

Translate, *then* Parse! A strong baseline for Cross-Lingual AMR Parsing

Sarah Uhrig* Yoalli Rezepka García* Juri Opitz Anette Frank

Dept. of Computational Linguistics

Heidelberg University

69120 Heidelberg

{surname}@cl.uni-heidelberg.de

Abstract

In cross-lingual Abstract Meaning Representation (AMR) parsing, researchers develop models that project sentences from various languages onto their AMRs to capture their essential semantic structures: given a sentence in any language, we aim to capture its core semantic content through concepts connected by manifold types of semantic relations. Methods typically leverage large silver training data to learn a single model that is able to project non-English sentences to AMRs. However, we find that a simple baseline tends to be overlooked: translating the sentences to English and projecting their AMR with a monolingual AMR parser (`translate+parse`, T+P). In this paper, we revisit this simple two-step baseline, and enhance it with a strong NMT system and a strong AMR parser. Our experiments show that T+P outperforms a recent state-of-the-art system across all tested languages: German, Italian, Spanish and Mandarin with +14.6, +12.6, +14.3 and +16.0 Smatch points.

1 Introduction

Abstract Meaning Representation (AMR), introduced by [Banarescu et al. \(2013\)](#), aims at representing the meaning of a sentence in a semantic graph format. Nodes represent entities, events and concepts, while (typed) edges express their relations.

AMR itself, as of now, is English-focused, e.g., predicate frames are linked to English PropBank ([Kingsbury and Palmer, 2002](#)). However, the abstract nature of AMR, and the fact that they are not explicitly linked to syntactic structure, make it appealing for extracting semantic structure of sentences in various languages. This insight led to the recent interest in a new task: **cross-lingual AMR parsing** ([Damonte and Cohen, 2018](#)). Here, researchers develop models to project sentences from different languages onto AMR graphs. Mod-

*Equal contribution.




Figure 1: Cross-lingual AMR parsing as introduced by [Damonte and Cohen \(2018\)](#).

els that have recently been proposed are typically trained on large-scale silver data and learn to directly project the non-English sentences onto their AMR graphs (see Figure 1) ([Damonte and Cohen, 2018](#); [Biloshmi et al., 2020](#)). However, there is an intuitive baseline that we argue has so-far received too little attention: `translate+parse`, T+P. It first translates a sentence to a pivot language and applies a mono-lingual parser for that language. In light of the rapid progress of both NMT and AMR parsing models for English, our hypothesis is that this baseline has become more effective and thus more realistic. Moreover, we argue that it could be beneficial to disentangle two key latent representations involved in the process of cross-lingual AMR parsing: i) one that translates between two natural languages and ii) one that translates between a natural language and a meaning representation. This way, the cross-lingual AMR construction process is more transparent and can be better analyzed.

In our work we test these hypotheses by translating the source language sentences into English with a strong NMT system, and parse the resulting English sentences using a strong AMR parser. We show that our baseline delivers strong performance in cross-lingual AMR parsing across all considered languages, outperforming task-focused state-of-the-art models in all settings. We also discuss **fairer evaluation of cross-lingual AMR parsing** and **relevant implications** of this work for research into cross-lingual AMR parsing.

We will release all code under public license.¹

2 Related work

Cross-lingual AMR parsing Cross-lingual AMR parsing was introduced by Damonte and Cohen (2018). They trained an alignment-based AMR parser model that leverages large amounts of parallel silver AMR data obtained through annotation projection from a curated parallel corpus. The authors also discussed `translate+parse` (T+P) as a baseline using either the NMT systems Google translate and Nematus (Sennrich et al., 2017), or the SMT system Moses (Koehn et al., 2007), together with a mono-lingual transition-based parser (Damonte et al., 2017). However, their best T+P approach was Google Translate (GT) – which cannot be fully replicated by other researchers since both training data and model structure are hidden. Given the recent advances in NMT (Barrault et al., 2019, 2020) and mono-lingual AMR parsing (Xu et al., 2020), where parsers now achieve scores on par with human IAA assessments (c.f. Banarescu et al. (2013)), we show that time is ripe to put more spotlight on T+P.

