

Learning compositional structures for semantic graph parsing

Jonas Groschwitz

Saarland University

jonasg@coli.uni-saarland.de

Meaghan Fowlie

Utrecht University

m.fowlie@uu.nl

Alexander Koller

Saarland University

koller@coli.uni-saarland.de

Abstract

AM dependency parsing is a method for neural semantic graph parsing that exploits the principle of compositionality. While AM dependency parsers have been shown to be fast and accurate across several graphbanks, they require explicit annotations of the compositional tree structures for training. In the past, these were obtained using complex graphbank-specific heuristics written by experts. Here we show how they can instead be trained directly on the graphs with a neural latent-variable model, drastically reducing the amount and complexity of manual heuristics. We demonstrate that our model picks up on several linguistic phenomena on its own and achieves comparable accuracy to supervised training, greatly facilitating the use of AM dependency parsing for new sembanks.

1 Introduction

It is generally accepted in linguistic semantics that meaning is *compositional*, i.e. that the meaning representation for a sentence can be computed by evaluating a tree bottom-up. A compositional parsing model not only reflects this insight, but has practical advantages such as in compositional generalisation (e.g. [Herzig and Berant 2020](#)), i.e. systematically generalizing from limited data.

However, in developing a compositional semantic parser, one faces the task of figuring out what exactly the compositional structures – i.e. the trees that link the sentence and the meaning representation – should look like. This is challenging even for expert linguists; for instance, ([Copestake et al., 2001](#)) report that 90% of the development time of the English Resource Grammar ([Copestake and Flickinger, 2000](#)) went into the development of the syntax-semantics interface.

Compositional semantic parsers which are learned from data face an analogous problem: to

train a such a parser, the compositional structures must be made explicit. However, these structures are not annotated in most sembanks. For instance, the *AM (Apply-Modify) dependency parser* of [Groschwitz et al. \(2018\)](#) uses a neural model to predict *AM dependency trees*, compositional structures that evaluate to semantic graphs. Their parser achieves high accuracy ([Lindemann et al., 2019](#)) and parsing speed ([Lindemann et al., 2020](#)) across a variety of English semantic graphbanks. To obtain an AM dependency tree for each graph in the corpus, they use hand-written graphbank-specific heuristics. These heuristics cost significant time and expert knowledge to create, limiting the ability of the AM parser to scale to new sembanks.

In this paper, we drastically reduce the need for hand-written heuristics for training the AM dependency parser. We first present a graphbank-independent method to compactly represent the relevant compositional structures of a graph in a tree automaton. We then train a neural AM dependency parser directly on these tree automata. Our code is available at github.com/coli-saar/am-parser.

We evaluate the consistency and usefulness of the learned compositional structures in two ways. We first evaluate the accuracy of the trained AM dependency parsers, across four graphbanks, and find that it is on par with an AM dependency parser that was trained on the hand-designed compositional structures of [Lindemann et al. \(2019\)](#). We then analyze the compositional structures which our algorithm produced, and find that they are linguistically consistent and meaningful. We expect that our methods will facilitate the design of compositional models of semantics in the future.

2 Related work

Compositional semantic graph parsers other than AM dependency parsers, like [Artzi et al. \(2015\)](#),

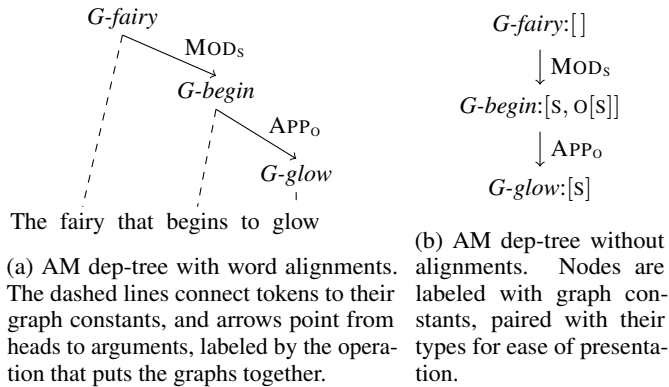


Figure 1: AM dep-trees and graphs for *the fairy that begins to glow*. We usually write our example AM dep-trees without alignments as in (b). We include node names where helpful, as in (c), where e.g. *b* is labeled *begin*.




