

Compositional Semantic Parsing Across Graphbanks

Matthias Lindemann* and Jonas Groschwitz* and Alexander Koller

Department of Language Science and Technology
Saarland University

{mlinde|jonasg|koller}@coli.uni-saarland.de

Abstract

Most semantic parsers that map sentences to graph-based meaning representations are hand-designed for specific graphbanks. We present a compositional neural semantic parser which achieves, for the first time, competitive accuracies across a diverse range of graphbanks. Incorporating BERT embeddings and multi-task learning improves the accuracy further, setting new states of the art on DM, PAS, PSD, AMR 2015 and EDS.

1 Introduction

Over the past few years, a wide variety of semantic *graphbanks* have become available. Although these corpora all pair natural-language sentences with graph-based semantic representations, they differ greatly in the design of these graphs (Kuhlmann and Oepen, 2016). Some, in particular the DM, PAS, and PSD corpora of the SemEval shared task on Semantic Dependency Parsing (Oepen et al., 2015), use the tokens of the sentence as nodes and connect them with semantic relations. By contrast, the AMRBank (Banarescu et al., 2013) represents the meaning of each word using a nontrivial concept graph; the EDS graphbank (Flickinger et al., 2017) encodes MRS representations (Copestake et al., 2005) as graphs with a many-to-many relation between tokens and nodes. In EDS, graph nodes are explicitly aligned with the tokens; in AMR, the alignments are implicit. The graphbanks also exhibit structural differences in their modeling of e.g. coordination or copula.

Because of these differences in annotation schemes, the best performing semantic parsers are typically designed for one or very few specific graphbanks. For instance, the currently best system for DM, PAS, and PSD (Dozat and Manning,

2018) assumes dependency graphs and cannot be directly applied to EDS or AMR. Conversely, top AMR parsers (Lyu and Titov, 2018) invest heavily into identifying AMR-specific alignments and concepts, which may not be useful in other graphbanks. Hershcovich et al. (2018) parse across different semantic graphbanks (UCCA, DM, AMR), but focus on UCCA and do poorly on DM. The system of Buys and Blunsom (2017) set a state of the art on EDS at the time, but does poorly on AMR.

In this paper, we present a single semantic parser that does very well across all of DM, PAS, PSD, EDS and AMR (2015 and 2017). Our system is based on the compositional neural AMR parser of Groschwitz et al. (2018), which represents each graph with its compositional tree structure and learns to predict it through neural dependency parsing and supertagging. We show how to heuristically compute the latent compositional structures of the graphs of DM, PAS, PSD, and EDS. This base parser already performs near the state of the art across all six graphbanks. We improve it further by using pretrained BERT embeddings (Devlin et al., 2019) and multi-task learning. With this, we set new states of the art on DM, PAS, PSD, AMR 2015, as well as (among systems that do not use specialized knowledge about the corpus) on EDS.

2 Semantic parsing with the AM algebra

The *Apply-Modify (AM) Algebra* (Groschwitz et al., 2017; Groschwitz, 2019) builds graphs from smaller graph fragments called *as-graphs*. Fig. 1b shows some as-graphs from which the AMR in Fig. 1a can be constructed. Take for example the graph G_{want} . Some of its nodes are marked with red *sources*, here S and O. These represent ‘argument slots’ to be filled. The O-source in G_{want} is *annotated* with *type* [S], which will be explained below. Further, in each as-graph, one node is marked as a

*Equal contribution

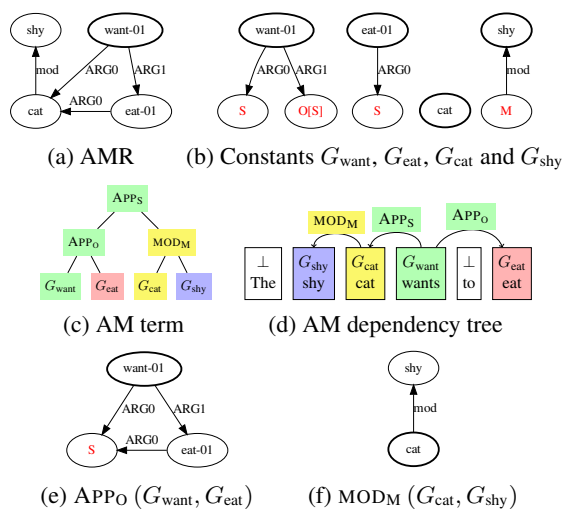


Figure 1: AMR for *The shy cat wants to eat* with its AM analysis.

special *root source*, drawn here with a bold outline.

There are two operations in the AM Algebra that combine as-graphs. First, the *apply* operation APP_X for a source X , as in $APP_O(G_{want}, G_{eat})$ with the result shown in Fig. 1e. The operation combines two as-graphs, a head and an argument, by filling the head’s X -source with the root of the argument. Nodes in both graphs with the same source are unified, i.e. here the two nodes marked with an S-source become one node. The type annotation [S] at the O-source of G_{want} requests the argument to have an S-source (which G_{eat} has). If the argument does not fulfill the request at a source’s annotation, the operation is not *well-typed* and thus not allowed.