Biloshmi et al. (2020) address the problem from complementary perspectives: i) they train a system that projects AMR graphs from parsed English sentences to target sentences via a parallel corpus, yielding *gold non-English sentences* and *silver AMRs*. Conversely, ii) they train a system that employs an NMT system to translate English sentences from a human-annotated AMR dataset to another language, yielding pairs of *silver non-English sentences* and *gold AMRs*. This alleviates the dependency on external AMR aligners.

(Mono-lingual) AMR parsing Mono-lingual AMR parsing equally made big strides in recent years, so that today AMR parsers deliver benchmark scores that are on-par with measured human IAA. The latest step forward was achieved with neural sequence-to-sequence models pre-trained on large-scale MT benchmark data (Roberts et al., 2020; Xu et al., 2020) or are fine-tuning self-supervised seq-to-seq language models such as T5 or BART (Lewis et al., 2019; Bevilacqua et al., 2021). Previous models perform parsing based on different techniques, e.g., predicting latent alignments jointly with nodes (Lyu and Titov, 2018), or

¹<https://github.com/Heidelberg-NLP/simple-xamr>

via an iterative BFS writing traversal (Cai and Lam, 2019, 2020).

3 Translate, then parse!

Our pipeline model contains two components:

Sent-to-Sent: NMT system We use Helsinki-NLP’s *Opus-MT models* (Tiedemann and Thottungal, 2020) to translate the sentences to English. The models are freely accessible² and provide high scores on public evaluation benchmarks.³

Sent-to-AMR: AMR parser For parsing English target sentences to AMR, we use the parser from *amrlib*⁴, which consists of a T5 language model (Roberts et al., 2020) that has been fine-tuned on English sentences and their AMRs.

4 Experiments

Data We employ the cross-lingual AMR parsing benchmark *LDC2020T07*. It was built from the test split of the English mono-lingual LDC2017T10 data by translating its sentences to four languages: German, Spanish, Italian and Mandarin Chinese. This amounts to a total of 5,484 AMR-sentence pairs, or 1,371 AMR-sentence pairs per language.

Baselines For all languages (German, Spanish, Italian and Mandarin Chinese), we compare against i) AMREAGER (Damonte and Cohen, 2018), and ii) XL-AMR (Biloshmi et al., 2020).

Evaluation metrics Our main evaluation metric is Smatch F1 (Cai and Knight, 2013). The Smatch metric aligns the predicted graph with the gold graph and computes an F1 score that measures normalized triple overlap. Additionally, we calculate F1 scores for finer-grained core semantic sub-tasks (Damonte et al. (2017)).⁵ In our analyses (§4.2), we also study results with S2MATCH (Opitz et al., 2020), that offers a potentially fairer evaluation in cross-lingual AMR parsing, since it does not penalize allowed paraphrases that may emerge, e.g., due

²They are implemented in *EasyNMT*, a SOTA NMT package: <https://github.com/UKPLab/EasyNMT>

³See <https://huggingface.co/Helsinki-NLP> for scores on benchmarks.

⁴<https://github.com/bjascob/amrlib>

⁵i) Unlabeled: score without node labels, ii) No WSD: score w/o predicate sense disambiguation, iii) Reentrancies: score on re-entrant nodes (coreference); iv) Concepts: score on concept nodes; v) Named Ent.: indicating NER performance; vi) negation: polarity detection performance; vii) SRL: semantic role labeling performance.

to the non-monotonous nature of translation (e.g., *huckleberry* → *Heidelbeere* (DE) → *blueberry*).

4.1 Main results

Results are displayed in Table 1. Overall, our `translate+parse` baseline outperforms previous work by large margins. In all assessed semantic categories, T+P outperforms XL-AMR models by more than 10 Smatch points. The smallest improvement obtained is achieved in IT with +12.6 points.