Figure 2: Graph constants

Peng et al. (2015) and Chen et al. (2018), use CCG and HRG based grammars to parse AMR and EDS (Flickinger et al., 2017). They use a combination of heuristics, hand-annotated compositional structures and sampling to obtain training data for their parsers, in contrast to our joint neural technique. None of these approaches use slot names that carry meaning; to the best of our knowledge this work is the first to learn them from data.

Fancellu et al. (2019) use DAG grammars for compositional parsing of Discourse Representation Structures (DRS). Their algorithm for extracting the compositional structure of a graph is deterministic and graphbank-independent, but comes at a cost: for example, rules for heads require different versions depending on how often the head is modified, reducing the reusability of the rule.

Maillard et al. (2019) and Havrylov et al. (2019) learn compositional, continuous-space neural sentence encodings using latent tree structures. Their tasks are different: they learn to predict continuous-space embeddings; we learn to predict symbolic compositional structures. Similar observations hold for self-attention (Vaswani et al., 2017; Kitaev and Klein, 2018).

3 AM dependency parsing

Compositional semantic graph parsing methods do not predict a graph directly, but rather predict a compositional structure which in turn determines the graph. Groschwitz et al. (2018) represent the compositional structure of a graph with *AM dependency trees* (AM dep-trees for short) like the one in Fig. 1a. It describes the way the meanings of the words – the graph fragments in Fig. 2 – combine to form the semantic graph in Fig. 1c, here an AMR

(Banarescu et al., 2013). The AM dep-tree edges are labeled with graph-combining operations, taken from the *Apply-Modify (AM) algebra* (Groschwitz et al., 2017; Groschwitz, 2019).

Graphs are built out of fragments called *graph constants* (Fig. 2). Each graph constant has a *root*, marked with a rectangular outline, and may have special node markers called *sources* (Courcelle and Engelfriet, 2012), drawn in red, which mark the empty slots where other graphs will be inserted.

In Fig. 1a, the APP_O operation plugs the root of *G-glow* into the O source of *G-begin*. Because *G-begin* and *G-glow* both have an S-source, APP_O merges these nodes, creating a *reentrancy*, i.e. an undirected cycle, and yielding Fig. 1d, which is in turn attached at S to the root of *G-fairy* by MOD_S . APP fills a source of a head with an argument while MOD uses a source of a modifier to connect it to a head; both operations keep the root of the head.

Types The [S] annotation at the O-source of *G-begin* in Fig. 2 is a *request* as to what the *type* of the O argument of *G-begin* should be. The type of a graph is the set of its sources with their request annotations, so the request [S] means that the source set of the argument must be {S}. Because this is true of *G-glow*, the AM dependency tree is *well-typed*; otherwise the tree could not be evaluated to a graph. Thus, the graph constants lexically specify the semantic valency of each word as well as reentrancies due to e.g. control.

If a graph has no sources, we say it has the *empty type* []; if a source in a graph printed here has no annotation, it is assumed to have the empty request (i.e. its argument must have no sources).

Parsing Groschwitz et al. (2018) use a neural supertagger and dependency parser to predict scores

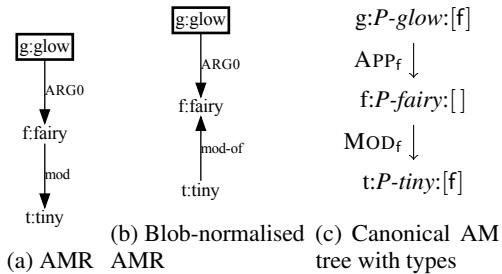


Figure 3: *The tiny fairy glows.*

for graph constants and edges respectively. Computing the highest scoring well-typed AM dep-tree is NP-hard; we use their fixed-tree approximate decoder here.

4 Decomposition algorithm

The central challenge of compositional methods lies in the fact that the compositional structures are not provided in the graphbanks. Existing AM parsers (Groschwitz et al., 2018; Lindemann et al., 2019, 2020) use hand-built heuristics to extract AM dep-trees for supervised training from the graphs in the graphbank. These heuristics require extensive expert work, including graphbank-specific decisions for source allocations and graphbank- and phenomenon-specific patterns to extract type requests for reentrancies. In this section we present a simpler yet more complete method for obtaining the basic structure of an AM dep-tree for a given semantic graph G (for *decomposing* the graph), with much reduced reliance on heuristics. We will learn meaningful source names jointly with training the parser in §5 and §6.