The second operation is *modify*, as in $MOD_M(G_{cat}, G_{shy})$ with the result shown in Fig. 1f. Here, G_{cat} is the head and G_{shy} the modifier, and in the operation, G_{shy} attaches with its M source at G_{cat} and loses its own root. We obtain the final graph with the APP_S operation at the top of the term in Fig. 1c, combining the two partial results we have built so far.

AM dependency parsing. By tracking the “semantic heads” of each subtree of an AM term as in Fig. 1c, we can encode AM terms as *AM dependency trees* (Fig. 1d): whenever the AM term combines two graphs with some operation, we add a dependency edge from one semantic head to the other (Groschwitz et al., 2018).

We can then parse a sentence into a graph by predicting an as-graph (or the absence of one, written ‘ \perp ’) for each token in the sentence, along with a well-typed AM dependency tree that connects them. This AM dependency tree evaluates deterministi-

cally to a graph. Groschwitz et al. (2018) show how to perform accurate AMR parsing by training a neural supertagger to predict as-graphs for the words and a neural dependency (tree) parser to predict the AM dependency trees. Here we use their basic models for predicting edge and supertag scores. Computing the highest-scoring well-typed AM dependency is NP-complete; we use Groschwitz et al.’s fixed-tree parser to compute it approximatively.

3 Decomposing the graphbanks

A central challenge with AM dependency parsing is that the AM dependency trees in the training corpus are latent: Strings are annotated with graphs (Fig. 1a), but we need the supertags and AM dependency trees (Fig. 1d).

Groschwitz et al. (2018) describe a heuristic algorithm to obtain AM dependency trees for AMRs (*decomposition*). They first **align** each node in the graph with a word token; then **group** the edges together with either their source or target nodes, depending on the edge label; choose a **source name** for the open slot at the other end of each attached edge; and match reentrancy patterns to determine **annotations** for each source. The dependency edges follow from these decisions.

Groschwitz et al. worked these steps out only for AMR. Here we extend their work to DM, PAS, PSD, and EDS (see Figure 2); this is the central technical contribution of this paper.

3.1 The graphbanks

Before we discuss the decomposition process, let us examine the key similarities and differences of AMR, DM, PAS, PSD and EDS. Most obvious is that DM, PAS and PSD are dependency graphs (Figure 2a-c) where the nodes of the graphs are the words of the sentences, while EDS (Figure 2d) and AMR use nodes related to, but separate from the words. Node-to-word alignments are given in EDS, but not in AMR, where predicting them is hard (Lyu and Titov, 2018).

In all graphbanks we consider here, the edges express semantic relations between the nodes. Several similarities exist: in our example, all graphbanks have edges from “wants” and “eat” to “cat” that indicate that the cat is both the wanter and the eater. These are for example the ‘ARG0’ edges in AMR and the ‘ACT-arg’ edges in PSD. In fact, all five graphs show a triangle structure between “want”, “eat” and “cat” that is characteristic of control verbs.

