Handbook of philosophical logic*, pages 125–394. Springer, 2011.

[9] Ioannis Konstas, Srinivasan Iyer, Mark Yatskar, Yejin Choi, and Luke Zettlemoyer. Neural amr: Sequence-to-sequence models for parsing and generation. *arXiv preprint arXiv:1704.08381*, 2017.

[10] Matthias Lindemann, Jonas Groschwitz, and Alexander Koller. Fast semantic parsing with well-typedness guarantees. *arXiv preprint arXiv:2009.07365*, 2020.

[11] Jiangming Liu, Shay B Cohen, and Mirella Lapata. Discourse representation structure parsing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 429–439, 2018.

[12] Manuel Mager, Ramón Fernandez Astudillo, Tahira Naseem, Md Arafat Sultan, Young-Suk Lee, Radu Florian, and Salim Roukos. Gpt-too: A language-model-first approach for amr-to-text generation. *arXiv preprint arXiv:2005.09123*, 2020.

[13] Emma Manning, Shira Wein, and Nathan Schneider. A Human Evaluation of AMR-to-English Generation Systems. In *Proceedings of the 28th International Conference on Computational Linguistics (COLING 2020)*, Online, 2020. to appear.

[14] Juri Opitz and Anette Frank. Towards a decomposable metric for explainable evaluation of text generation from amr. *arXiv preprint arXiv:2008.08896*, 2020.

[15] Juri Opitz, Letitia Parcalabescu, and Anette Frank. Amr similarity metrics from principles. *arXiv preprint arXiv:2001.10929*, 2020.

[16] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318, 2002.

[17] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2019.

[18] Linfeng Song, Yue Zhang, Zhiguo Wang, and Daniel Gildea. A graph-to-sequence model for amr-to-text generation. *arXiv preprint arXiv:1805.02473*, 2018.

[19] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112, 2014.

[20] Rik Van Noord and Johan Bos. Neural semantic parsing by character-based translation: Experiments with abstract meaning representations. *arXiv preprint arXiv:1705.09980*, 2017.

[21] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.

[22] Chuan Wang, 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*, pages 366–375, 2015.

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

[24] Sheng Zhang, Xutai Ma, Kevin Duh, and Benjamin Van Durme. Amr parsing as sequence-to-graph transduction. *arXiv preprint arXiv:1905.08704*, 2019.