Notation. We treat graphs as a quadruple $G = \langle N_G, r_G, E_G, L_G \rangle$, where the nodes N_G are arbitrary objects (in the examples here we use lower-case letters), $r_G \in N_G$ is the root, $E_G \subseteq N_G \times N_G$ is a set of directed edges, and L_G is the labelling function for the nodes and edges. For example in Fig. 3a, the node g is labeled “glow”. The node identities are not relevant for graph identity or evaluation measures, but allow us to refer to specific nodes during decomposition. We formalize AM dep-trees as similar quadruples. Note that our example graphs are all AMRs, but our algorithms apply unchanged to all graphbanks

4.1 Basic transformation to AM dep-trees

Let us first consider the case where the semantic graph G has no reentrancies, like in Fig. 3a. The first step in obtaining the AM dep-tree for G is




Figure 4: Canonical constants.




Figure 5: Non-canonical constants with placeholder sources.

to obtain the basic shape of the constants. We let each graph constant contain exactly one labeled node. Each edge belongs to the constant of exactly one node. The edges in the constant of a node are called its **blob** (Groschwitz et al., 2017); the blobs partition the edge set of the graph. For example, the blobs of the AMR in Fig. 3a are g plus the ‘ARG0’ edge, t plus the ‘mod’ edge, and f . We normalise edges so that they point away from the node to whose blob they belong, like in Fig. 3b, where the ‘mod’ edge is reversed and grouped with the node t to match P -tiny in Fig. 4. We add an *-of* suffix to the label of reversed edges. From here on, we assume all graph edges to be normalised this way.

Heuristics for this partition of edges into blobs are simple yet effective. Thus, this is the only part of this method where we still rely on graphbank-specific heuristics. (We use the same blob heuristics as Lindemann et al. (2019) in our experiments).

Once the decision of which edge goes in which blob is made, we obtain **canonical constants**, which are single node constants using **placeholder source names** and the empty request at every source; see e.g. P -glow in Fig. 4 (P for ‘placeholder’). Placeholder source names are graph-specific source names: for a given argument slot in a constant, let n be the node that eventually fills it in G ; we write n for the placeholder source in that slot. For example in the AM dep-tree in Fig. 3c the source f in P -glow (Fig. 4) gets filled by node f in the AMR in Fig. 3b. These placeholder sources are unique within the graph, allowing us to track source names through the AM dep-tree. When we restrict ourselves to the canonical constants, in a setting without reentrancies, the compositional structure is fully determined by the structure of the graph:

Lemma 4.1. *For a graph G without reentrancies, given a partition of G into blobs, there is exactly one AM dep-tree C_G with canonical constants that evaluates to G .*

We call this AM dep-tree the **canonical AM tree**




Figure 6: Analysis for *The fairy sparkles and glows*.

$C_G = \langle N_G, r_G, E_C, L_C \rangle$ of G . Fig. 3c shows the canonical AM tree for the graph in Fig. 3b, using the canonical constants in Fig. 4. The canonical AM tree uses the same nodes and root as G , and essentially the same edges, but all edges point away from the root, forming a tree. Each node is labeled with its canonical constant. Each edge $n \rightarrow m \in E_C$ is labeled APP_m if the corresponding edge in the graph has the same direction, and is labeled MOD_n if there is instead an edge $m \rightarrow n$ in G .

4.2 Reentrancies and types

Finding AM dep-trees for graphs *with* reentrancies, like in Fig. 6a, is more challenging. To solve the problem in its generality, we first *unroll* the graph as in Fig. 6b, representing the reentrancy at f not directly, but with a *reference node* with label $\text{REF-}f$. Merging this $\text{REF-}f$ node with the node f it refers to yields the original graph again. (See §4.3 for our unrolling algorithm.) An unrolled graph U shares its non- REF- nodes with the original graph G . REF- nodes are always leaves.

We then obtain a canonical AM-tree C_U for the unrolled graph U as in §4.1 (see Fig. 6c), but $\text{REF-}n$ nodes fill n -sources; e.g. x has an incoming APP_f edge here. C_U evaluates to U , not to G ; we obtain an AM dep-tree that evaluates to G through a process called *resolving* the reentrancies, which removes all REF- nodes and instead expresses the reentrancies with the AM type system.