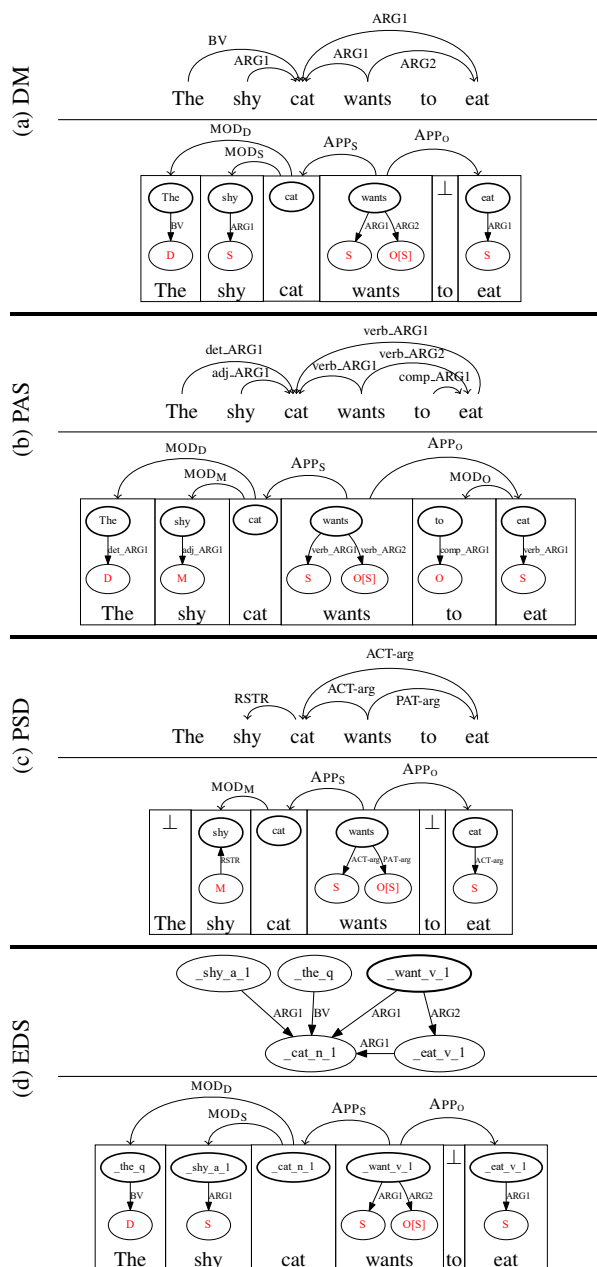


Figure 2: Semantic representations for *The shy cat wants to eat*, each with an AM dependency tree below.

Similarly, all graphs have an edge indicating that “shy” modifies “cat”, although edge label and edge direction vary. However, the graphbanks differ not only in edge directions and labels, but also structurally. For example, DM, PAS and EDS annotate determiners while AMR and PSD do not. Figure 3 shows a reentrancy structure for a copular “are” in PAS that is not present in AMR.

3.2 Our decomposition method

We adapt the decomposition procedure of Groschwitz et al. in the following ways. We sketch the most interesting points here; full details are in the supplementary materials.




Figure 3: AMR (a) and PAS (b) for *Giraffes are tall*.

Alignments are given in EDS and not necessary in DM, PAS, and PSD.

Grouping. We follow two principles in grouping edges with nodes: Edges between heads and arguments always belong with the head, and edges between heads and modifiers with the modifier (regardless of the direction into which the edge points). This yields supertags that generalize well, e.g. a noun has the same supertag no matter whether it has a determiner, whether it is modified by adjectives, whether is agent, and so on.

We find that for all graphbanks, just knowing the edge label is enough to group an edge properly. Thus, we manually decide for each of the 216 edge labels of all graphbanks whether the edges with this label are to be grouped with their target or source node. For instance, ‘ACT-arg’ edges in PSD and ‘verb_ARG1’ edges in PAS are argument-type edges grouped with their source node (they point from a verb to its agent). ‘RSTR’ edges in PSD and ‘adj_ARG1’ edges in PAS are modifier-type and grouped with the adjective; the former is grouped with its target node and the latter with its source. In DM, ‘ARG1’ edges can be both modifier- or argument-type (they are used for both adjectives and verbs); grouping them with their source node is the correct choice in both cases.

Source names. We largely reuse Groschwitz et al.’s source names, which are loosely inspired by (deep) syntactic relations, and map the edge labels of each graphbank to preferred source names. For example, in PSD we associate ‘ACT-arg’ edges with S sources (for “subject”). Some source names are new, such as D for determiners in DM, PAS and EDS (AMRs do not represent determiners).

Annotations. Groschwitz et al.’s algorithm for assigning annotations to sources carries over to the other graphbanks. For patterns that are the same across all graphbanks, such as the ‘triangle’ created by the control verb “want” in Figures 1 and 2, we can re-use the same pattern as for AMR. Thus, control verbs are identified automatically, and their sources are assigned annotations which enforce the appropriate argument sharing.

Interestingly, the original patterns are useful beyond their initial design. We found that for phe-




Figure 4: Coordination in (a) DM and (b, c) PSD.

nomena that cause reentrancies in the new graphbanks, but not in AMR – such as copula in PAS, c.f. Figure 3 – there was typically a suitable pattern designed for a different phenomenon in AMR. E.g. for copula in PAS, the control pattern works.

We thus only update patterns that depend on edge labels; for instance, coordinations in PAS are characterized through their ‘coord_ARGx’ edges.

Challenges with coordination. Coordination in DM (Fig. 4a) is hard to model in the AM algebra because the supertag for “and” would need to consist only of a single ‘and.c’ edge. We group the ‘and.c’ edge with its target node (*Mary*), creating extra supertags e.g. for coordinated and non-coordinated nouns.

In PSD, coordinated arguments (*John and Mary* in Fig. 4b) have an edge into each conjunct. This too is hard to model with the AM algebra because after building *John and Mary*, there can only be one node (the root source) where edges can be attached. We therefore rewrite the graph as shown in Fig. 4c in preprocessing and revert the transformation in postprocessing.

Non-decomposable graphs. While some encodings of graphs as trees are lossy (Agić et al., 2015), ours is not: when we obtain an AM dependency tree from a graph, that dependency tree evaluates uniquely to the original graph. However, not every graph in the training data can be decomposed into an AM dependency tree in the way described above. We mitigate the problem by making DM, PAS, and PSD graphs that have multiple roots connected by adding an artificial root node, and by removing ‘R-HNDL’ and ‘L-HNDL’ edges from EDS (2.3% of edges). We remove some reentrant edges in AMR as described in Groschwitz et al.