In some key semantic categories, the differences are extreme. For negation detection we obtain performance improvements that range from +26.5 points (IT) to +37.3 (DE). The named entity recognition improves by +20.3 points for German, +20.4 points for Spanish, +17.6 points for Italian.

4.2 Studies

Using a graded metric for evaluation When evaluating predicted AMRs against reference AMRs in cross-lingual AMR parsing, we are essentially comparing AMRs from sentences that are not exactly the same. This means that predicted concepts that are valid may get erroneously penalized by the evaluation metric. For instance, consider a German source sentence that contains *Heidelbeere*, and our cross-lingual AMR system predicts i) *huckleberry* or ii) *blueberry*. Depending on which concept is mentioned in the reference AMR graph (based on the unseen sentence from which the human SemBank annotator created this graph), only one of the two options will be viewed as correct, which results in unfair evaluation. To mitigate this, we propose to conduct the cross-lingual AMR evaluation using S2MATCH (Opitz et al., 2020), a metric that admits graded concept similarity. S2MATCH has a hyper-parameter τ that sets a threshold for sufficiently similar concept nodes across AMRs, using cosine-similarity. The alignment of similar concepts can increase the final score. The default τ is 0.5, but we also try 0.0 which is less strict and fosters dense alignment.

The results are displayed in Table 2. Interestingly, most score improvements are obtained for German (+3.9 points) and Mandarin Chinese (+5.2 points). We conjecture that this is because there is slightly less variety in EN- $\{ES, IT\}$ translations, than for EN-DE, and especially for EN-ZH. This is also visible from the results of our baseline XL-AMR, which we reevaluate using S2MATCH: Most gains are obtained for Mandarin Chinese with an improvement of more than 7 points F1

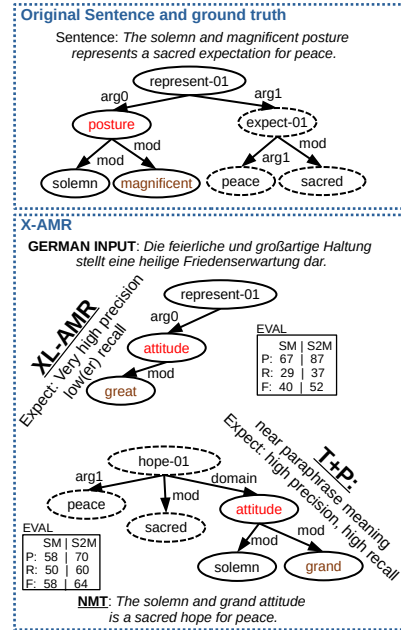


Figure 2: Evaluation example.

score. Inspecting test cases manually, we find many cases where S2MATCH made the evaluation fairer. For instance, the following *gold-pred* (DE: [German word]) concept tuples are ignored by SMATCH but considered by S2MATCH: *pledge-promise* (DE: ‘versprechen’); *write-compose* (DE: ‘verfasst’); *strong-resolute* (DE: ‘deutlich’); *spirit-ghost* (DE: ‘Geist’), etc. In all these cases the cross-lingual AMR system predicted the correct concept, but was penalized by SMATCH. A concrete example case, with lexical (see colored nodes) and structural (see dotted nodes) meaning-preserving divergences, is shown in Fig. 2.

For future work that applies cross-lingual AMR parsing evaluation, we recommend additional evaluation assessment with S2MATCH.

NMT quality The quality of our automatic translations is evaluated with two metrics: i) BLEU score (Papineni et al., 2002) and ii) S(entence)-BERT (Reimers and Gurevych, 2019), in order to assess surface-oriented as well as semantic similarity. For SBERT, we create sentence embeddings for both our translations and the English reference sentences and compute pair-wise cosine similarity.