Fig. 6e shows the result T of applying this resolution process to C_U in Fig. 6c. In T , the s and g sources of the graph P' -and (see Fig. 5) each have a request $[f]$ that signals that the f sources of P -sparkle and P -glow are still open when these graphs combine with P' -and, yielding the partial

Algorithm 1: Reentrancy resolution

```

1  $T \leftarrow$  the canonical AM-tree  $C_U$  of an
  unrolling  $U$  of  $G$ ;
2  $R \leftarrow \{n \in N_G \mid \exists \text{REF-}n \text{ node in } U\}$ ;
3 while  $R \neq \emptyset$ :
4   Pick a  $y \in R$  s.t. there is no  $x \in R$ ,
      $x \neq y$ , with  $y$  on an  $x$ -resolution path;
5   for  $p \in y$ -resolution paths:
6     for  $n \xrightarrow{\text{APP}} m \in p$ :
7       if  $m$  is  $y$  or labeled  $\text{REF-}y$ :
8         Add  $\beta(y)$  to the request at  $y$ 
          in  $\tau(n)$ ;
9       else:
10        Add  $y[\beta(y)]$  to the request
          at  $m$  in  $\tau(n)$ ;
11   Move the subtree of  $T$  rooted at  $y$  up to
     be an  $\text{APP}_y$  daughter of  $\text{RT}(y)$ , unless
      $\text{RT}(y) = y$ ;
12   Delete all  $\text{REF-}y$  nodes from  $T$ ;
13    $R \leftarrow R - \{y\}$ 
14 return  $T$ 

```

result in Fig. 6d. Since identical sources merge in the AM algebra, Fig. 6d has a single f -source slot. Into this slot, P -fairy is inserted to yield the original graph G in Fig. 6a, and we have obtained the reentrancy without using a REF- node. f is now a child of a in T ; we call a the **resolution target** of f , $\text{RT}(f)$. In general the resolution target of a node n is the lowest common ancestor of n and all nodes labeled $\text{REF-}n$.

Thus, to resolve the graph, we (a) add the necessary type requests to account for sources remaining open until they are merged at the resolution target and (b) make each node a dependent of its resolution target and remove all REF- nodes. Algorithm 1 describes this procedure. It uses the idea of an ***n-resolution path***, which is a path between a node n or a $\text{REF-}n$ node and its resolution target. In Fig. 6c, there are two f -resolution paths: one in blue between f and its resolution target a , and one in green between the $\text{REF-}f$ node x and its resolution target a . Further, $\tau(n)$ is the type of the graph constant in T for a node n and $\beta(n)$ is the type of the result of evaluating the subtree below n in T .

In the example, Algorithm 1 iterates over all edges in both resolution paths (Line 6; the order of these iterations does not impact the result). For the two bottom edges $s \xrightarrow{\text{APP}_f} f$ and $g \xrightarrow{\text{APP}_f} x$, Line 8 applies. Since the subtree rooted at f evaluates to

a constant with empty type, no actual changes are made here ($\beta(y)$ can be non-trivial from resolution paths handled previously). For the two upper edges $a \xrightarrow{\text{APP}_s} s$ and $a \xrightarrow{\text{APP}_g} g$, Line 10 applies, adding f to the requests at s and g in the constant at a . In Line 11, f gets moved up to become a child of its resolution target a and in Line 12 the REF- f node x gets removed, yielding T in Fig. 6e. Algorithm 1 is correct in the following precise sense:

Theorem 1. *Let G be a graph, let U be an unrolling of G , let C_U be the canonical AM-tree of U , and let T be the result of applying Algorithm 1 to C_U . Then T is a well-typed AM dep-tree that evaluates to G iff for all $y \in N_G$, for all y -resolution paths p in C ,*

1. *the bottom-most edge $n \rightarrow m$ of p (i.e. m is y or labeled REF- y) does not have a MOD label, and*
2. *for all y -resolution paths p in C , if $n \xrightarrow{\text{MOD}} m \in p$, $n, m \neq y$, then there is a directed path in G from n to y .*

Condition (1) captures the fact that moving MOD edges in the graph changes the evaluation result (the modifier would attach at a different node) and Condition (2) the fact that modifiers are not allowed to add sources to the type of the head they modify.

Algorithm 1 does not yield *all* possible AM dep-trees; in Appendix B, we present an algorithm that yields all possible AM dep-trees (with placeholder sources) for a graph. However, we find in practice that Algorithm 1 almost always finds the best linguistic analysis; i.e. reasons to deviate from Algorithm 1 are rare (we estimate that this affects about 1% of nodes and edges in the AM dep-tree). We leave handling these rare cases to future work.