We remove the remaining non-decomposable graphs from the training data: 8% of instances in DM, 6% each for PAS and PSD, 24% for EDS, and 10% for AMR. The high percentage of non-decomposable graphs in EDS stems from the fact that EDS can align multiple nodes to the same token, creating multi-node constants. If more than one of these nodes are arguments or are modified in the graph, this cannot be easily represented with

the AM algebra, and thus no valid AM dependency tree is available.

We do not remove graphs from the test data.

4 Evaluation

Data. We evaluate on the DM, PAS and PSD corpora of the SemEval 2015 shared task (Oepen et al., 2015), the EDS corpus (Flickinger et al., 2017) and the releases LDC2015E86 and LDC2017T10 of the AMRBank. All corpora are named entity tagged using Stanford CoreNLP. When tokenization, POS tags and lemmas are provided with the data (DM, PAS, PSD), we use those. Otherwise we employ CoreNLP. We use the same hyperparameters for all graphbanks, as detailed in the appendix.

Parser. We use the BiLSTM-based arc-factored dependency parsing model of Kiperwasser and Goldberg (2016). On the edge existence scores we use the hinge loss of the original K&G model, but we use cross-entropy loss on the edge label predictions; this improved the accuracy of our parser. We also experimented with the dependency parsing model of Dozat and Manning (2017), but this yielded lower accuracies than the K&G model.

We feed each word’s BiLSTM encoding into an MLP with one hidden layer to predict the supertags. We use separate BiLSTMs for the dependency parser and the supertagger but share embeddings. For every token, the BiLSTMs are fed a word embedding, the lemma, POS, and named entity tag. In the basic version of our experiments, we used pretrained GloVe embeddings (Pennington et al., 2014) along with trainable embeddings. In the other version we replace them by pretrained BERT embeddings (Devlin et al., 2019).

AMR and EDS use node labels which are non-trivially related to the words. Therefore, we split each of their supertags into a *delexicalized supertag* and a lexical label. For instance, instead of predicting the supertag G_{want} in Fig. 1b in its entirety, we predict the label “want-01” separately from the rest of the graph. We complement the neural label prediction with a copy function based on the word form and lemma (see supplementary materials).

We implemented this model and Groschwitz et al.’s fixed-tree decoder within the AllenNLP framework (Gardner et al., 2017). Our code is available at <https://github.com/coli-saar/am-parser>.

Results. Table 1 (upper part) shows the results of our basic semantic parser (with GloVe embeddings) on all six graphbanks (mean scores over five

	DM		PAS		PSD		EDS		AMR 2015	AMR 2017
	id F	ood F	id F	ood F	id F	ood F	Smatch F	EDM	Smatch F	Smatch F
Groschwitz et al. (2018)	-	-	-	-	-	-	-	-	70.2	71.0
Lyu and Titov (2018)	-	-	-	-	-	-	-	-	73.7	74.4 ± 0.16
Zhang et al. (2019)	-	-	-	-	-	-	-	-	-	76.3 ± 0.1
Peng et al. (2017) Basic	89.4	84.5	92.2	88.3	77.6	75.3	-	-	-	-
Dozat and Manning (2018)	93.7	88.9	94.0	90.8	81.0	79.4	-	-	-	-
Buyss and Blunsom (2017)	-	-	-	-	-	-	85.5	85.9	60.1	-
Chen et al. (2018)	-	-	-	-	-	-	90.9 ^{1,2}	90.4 ¹	-	-
This paper (GloVe)	90.4 ± 0.2	84.3 ± 0.2	91.4 ± 0.1	86.6 ± 0.1	78.1 ± 0.2	74.5 ± 0.2	87.6 ± 0.1	82.5 ± 0.1	69.2 ± 0.4	70.7 ± 0.2
This paper (BERT)	93.9 ± 0.1	90.3 ± 0.1	94.5 ± 0.1	92.5 ± 0.1	82.0 ± 0.1	81.5 ± 0.3	90.1 ± 0.1	84.9 ± 0.1	74.3 ± 0.2	75.3 ± 0.2
Peng et al. (2017) Freda1	90.0	84.9	92.3	88.3	78.1	75.8	-	-	-	-
Peng et al. (2017) Freda3	90.4	85.3	92.7	89.0	78.5	76.4	-	-	-	-
This paper, MTL (GloVe)	91.2 ± 0.1	85.7 ± 0.0	92.2 ± 0.2	88.0 ± 0.3	78.9 ± 0.3	76.2 ± 0.4	88.2 ± 0.1	83.3 ± 0.1	(70.4) ³ ± 0.2	71.2 ± 0.2
This paper, MTL (BERT)	94.1 ± 0.1	90.5 ± 0.1	94.7 ± 0.1	92.8 ± 0.1	82.1 ± 0.2	81.6 ± 0.1	90.4 ± 0.1	85.2 ± 0.1	(74.5) ³ ± 0.1	75.3 ± 0.1

Table 1: Semantic parsing accuracies (id = in domain test set; ood = out of domain test set).

runs and standard deviations). Our results are competitive across the board, and set a new state of the art for EDS Smatch scores (Cai and Knight, 2013) among EDS parsers which are not trained on gold syntax information. Our EDM score (Dridan and Oepen, 2011) on EDS is lower, partially because EDM evaluates the parser’s ability to align nodes with multi-token spans; our supertagger can only align nodes with individual tokens, and we add alignment spans heuristically.