Looking at the quality of our MT outputs (Table 4), we see that translation quality is generally quite high. The moderate BLEU scores seem to result more from variation in surface form than from incorrect translations, which is backed by the high cosine similarity scores across languages (and also

Metric	AMREAGER				XL-AMR				translate+parse			
	DE	ES	IT	ZH	DE	ES	IT	ZH	DE	ES	IT	ZH
SMATCH	39.1	42.1	43.2	34.6	53.0	58.0	58.1	43.1	67.6	72.3	70.7	59.1
Unlabeled	45.0	46.6	48.5	41.1	57.7	63.0	63.4	48.9	71.9	76.5	75.1	65.4
No WSD	39.2	42.2	42.5	34.7	53.2	58.4	58.4	43.2	67.9	72.7	71.1	60.4
Reentrancies	18.6	27.2	25.7	15.9	39.9	46.6	46.1	34.7	55.8	60.9	58.2	47.5
Concepts	44.9	53.3	52.3	39.9	58.0	65.9	64.7	48.0	71.4	78.1	75.6	63.3
Named Ent.	63.1	65.7	67.7	67.9	66.0	66.2	70.0	60.6	86.3	86.6	87.6	84.2
Negation	18.6	19.8	22.3	6.8	11.7	23.4	29.2	12.8	49.0	59.5	55.7	38.5
SRL	29.4	35.9	34.3	27.2	47.9	55.2	54.7	41.3	61.7	68.0	65.8	54.1

Table 1: F1 Smatch for two baselines and T+P. Best results in **bold**. Improvements > 20 points are underlined.

Metric		XL-AMR				translate+parse			
		DE	ES	IT	ZH	DE	ES	IT	ZH
$\tau=\frac{1}{2}$	S2M P	59.7(+4.3)	63.9(+3.9)	64.7(+4.7)	49.2(+4.7)	74.1(+3.1)	78.3(+2.5)	77.0(+2.8)	65.8(+4.0)
	S2M R	54.8(+4.0)	59.9(+3.7)	59.4(+3.7)	47.3(+5.0)	67.3(+2.9)	71.5(+2.4)	70.0(+2.5)	60.4(+3.7)
	S2M F1	57.1(+4.1)	61.8(+3.1)	62.0(+3.9)	48.2(+5.1)	70.5(+2.9)	74.7(+2.4)	73.4(+2.7)	63.0(+3.9)
$\tau=0.0$	S2M P	61.5(+6.1)	65.4(+5.4)	66.2(+6.2)	51.3(+6.6)	75.2(+4.2)	79.1(+3.3)	77.9(+3.7)	67.1(+5.3)
	S2M R	56.4(+5.6)	61.2(+5.0)	60.8(+5.1)	49.4(+7.1)	68.2(+3.8)	72.1(+3.0)	70.8(+3.3)	61.6(+4.9)
	S2M F1	58.9(+5.9)	63.2(+5.2)	63.4(+5.3)	50.4(+7.3)	71.5(+3.9)	75.4(+3.1)	74.2(+3.5)	64.3(+5.2)

Table 2: Evaluation using a graded metric. **Bold**: Largest improvement of a parser using fairer graded evaluation.

	XL-AMR						translate+parse					
	D-ES	D-I	D-Z	ES-I	ES-Z	I-Z	D-ES	D-I	D-Z	ES-I	ES-Z	I-Z
SMATCH	52.3	52.2	40.9	58.5	42.7	42.8	74.3	73.7	61.6	79.0	63.7	63.1
S2MATCH	58.7	58.6	48.7	63.9	50.2	50.4	77.7	77.2	66.7	81.8	68.6	68.0
Unl.	57.0	57.2	46.6	63.3	48.0	48.7	77.0	76.7	65.7	81.6	67.8	67.1
NoWSD	52.3	52.4	41.0	58.7	42.7	42.8	74.4	73.8	61.7	79.0	63.8	63.1
Conc.	57.7	56.8	45.0	65.4	47.7	47.4	75.1	74.2	63.1	80.0	65.6	64.8
NER	63.9	64.3	56.6	68.2	56.6	57.9	90.1	90.1	83.8	91.2	84.0	84.2
Neg.	7.6	9.7	6.6	44.0	12.8	11.6	61.3	55.4	42.3	69.4	46.2	47.8
Reent.	40.7	41.5	34.6	49.8	36.7	37.2	64.2	63.2	50.4	69.6	51.8	51.0
SRL	47.4	47.5	39.5	55.8	41.7	41.7	69.7	68.9	56.5	75.0	58.3	57.6

Table 3: Semantic consistency over language pairs measured with SMATCH, S2MATCH ($\tau=0$). **Bold/italics** : highest/lowest score for language pairs.