4.3 Unrolling the graph

To obtain an unrolled graph U , we use Algorithm 2. The idea is to simply expand G through breadth-first search, creating REF-nodes when we encounter a node a second time. We use separate queues F and B for forward and backward traversal of edges, allowing us to avoid traversing edges backwards wherever possible, since that would yield MOD edges in the canonical AM-tree C_U , which can be problematic for the conditions of Theorem 1. And indeed, we can show that whenever there is an unrolled graph U satisfying the conditions of Theorem 1, Algorithm 2 returns one.

Algorithm 2 does not specify the order in which the incident edges of each node n are added to the

Algorithm 2: Unrolling

Input: Graph G

```

1  $F, B \leftarrow$  empty FIFO queues;
2  $U \leftarrow$  empty graph;
3 add  $r_G$  to  $U$ , add outgoing edges of  $r_G$  to  $F$ 
  and incoming edges of  $r_G$  to  $B$ ;
4 while  $F \cup B \neq \emptyset$ :
5   if  $F \neq \emptyset$ : // traverse forward
6      $e \leftarrow F.\text{pop}$ ;
7      $n \leftarrow e.\text{target}$ ;
8   else: // traverse backward
9      $e \leftarrow B.\text{pop}$ ;
10     $n \leftarrow e.\text{origin}$ ;
11   Mark  $e$  as traversed;
12   if  $n \notin N_U$ :
13     add  $n, e$  to  $U$ ;
14     add untraversed outgoing edges of  $n$ 
      to  $F$  and incoming to  $B$ 
15   else:
16     add new  $x$  to  $N_U$ ;  $L(x) = \text{REF-}n$ ;
17     add  $e'$  to  $E_U$  where  $e'$  is just like  $e$ 
      except with  $x$  in place of  $n$ 
18 return  $U$ 

```

queues, leaving an element of choice. However, we find that nearly all of these choices are unified later in the resolution process; meaningful choices are rare. For example in Fig. 6b, f and x may be switched, but Algorithm 1 always yields the AM dep-tree in Fig. 6e. In practice, we execute Algorithm 2 with arbitrary queueing order, and follow it with Algorithm 1. The AM dep-tree we obtain is guaranteed to be a decomposition of the original graph whenever one exists:

Theorem 2. *Let G be a graph partitioned into blobs. If there is a well-typed AM dep-tree T , using that blob partition, that evaluates to G , then Algorithm 2 (with any queueing order) and Algorithm 1 yield such a tree.*

5 Tree automata for source names

We have now seen how, for any graph G , we obtain a unique AM dependency tree T . This tree represents the compositional structure of G , but it still contains placeholder source names. We will now show how to automatically choose source names. These names should be consistent across the trees for different sentences; this yields reusable graph constants, which capture linguistic generalizations and permit more accurate parsing. But the source

names must also remain consistent *within* each tree to ensure that the tree still evaluates correctly to G ; for instance, if we replace the placeholder source f in P -glow in Fig. 6e by O , but we replace f in P' -and by S , then the AM dep-tree would not be well-typed because the request is not satisfied.

We therefore proceed in two steps. In this section, we represent all internally consistent source assignments compactly with a tree automaton. In §6, we then learn to select globally reusable source names jointly with training the neural parser.

Tree automata. A (bottom-up) tree automaton (Comon et al., 2007) is a device for compactly describing a language (set) of trees. It processes a tree bottom-up, starting at the leaves, and nondeterministically assigns states from a finite set to the nodes. A rule in a tree automaton has the general shape $f(q_1, \dots, q_n) \rightarrow q$. If the automaton can assign the states q_1, \dots, q_n to the children of a node π with node label f , this rule allows it to assign the state q to π . The automaton *accepts* a tree if it can assign a *final state* to the root node. Tree automata can be seen as generalisation of parse charts.

General construction. Given an AM dependency tree T with placeholders, we construct a tree automaton that accepts all well-typed variants of T with consistent source assignments. More specifically, let \mathcal{S} be a finite set of reusable source names; we will use $\mathcal{S} = \{S, O, M\}$ here, evoking subject, object, and modifier. The automaton will keep track of *source name assignments*, i.e. of partial functions ϕ from placeholder source names into \mathcal{S} . Its rules will ensure that the functions ϕ assign source names consistently.