To test the impact of the grouping and source-naming heuristics from Section 3.2, we experimented with randomized heuristics on DM. The F-score dropped by up to 18 points.

BERT. The use of BERT embeddings is highly effective across the board. We set a new state of the art (without gold syntax) on all graphbanks except AMR 2017; note that Zhang et al. (2019) also use BERT. The improvement is particularly pronounced in the out-of-domain evaluations, illustrating BERT’s ability to transfer across domains.

Multi-task learning. Multi-task learning has been shown to substantially improve accuracy on various semantic parsing tasks (Stanovsky and Dagan, 2018; Hershovich et al., 2018; Peng et al., 2018). It is particularly easy to apply here, because we have converted all graphbanks into a uniform format (supertags and AM dependency trees).

We explored several multi-task approaches during development, namely Freda (Daumé III, 2007; Peng et al., 2017), the Freda generalization of Lu et al. (2016) and the method of Stymne et al. (2018). We found Freda to work best and use it for evalua-

tion. Our setup compares most directly to Peng et al.’s “Freda1” model, concatenating the output of a graphbank-specific BiLSTM with that of a shared BiLSTM, using graphbank-specific MLPs for supertags and edges, and sharing input embeddings.

We pooled all corpora into a multi-task training set except for AMR 2015, since it is a subset of AMR 2017. We also added the English Universal Dependency treebanks (Nivre et al., 2018) to our training set (without any supertags). The results on the test dataset are shown in Table 1 (bottom). With GloVe, multi-task learning led to substantial improvements; with BERT the improvements are smaller but still noticeable.

5 Conclusion

We have shown how to perform accurate semantic parsing across a diverse range of graphbanks. We achieve this by training a compositional neural parser on graphbank-specific tree decompositions of the annotated graphs and combining it with BERT and multi-task learning.

In the future, we would like to extend our approach to sembanks which are annotated with different types of semantic representation, e.g. SQL (Yu et al., 2018) or DRT (Abzianidze et al., 2017). Furthermore, one limitation of our approach is that the latent AM dependency trees are determined by heuristics, which must be redeveloped for each new graphbank. We will explore latent-variable models to learn the dependency trees automatically.

Acknowledgements

We thank Stephan Oepen, Weiwei Sun and Meaghan Fowlie for helpful discussions and the reviewers for their insightful comments. This work was supported by DFG grant KO 2916/2-2.

¹Uses gold syntax information from the HPSG DeepBank annotations at training time.

²Weiwei Sun, p.c.

³Not comparable to other AMR 2015 results because training data contained AMR 2017.

References

- Lasha Abzianidze, Johannes Bjerva, Kilian Evang, Hessel Haagsma, Rik van Noord, Pierre Ludmann, Duc-Duy Nguyen, and Johan Bos. 2017. [The parallel meaning bank: Towards a multilingual corpus of translations annotated with compositional meaning representations](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics*.
- Željko Agić, Alexander Koller, and Stephan Oepen. 2015. Semantic dependency graph parsing using tree approximations. In *Proceedings of the 14th International Conference on Computational Semantics (IWCS)*.
- 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*.
- Jan Buys and Phil Blunsom. 2017. [Robust incremental neural semantic graph parsing](#). In *Proceedings of the 55th Annual Meeting of the 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*.
- 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.
- Ann Copestake, Dan Flickinger, Carl Pollard, and Ivan A Sag. 2005. Minimal recursion semantics: An introduction. *Research on language and computation*, 3(2-3):281–332.
- Hal Daumé III. 2007. [Frustratingly easy domain adaptation](#). In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Timothy Dozat and Christopher D. Manning. 2017. Deep biaffine attention for neural dependency parsing. In *ICLR*.
- Timothy Dozat and Christopher D. Manning. 2018. [Simpler but more accurate semantic dependency parsing](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*.
- Rebecca Dridan and Stephan Oepen. 2011. Parser evaluation using elementary dependency matching. In *Proceedings of the 12th International Conference on Parsing Technologies*, pages 225–230.
- Dan Flickinger, Jan Hajič, Angelina Ivanova, Marco Kuhlmann, Yusuke Miyao, Stephan Oepen, and Daniel Zeman. 2017. [Open SDP 1.2](#). LINDAT/CLARIN digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.
- Matthew Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson H S Liu, Matthew Peters, Michael Schmitz, and Luke S. Zettlemoyer. 2017. A deep semantic natural language processing platform.
- Jonas Groschwitz. 2019. [Methods for taking semantic graphs apart and putting them back together again](#). Ph.D. thesis, Macquarie University and Saarland University.
- Jonas Groschwitz, Meaghan Fowlie, Mark Johnson, and Alexander Koller. 2017. [A constrained graph algebra for semantic parsing with AMRs](#). In *Proceedings of the 12th International Conference on Computational Semantics (IWCS)*.
- Jonas Groschwitz, Matthias Lindemann, Meaghan Fowlie, Mark Johnson, and Alexander Koller. 2018. [AMR Dependency Parsing with a Typed Semantic Algebra](#). In *Proceedings of ACL*.
- Daniel Hershcovich, Omri Abend, and Ari Rappoport. 2018. [Multitask parsing across semantic representations](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*.
- Eliyahu Kiperwasser and Yoav Goldberg. 2016. [Simple and Accurate Dependency Parsing Using Bidirectional LSTM Feature Representations](#). *Transactions of the Association for Computational Linguistics*, 4:313–327.
- Marco Kuhlmann and Stephan Oepen. 2016. [Towards a catalogue of linguistic graph banks](#). *Computational Linguistics*, 42(4):819–827.
- Wei Lu, Hai Leong Chieu, and Jonathan Løfgren. 2016. A general regularization framework for domain adaptation. In *Proceedings of EMNLP*.
- 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*.
- Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014. The stanford corenlp natural language processing toolkit. In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*.