	DE	ES	IT	ZH	mean
BLEU	0.41	0.49	0.46	0.23	0.40
SBERT (cosim)	0.93	0.95	0.94	0.88	0.92

Table 4: BLEU and SBERT MT quality assessment.

highlights the need for a fairer and graded AMR evaluation as proposed above.⁶ Finally, comparing the different source languages, there seems to be a higher quality in the translations from German, Spanish, and Italian, compared to Mandarin Chinese. This is not only reflected in the BLEU scores, but also in the SBERT cosine scores, which suggest a higher semantic similarity between our translations from DE, ES, IT and the reference sentences.

Semantic cross-lingual consistency of cross-lingual AMR systems A cross-lingual AMR system should be expected to deliver the same or highly similar AMRs for two sentences from dif-

⁶This is also supported by a manual analysis of samples.

ferent languages, if the sentences carry the same meaning. We may say that a system is *semantically consistent* if it complies to this expectation.

To measure the degree of consistency, we evaluate system outputs of a cross-lingual AMR system for input language X against the outputs of the same system when fed sentences in language Y, from a *parallel* dataset (X,Y) of sentences in languages X and Y. In the standard evaluation, we computed $EVAL(system(X), A)$ and $EVAL(system(Y), A)$, where A are target AMRs. In this experiment, we instead calculate $EVAL(system(X), system(Y))$, assessing the degree of consistency of a *system*.

The results are provided in Table 3, where we see a very clear picture that holds true both for our joint baseline (XL-AMR) and our T+P approach and all examined semantic categories: the highest consistency is achieved for Spanish-Italian (ES-I, XL-AMR: 63.9 S2MATCH; T+P: 81.8 S2MATCH), while the lowest consistency is achieved for German and Mandarin Chinese (D-Z, XL-AMR: 48.7 S2MATCH; T+P: 66.7 S2MATCH). When directly comparing the parsing systems, overall T+P appears to offer better consistency in all categories, especially negation. However, the substantial variance between languages may indicate that either i) there is a great necessity for making cross-lingual parsers more robust or, ii), that AMR representations, as constructed from English, may be better prepared to represent (besides English) Spanish and Italian language, than, e.g, German or Chinese.

5 Discussion

We believe that the surprising effectiveness of `translate+parse` touches upon a key question: **to what degree can AMR be considered an interlingua?** On one hand, [Banarescu et al. \(2013\)](#) explicitly state that AMR ‘is not designed as an interlingua’. Indeed, AMRs created for English sentences do have a flavour of English, since they are partially grounded in English PropBank ([Kingsbury and Palmer, 2002](#)). But linking AMRs to a PropBank of another language, e.g., Brazilian ([Duran and Aluísio, 2012](#)) or Arabic ([Palmer et al., 2008](#)), and parsing non-English sentences into corresponding AMRs, would not solve, but only displace the problem of being tied to a specific language’s lexical semantic inventory.⁷ On the other hand, AMR *does* contain abstract meaning components that represent language phenomena we may consider as universals: negation, occurrence of named entities, semantic events and their related participants, as well as semantic relations such as Possession, Purpose or Instrument.⁸ We argue that this abstract structure again pushes AMR more towards an interlingua. Hence, the emergent interest in cross-lingual (A)MR ([Oepen et al., 2020](#); [Fan and Gardent, 2020](#); [Sheth et al., 2021](#); [Sherborne and Lapata, 2021](#)) is well justified. However, even if AMR’s *inventory* may favor an interlingual representation, we cannot, in general, expect a homomorphism of AMRs constructed from semantically equivalent sentences in various languages, given wide-spread phenomena that can preclude a uniform AMR representation, such as constructions involving head-switching phenomena or differences in lexical meaning.