We start by binarizing T into a binary tree B , whose leaves are the graph constants in T and whose internal nodes correspond to the edges of T ; the binarized tree for the dependency tree in Fig. 7a is shown in Fig. 7b. We then construct a tree automaton A_B that accepts binarized trees which are isomorphic to B , but whose node labels have been replaced by graph constants and operations with reusable source names. The states of A_B are of the form $\langle \pi, \phi \rangle$, where ϕ is a source name assignment and π is the address of a node in B . Node addresses $\pi \in \mathbb{N}^*$ are defined recursively: the root has the empty address ϵ , and the i -th child of a node at address π has address πi . The final states are all states with $\pi = \epsilon$, indicating that we have reached the root.

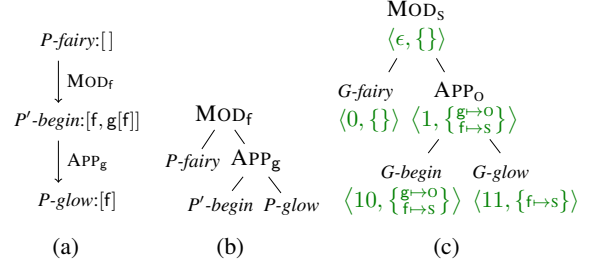


Figure 7: (a) AM dep-tree with placeholder sources for the graph in Fig. 1c, (b) its binarization B and (c) example automaton run (states in green).

Rules. The automaton A_B has two kinds of rules. *Leaf rules* choose injective source name assignments for constants; there is one rule for every possible assignment at each constant. That is, for every graph constant H at an address π in B , the automaton A_B contains all rules of the form

$$G \mapsto \langle \pi, \phi \rangle$$

where ϕ is an injective map from the placeholder sources in H to \mathcal{S} , and G is the graph constant identical to H except that each placeholder source s in H has been replaced by $\phi(s)$.

For example, the automaton for Fig. 7b contains the following rule:

$$G\text{-begin} \rightarrow \langle 00, \{g \mapsto O, f \mapsto S\} \rangle$$

Note that this rule uses the node label G -begin with the reusable source names, not the graph constant P' -begin in B with the placeholders.

In addition, *operation rules* percolate source assignments from children to parents. Let APP_x for some placeholder source x be the operation at address π in B . Then A_B contains all rules of the form

$$APP_{\phi_1(x)}(\langle \pi 0, \phi_1 \rangle, \langle \pi 1, \phi_2 \rangle) \rightarrow \langle \pi, \phi_1 \rangle$$

as long as ϕ_1 and ϕ_2 are identical where their domains overlap, i.e. they assign consistent source names to the placeholders. The rule passes ϕ_1 on to its parent. The assignments in ϕ_2 are either redundant, because of overlap with ϕ_1 , or they are no longer relevant because they were filled by operations further below in the tree. The MOD case works out similarly.

In the example, A_B contains the rule

$$APP_O(\langle 10, \phi_b \rangle, \langle 11, \phi_g \rangle) \rightarrow \langle 1, \phi_b \rangle$$

where $\phi_b = \{g \mapsto O, f \mapsto S\}$ and $\phi_g = \{f \mapsto S\}$, because ϕ_b and ϕ_g agree on f . A complete accepting run of the automaton is shown in Fig. 7c.

The automaton A_B thus constructed accepts the binarizations of all well-typed AM dependency trees with sources in \mathcal{S} that match T .

6 Joint learning of compositional structure and parser

As a final step, we train the neural parser of Groschwitz et al. (2018) directly on the tree automata. For each position i in the sentence, the parser predicts a score $c(G, i)$ for each graph constant G , and for each pair i, j of positions and operation ℓ , it predicts an edge score $c(i \xrightarrow{\ell} j)$.

The tree automata are factored the same way, in that they have one rule per graph constant and per dependency edge. As a result, we get a one-to-one correspondence between parser scores and automaton rules when aligning automata rules to words via the words’ alignments to graph nodes.

We thus take the neural parser scores as *rule weights* $c(r)$ for rules r in the automaton. In a weighted tree automaton, the weight of a tree is defined as the product of the weights of all rules that built it. The *inside score* I of the tree automaton is the sum of the weights of all the trees it accepts. Computing this sum naively would be intractable, but the inside score can be computed efficiently with dynamic programming. Our training objective is to maximize the sum of the log inside scores of all automata in the corpus.