Joakim Nivre, Mitchell Abrams, Željko Agić, Lars Ahrenberg, Lene Antonsen, Katya Aplonova, Maria Jesus Aranzabe, et al. 2018. [Universal dependencies 2.3](#). LINDAT/CLARIN digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.

Stephan Oepen, Marco Kuhlmann, Yusuke Miyao, Daniel Zeman, Silvie Cinková, Dan Flickinger, Jan Hajič, and Zdeňka Urešová. 2015. [Semeval 2015 task 18: Broad-coverage semantic dependency parsing](#). In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*.

Hao Peng, Sam Thomson, and Noah A. Smith. 2017. [Deep Multitask Learning for Semantic Dependency Parsing](#). In *Proceedings of ACL*.

Hao Peng, Sam Thomson, Swabha Swayamdipta, and Noah A. Smith. 2018. [Learning joint semantic parsers from disjoint data](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1492–1502.

Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*.

Gabriel Stanovsky and Ido Dagan. 2018. [Semantics as a foreign language](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*.

Sara Stymne, Miryam de Lhoneux, Aaron Smith, and Joakim Nivre. 2018. [Parser training with heterogeneous treebanks](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*.

Tao Yu, Rui Zhang, Kai Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, Zilin Zhang, and Dragomir Radev. 2018. [Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-sql task](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*.

Shen Zhang, Xutai Ma, Kevin Duh, and Benjamin Van Durme. 2019. AMR parsing as sequence-to-graph transduction. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*.

A Edge Attachment and Source Heuristics

This section presents details of the heuristics discussed in Section 3 of the main paper, concerning grouping (edge attachment) and source heuristics.

Tables 2, 3, 4 and 5 show our edge attachment and source assignment heuristics for DM, PAS, PSD and EDS respectively. The heuristics are broken down by the edge’s label in the ‘Label’ column (* is a wildcard matching any string). A checkmark (✓) in the ‘To Origin’ column means that all edges with this label are attached to their origin node, a cross (✗) means the edge is attached to its target node. In EDS, all edges go with their origin in principle, but in order to improve decomposability, we attach an edge to its target if the target node has label *undef_q* or *nominalization*.

The graphbanks differ in the directionality of edges; in particular, modifier relations sometimes point from the head to the modifier (PSD, AMR) and sometimes from the modifier to the head (DM, PAS, EDS). Our edge assignment heuristics account for that, following the principles for grouping. In DM, for instance, we treat the BV edge (pointing from a determiner to its head noun) as a modifier edge, and thus, it belongs to the determiner, which happens to be at the origin of the edge.

The ‘Source’ column specifies which source is assigned to an empty node attached to an as-graph, depending on the label of the edge with which the node is attached. If an as-graph has multiple attached edges with the same label (or labels that map to the same source), i.e. multiple nodes would obtain the same source, we disambiguate the sources by sorting the nodes with the same source in an arbitrary order and appending ‘2’ to the source at the second node, ‘3’ to the source at the third node and so on (the source at the first node remains unchanged). In PSD, where this happens particularly often, we order the nodes with the same source in their word order rather than arbitrarily, to get more consistent AM dependency trees. For example, if in PSD there are two nodes that are attached to an as-graph with ‘CONJ.member’ edges (such as in Figure 4c in the main paper), the edge going to the left gets assigned an OP source and the edge going to the right an OP2 source.

Passive and object promotion. Following Groschwitz et al. (2018), we allow some source names to be changed or swapped in an as-graph constant after their original assignments. That is, after we build a constant according to the edge grouping and source assignments described above, we generate multiple variants of the constant that have different source names. We allow

Label	To Origin	Source
ARG1	✓	S
ARG2	✓	O
compound	✓	comp
BV	✓	D
poss	✓	poss
*_c	✗	coord
ARG3	✓	O2
mwe	✗	comp
conj	✗	coord
plus	✗	M
ARG4	✓	O3
all other edges	✓	M

Table 2: Heuristics for DM.