Such a middle-ground is indicated by our results: (Too) much divergence may be involved when mapping non-English sentences to original EN-AMRs *directly*, which is penalized by the strict(er) SMATCH metric. We show that evaluation with the softer S2MATCH metric admits small deviations in the conceptual inventory of different languages. The fact that our *indirect* two-step approach `T+P` shows very strong performance also strengthens the view that AMR is not fully an interlingua. The better performance of `T+P` may in part be due to a capacity of strong NMT systems to neutralize some amount of inter-lingual divergence, so that evalu-

⁷Potentially, this may be mitigated in the future by linking AMR to x-lingual PropBanks ([Akbik et al., 2015](#))

⁸C.f. [Xue et al. \(2014\)](#).

ation against EN-AMRs can yield better results in this setting.

Note that in our `T+P` approach two important intermediate (latent) representations are clearly separated: one in the NMT model (that builds a bridge between two natural languages) and one in the parser (that builds a bridge between English and a language of meaning with a flavor of English). By analyzing divergences between source and target in the `T` step, we can uncover aspects of semantic representations that are *not* isomorphic between languages, and which – by transfer via translation – may be neutralized to match the pivot-flavored AMR structure. Hence, the `T+P` approach offers an ideal framework for studying interlingual similarities and divergences in cross-lingual AMR parsing, by comparing the structural-semantic divergences of non-English sentences and their translated English counterparts (aka translational divergences), with the aim of identifying structural-semantic differences between languages that can affect the cross-lingual mapping of sentences into a uniform interlingual AMR.⁹

6 Conclusion

We revisited `translate+parse`, an intuitive baseline for cross-lingual AMR parsing. Equipped with a recent NMT system and a monolingual AMR parser, `T+P` outperforms other approaches by large margins across all evaluation settings. We propose to employ a graded metric for fairer evaluation of cross-lingual AMR parsing. Our work can serve as a strong baseline for future development of cross-lingual AMR parsers. Finally, the `T+P` approach provides an ideal platform for deeper assessment, analysis, and break-down of potential interlingual aspects of AMR.

Acknowledgments

This work has been partially funded by the DFG within the project ACCEPT as part of the Priority Program “Robust Argumentation Machines” (SPP-1999). We thank our anonymous reviewers for constructive comments and valuable feedback.

⁹For example, we may train a system to parse Non-EN sentences to EN-(flavored)AMR graphs, and compare them to AMRs we obtain from *translated EN sentences* targeting the same EN-AMRs. Divergences we find between AMRs predicted in these settings can indicate phenomena leading to non-isomorphic AMRs that require attention when aiming for a true interlingual AMR formalism.