The arithmetic structure of computing the inside scores is complex and varies from automaton to automaton, which would make batching difficult. We solve this with the chain rule as follows:

$$\begin{aligned} \nabla_{\theta} \log I &= \frac{1}{I} \nabla I = \frac{1}{I} \sum_{r \in A} \frac{\partial}{\partial c(r)} I \nabla_{\theta} c(r) \\ &= \frac{1}{I} \sum_{r \in A} \alpha(r) \nabla_{\theta} c(r), \end{aligned}$$

where θ are the parameters of the neural parser, which determine $c(r)$, and $\alpha(r)$ is the *outer weight* of the rule r (Eisner, 2016), i.e. the total weight of trees that use r divided by $c(r)$. The outer weight can be effectively computed with the inside-outside algorithm (Baker, 1979). This occurs outside of the gradient, so we do not need to backpropagate into it. Since the scores $c(r)$ are direct outputs of the neural parser, their gradients can be batched straightforwardly.

Method	DM	PAS	PSD	AMR
random trees	81.1	79.0	67.8	70.8
random weights	93.0	94.4	80.0	75.0
EM weights	93.8	94.3	81.7	75.2
joint neural model (§6)	94.5	94.8	82.7	76.5

Table 1: Baseline comparisons on the development sets (3 source names in all experiments).

7 Evaluation

7.1 Setup

We evaluate parsing accuracy on the graphbanks DM, PAS, and PSD from the SemEval 2015 shared task on Semantic Dependency Parsing (SDP, Oepen et al. (2015)) and on the AMRBank LDC2017T10 (Banarescu et al., 2013). We follow Lindemann et al. (2019) in the choice of neural architecture, in particular using BERT (Devlin et al., 2019) embeddings, and in the choice of decoder, hyperparameters and pre- and postprocessing (we train the model of §6 for 100 instead of 40 epochs, since it is slower to converge than supervised training). When a graph G is *non-decomposable* using our blob partition, i.e. if there is no well-typed AM dep-tree T that evaluates to G , and so the condition of Theorem 2 does not hold, then we remove that graph from the training set. (This does not affect coverage at evaluation time.) This occurs rarely, affecting e.g. about 1.6% of graphs in the PSD training set.

Like (Lindemann et al., 2019), we use the heuristic AMR alignments of (Groschwitz et al., 2018). These alignments can yield multi-node constants. In those cases, we first run the algorithm of Section 4 to obtain an AM tree with placeholder source names, and then consolidate those constants that are aligned to the same word into one constant, effectively collapsing segments of the AM tree into a single constant. We then construct the tree automata of Section 5 as normal.

7.2 Results

We consider three baselines. Each of these chooses a single tree for each training instance from the tree automata and performs supervised training. The **random trees** baseline samples a tree for each sentence from its automaton, uniformly at random. In the **random weights** baseline, we fix a random weight for each graph constant and edge label, globally across the corpus, and select the highest-scoring tree for each sentence. The **EM weights** baseline instead optimizes these global weights with the inside-outside algorithm.

	DM		PAS		id F	PSD		AMR 17 Smatch F
	id F	ood F	id F	ood F		ood F		
He and Choi (2020)	94.6	90.8	96.1	94.4	86.8	79.5	-	
FG'20	94.4	91.0	95.1	93.4	82.6	82.0	-	
Bevilacqua et al. (2021)	-	-	-	-	-	-	84.5	
L'19, w/o MTL	93.9±0.1	90.3±0.1	94.5±0.1	92.5±0.1	82.0±0.1	81.5±0.3	76.3±0.2	
This work	94.2±0.0	90.2±0.1	94.6±0.0	92.7±0.1	81.4±0.1 (75.8±0.1)	80.7±0.4 (74.1±0.1)	75.1±0.2 (74.2±0.3)	

Table 2: Semantic parsing accuracies (id = in domain test set; ood = out of domain test set). Results for our work are averages of three runs with standard deviations. L'19 are results of Lindemann et al. (2019) with fixed tree decoder (incl. post-processing bugfix for AMR as per Lindemann et al. (2020)). FG'20 is Fernández-González and Gómez-Rodríguez (2020).

Table 1 compares the baselines and the joint neural method. Random trees perform worst – consistency across the corpus matters. The difference between random weights and EM is surprisingly small, despite the EM algorithm converging well. The joint neural learning outperforms the baselines on all graphbanks; we analyze this in § 8. We also experimented with different numbers of sources, finding 3 to work best for DM, PAS and AMR, and 4 for PSD (all results in Appendix C).