Edge label	To Origin	Source
det_ARG1	✓	D
punct_ARG1	✓	punct
coord_ARG _{<i>i</i>}	✓	op _{<i>i</i>}
verb_ARG1	✓	S
*_ARG1	✓	O
*_ARG2	✓	O
all other edges	✓	M

Table 3: Heuristics for PAS.

- object promotion, e.g. instead of an O3 source we may also use an O2 or O source, as long as they don’t exist yet in the constant,
- unaccusative subjects, i.e. an O source may become an S source if no S source is present yet in the constant, and
- passive, i.e. switching O and S sources.

This allows more graphs to be decomposed, by allowing e.g. the coordination of a verb in active and a verb in passive, or the raising of unaccusative subjects. We also follow [Groschwitz et al. \(2018\)](#) in the following (quoted directly from their Section 4.2): “To make our as-graphs more consistent, we prefer constants that promote objects as far as possible, use unaccusative subjects, and no passive alternation, but still allow constants that do not satisfy these conditions if necessary.”

Reentrancy heuristics. We update the reentrancy patterns of [Groschwitz et al. \(2017\)](#) in the following way. No coordination node patterns are allowed in DM (since DM uses edges for coordination); Coordination nodes in PAS are characterized via their coord_ARG_{*i*} edges. In PSD and EDS, any node that has two arguments that themselves have a common argument can be a coordination node.

Raising in PAS is done with the coordination pattern; in the others, a node v where one argument w has an S source can be a raising node (that

Label	To Origin	Source
ACT-arg	✓	S
PAT-arg	✓	O
*-arg (except ACT, PAT)	✓	OO
*.member	✓	op
CPR	✓	M
all other edges	✗	M

Table 4: Heuristics for PSD.

Label	To Origin	Source
ARG1	✓	S
ARG2	✓	O
BV	✓	D
R-INDEX	✓	op1
L-INDEX	✓	op2
R-HNDL	✓	op1
L-HNDL	✓	op2
ARG*	✓	O
all other edges	✓	M

Table 5: Heuristics for EDS.

is, we add an [S]-annotation at the source that the v -constant has at node w), as long as the edge between v and w has label

- ARG1 or ARG2 in DM,
- PAT-arg in PSD, or
- any label in EDS.

We use the same ‘raising’-style pattern for comparatives in PSD, where we use no condition on the source that is ‘passed along’, but the edge from v to w must have label ‘CPR’.

Randomized heuristics. The randomized heuristics we experimented with on the DM set choose edge grouping (to target or to origin) and source names for each edge label independently uniformly at random (but consistently across the corpus).

B Training and Parsing Details

We reimplemented the graph-based parser of [Kiperwasser and Goldberg \(2016\)](#) in AllenNLP. We deviate from the original implementation in the following:

- We use a cross-entropy loss instead of a hinge loss on the edge *label* predictions.
- We follow [Groschwitz et al. \(2018\)](#) in using the Chu-Liu-Edmonds algorithm instead of Eisner’s algorithm.
- We don’t perform word dropout but regular dropout on the input.

We add a supertagger consisting of a separate BiLSTM, from whose states we predict delexicalized graph fragments and lexical labels with an MLP. Learned embeddings are shared between the BiLSTM of the supertagger and the dependency parser.

The hyperparameters are collected in table 6. We train the parser for 40 epochs and pick the model with the highest performance on the development set (measured in Smatch for EDS, not in EDM). We perform early stopping with patience of 10 epochs. Every lemma (and word in the case of using GloVe embeddings) that occurs fewer than 7 times is treated as unknown.

We use BucketIterators (padding noise 0.1) and the methods implemented in AllenNLP for performing padding and masking.

Training with GloVe We use the 200-dimensional version of GloVe (6B.200d) along with 100-dimensional trainable embeddings. We use two layers in the BiLSTMs and train with a batch size of 48.

Training with BERT When using BERT, we replace both the GloVe embeddings and the learned word embedding with BERT. Since BERT does not provide embeddings for the artificial root of the dependency tree, we learn a separate embedding. In some graphbanks (DM, PAS, PSD), we also have an artificial word at the end of each sentence, that is used to connect the graphs. From BERT’s perspective, the artificial word is a period symbol.

When training with BERT, we use a batch size of 64 and only one layer in the BiLSTMs. We use the “large-uncased” model as available through AllenNLP and don’t fine-tune BERT.

MTL In our Freda experiments, we have one LSTM per graphbank and one that is shared between the graphbanks. When we compute scores for a sentence, we run it through its graphbank-specific LSTM and the shared one. We concatenate the outputs and feed it to graphbank-specific MLPs. Again, we have separate LSTM for the edge model (input to edge existence and edge label MLP) and the supertagging model. In effect, we have two LSTMs that are shared over the graphbanks: one for the edge model and one for the supertagging model.