References

- 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.
- 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.
- Loïc Barrault, Magdalena Biesialska, Ondřej Bojar, Marta R. Costa-jussà, Christian Federmann, Yvette Graham, Roman Grundkiewicz, Barry Haddow, Matthias Huck, Eric Joanis, Tom Kocmi, Philipp Koehn, Chi-kiu Lo, Nikola Ljubešić, Christof Monz, Makoto Morishita, Masaaki Nagata, Toshiaki Nakazawa, Santanu Pal, Matt Post, and Marcos Zampieri. 2020. [Findings of the 2020 conference on machine translation \(WMT20\)](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 1–55, Online. Association for Computational Linguistics.
- Loïc Barrault, Ondřej Bojar, Marta R. Costa-jussà, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Philipp Koehn, Shervin Malmasi, Christof Monz, Mathias Müller, Santanu Pal, Matt Post, and Marcos Zampieri. 2019. [Findings of the 2019 conference on machine translation \(WMT19\)](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 1–61, Florence, Italy. Association for Computational Linguistics.
- Michele Bevilacqua, Rexhina Blloshmi, and Roberto Navigli. 2021. [One SPRING to Rule Them Both: Symmetric AMR Semantic Parsing and Generation without a Complex Pipeline](#). *AAAI*.
- 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.
- Deng Cai and Wai Lam. 2019. [Core semantic first: A top-down approach for AMR parsing](#). 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 3799–3809, Hong Kong, China. Association for Computational Linguistics.
- 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.
- Shu Cai and Kevin Knight. 2013. [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, Sofia, Bulgaria. Association for Computational Linguistics.
- Marco Damonte and Shay B. Cohen. 2018. [Cross-lingual Abstract Meaning Representation parsing](#). 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 1146–1155, New Orleans, Louisiana. Association for Computational Linguistics.
- Marco Damonte, Shay B. Cohen, and Giorgio Satta. 2017. [An incremental parser for Abstract Meaning Representation](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 536–546, Valencia, Spain. Association for Computational Linguistics.
- Magali Sanches Duran and Sandra Maria Aluísio. 2012. [Propbank-br: a Brazilian treebank annotated with semantic role labels](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 1862–1867, Istanbul, Turkey. European Language Resources Association (ELRA).
- Angela Fan and Claire Gardent. 2020. [Multilingual AMR-to-text generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2889–2901, Online. Association for Computational Linguistics.
- Paul Kingsbury and Martha Palmer. 2002. [From TreeBank to PropBank](#). In *Proceedings of the Third International Conference on Language Resources and Evaluation (LREC'02)*, Las Palmas, Canary Islands - Spain. European Language Resources Association (ELRA).
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. [Moses: Open source toolkit for statistical machine translation](#). In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180, Prague, Czech Republic. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer

- Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- 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.
- Stephan Oepen, Omri Abend, Lasha Abzianidze, Johan Bos, Jan Hajič, 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.
- Juri Opitz, Letitia Parcalabescu, and Anette Frank. 2020. **AMR similarity metrics from principles**. *Transactions of the Association for Computational Linguistics*, 8:522–538.
- Martha Palmer, Olga Babko-Malaya, Ann Bies, Mona T Diab, Mohamed Maamouri, Aous Mansouri, and Wajdi Zaghouani. 2008. A pilot arabic propbank. In *LREC*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. **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, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. **Sentence-BERT: Sentence embeddings using Siamese BERT networks**. 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 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. How much knowledge can you pack into the parameters of a language model? *arXiv preprint arXiv:2002.08910*.
- Rico Sennrich, Orhan Firat, Kyunghyun Cho, Alexandra Birch, Barry Haddow, Julian Hirschler, Marcin Junczys-Dowmunt, Samuel Läubli, Antonio Valerio Miceli Barone, Jozef Mokry, et al. 2017. Nemat: a toolkit for neural machine translation. *arXiv preprint arXiv:1703.04357*.
- Tom Sherborne and Mirella Lapata. 2021. **Zero-shot cross-lingual semantic parsing**.
- Janaki Sheth, Young-Suk Lee, Ramón Fernandez As-tudillo, Tahira Naseem, Radu Florian, Salim Roukos, and Todd Ward. 2021. **Bootstrapping multilingual AMR with contextual word alignments**. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 394–404, Online. Association for Computational Linguistics.
- Jörg Tiedemann and Santhosh Thottingal. 2020. Opusmt—building open translation services for the world. In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, pages 479–480.
- Dongqin Xu, Junhui Li, Muhua Zhu, Min Zhang, and Guodong Zhou. 2020. **Improving AMR parsing with sequence-to-sequence pre-training**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2501–2511, Online. Association for Computational Linguistics.
- Nianwen Xue, Ondřej Bojar, Jan Hajič, Martha Palmer, Zdeňka Urešová, and Xiuhong Zhang. 2014. **Not an interlingua, but close: Comparison of English AMRs to Chinese and Czech**. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*, pages 1765–1772, Reykjavik, Iceland. European Language Resources Association (ELRA).