All LSTMs have the hyperparameters detailed in table 6. In the case of UD, we don’t use a graphbank-specific supertagger because there are

Activation function	tanh
Optimizer	Adam
Learning rate	0.001
Epochs	40
Dim of lemma embeddings	64
Dim of POS embeddings	32
Dim of NE embeddings	16
Minimum lemma frequency	7
Hidden layers in all MLPs	1
Hidden units in LSTM (per direction)	256
Hidden units in edge existence MLP	256
Hidden units in edge label MLP	256
Hidden units in supertagger MLP	1024
Hidden units in lexical label tagger MLP	1024
Layer dropout in LSTMs	0.3
Recurrent dropout in LSTMs	0.4
Input dropout	0.3
Dropout in edge existence MLP	0.0
Dropout in edge label MLP	0.0
Dropout in supertagger MLP	0.4
Dropout in lexical label tagger MLP	0.4

Table 6: Common hyperparameters used in all experiments.

no supertags for UD. We don’t pool the UD treebanks together.

In the MTL setup, we select the epoch with the highest development F-score for DM for evaluation on all test sets.

Parsing We follow Groschwitz et al. (2018) in predicting the best unlabeled dependency tree with the Chu-Liu-Edmonds algorithm and then run their fixed-tree decoder restricted to the 6 best supertags. This computes the best well-typed AM dependency tree with the same shape as the unlabeled tree.

Parsing is usually relatively fast (between 30 seconds and 2 minutes for the test corpora) but very slow for a few sentences very long sentences in the AMR test corpora. Therefore, we set a timeout. If parsing with k supertags is not completed within 30 minutes, we retry with $k - 1$ supertags. If $k = 0$, we use a dummy graph with a single node. This happened 4 times over different runs on AMR with the basic version of the parser and once when using BERT.

Copy function In order to predict the lexical label for EDS and AMR, we predict only the difference to its lemma or word form. For instance, if the lexical label is “want-01”, we try to predict `$LEMMA$-01` instead at the word in question, e.g. wanted, and restore the full form of the lexical label in postprocessing.

C Details of Preprocessing and Postprocessing

DM, PAS and PSD We handle disconnected graphs with components that contain more than one node by adding an artificial word to the end of the sentence. We draw an edge from this word to one node in every weakly connected component of the graph. We select this node by invoking Stanford CoreNLP (Manning et al., 2014) to find the head of the span the component comprises.

Disconnected components that only contain one word are treated as words without semantic contribution, which we attach to the artificial root (position 0) with an IGNORE-edge.

Since the node labels in these graphbanks are the words of the sentences, we simply copy the words over to the graph.

We use the evaluation toolkit that was developed for the shared task: <https://github.com/semantic-dependency-parsing/toolkit>.

EDS We only consider connected EDS graphs (98.5% of the corpus) and follow Buys and Blunsom (2017) regarding options for the tokenizer except for hyphenated words, which we split. Since EDS nodes are aligned with (character) spans in the sentence, we make use of this information in the decomposition. In our approach, however, we require every graph constant to stem from exactly one token. In order to enforce this, we assign nodes belonging to a multi-token span to an atomic span whose nodes are incident. For consistency, we perform this from left to right. We try to avoid creating graph constants that would require more than one root source. Where this fails, the graph cannot be decomposed.

We delete R-HNDL and L-HNDL edges only if this does not make the graph disconnected. Thus, we need heuristics for them (see table 5).

Before delexicalizing graphs constants, we need to identify lexical nodes. A node is considered lexical if has an incoming c-arg edge or if its label is similar to the aligned word, its lemma or its modified lemma. We compute the modified lemma by a few hand-written rules from the CoreNLP lemma. For instance, "Tuesday" is mapped to "Tue". We also re-inflect adverbs (as identified by the POS tagger) to their respective adjectives if possible, e.g. "interestingly" becomes "interesting". We perform this step in order to be able to represent the lexical label of more graph constants as function of the

word which they belong to. The modified lemma is not used as input to the neural network.

When performing the delexicalization, we replace the character span information with placeholders indicating if this span is atomic (comprises a single word) or not. We restore the span information for every node with a very simple heuristic in postprocessing: If the span is atomic, we simply look up the character span in the original string. For nodes with complex spans, we compute the minimum of beginnings and the maximum of endings of its children. In terms of evaluation, the span information is relevant only for EDM. Comparing the graphs that we restore from our training data to the gold standard, we find that the upper bound is at 89.7 EDM F-score. The upper bound in terms of Smatch is at 96.9 F-score.

We use EDM in an implementation by Buys and Blunsom (2017).

UD Since UD POS tags are different from the English PTB tagset, we use CoreNLP to tag the UD treebanks. We use the English treebanks EWT, GUM, ParTUT and LinES (Nivre et al., 2018).

AMR We use the pre- and postprocessing pipeline of Groschwitz et al. (2018). We conflate named entities in preprocessing. For instance, "New York" is conflated to one token "New_York". When such a graph constant is predicted, we restore the named entity prior to evaluation.