Table 2 compares the accuracy of our joint model to Lindemann et al. (2019) and to the state of the art on the respective graphbanks. Our model is competitive with the state of the art on most graphbanks. In particular, our parsing accuracy is on par with Lindemann et al. (2019), who perform supervised training with hand-crafted heuristics. This indicates that our model learns appropriate source names.

Graphbank-specific pre- and processing. The pre- and postprocessing steps of (Lindemann et al., 2019) we use still rely on two graphbank-specific heuristics, that directly relate to AM dependency trees: in PSD, it includes a simple but effective step to make coordination structures more compatible with the specific flavor of application and modification of AM dependency trees. In AMR it includes a step to remove some edges related to coreference (a non-compositional source of reentrancy).

We include in brackets the results *without* those two preprocessing steps. The drop in performance for PSD indicates that while for the most part our method is graphbank-independent, not all shapes of graphs are equally suited for AM dependency-parsing and some preprocessing to bring the graph ‘into shape’ can still be important. For AMR, keeping the co-reference based edges leads to AM trees that resolve those reentrancies with the AM type system. That is, the algorithm ‘invents’ ad-hoc compositional explanations for a

non-compositional phenomenon, yielding graph constants with type annotations that do not generalize well. The corresponding drop in performance indicates that extending AM dependency parsing to handle coreference will be an important future step when parsing AMR; some work in that direction has already been undertaken (Anikina et al., 2020).

8 Linguistic Analysis

As AM parsing is inherently interpretable, we can explore linguistic properties of the learned graph constants and trees. We find that the neural method makes use of both syntax and semantics.

We compute for each sentence in the training set the best tree from its tree automaton, according to the neural weights of the best performing epoch. We then sample trees from this set for hand-analysis (see Appendix A), to examine whether the model learned consistent sources for subjects and objects. We find that while the EM method uses highly consistent graph constants and AM operations, the neural method, which has access to the strings, sacrifices some graph constant and operation consistency in favour of *syntactic* consistency.

Syntactic Subjects and Objects. In the active sentence *The fairy charms the elf*, the phrase *the fairy* is the syntactic subject and *the elf* the syntactic object. In the passive *The elf is charmed (by the fairy)*, the phrase *the elf* is now the syntactic subject, even though in both sentences, the fairy is the charmer and the elf the charmee. Similarly, *the fairy* is the syntactic subject in the intransitive sentence *The fairy glows*.

Intra-Phenomenon Consistency. For both the EM and neural method, we found completely consistent source allocations for active transitive verbs in all four sembanks. These source allocations were also the overwhelming favourite graph constants for two-argument predicates (72-92%), and

the most common sources used by Apply operations (94-98%). For example, in AMR, the graph constant template in Fig. 8a appears 26,653 times in the neural parser output. 74% of these used sources $x = S_1$ and $y = S_2$ (from $\mathcal{S} = \{S_1, S_2, S_3\}$). All active transitive sentences in our sample used this source allocation, so we call this the *active* graph constant (e.g. $G\text{-charm}$ in Fig. 2) and refer to the sources S_1 and S_2 as S and O respectively, for *subject* and *object*. All four sembanks showed this kind of consistency; when we refer to S and O sources below, we mean whichever two sources displayed the same behaviour as S_1 and S_2 in AMR.

All four graphbanks are also highly consistent in their modifiers: classical modifiers such as adjectives are nearly universally adjoined with one consistent source – we refer to it as M – and MOD_M is the overwhelming favourite (90-99%) for MOD operations.

Cross-Phenomenon Consistency. We call a parser *syntactically consistent* if its syntactic subjects fill the S slot, regardless of their semantic role. A syntactically consistent parser would acquire the AMR in Fig. 8c from the active sentence by the analysis in Fig. 8b, and from the passive sentence by the analysis in Fig. 8d, with the passive constant $G\text{-charmP}$ from Fig. 2.

The neural parser is syntactically consistent: in all sembanks, it uses the same source S for syntactic subjects in passives as for actives. EM, conversely, prefers to use the same graph constants for active and passives, flipping the APP edges to produce syntactically inconsistent trees as in Fig. 8e. Single-argument predicates are also syntactically consistent in the neural model, using S for subjects and O for objects, while EM picks one source. The heuristics in Lindemann et al. (2019) have passive constants, but use them only when forced to, e.g. when coordinating active and passive.

Finally, we compute the entropy of the graph constants for the best trees of the training set as $\sum_G f(g) \ln f(G)$, where $f(G)$ is the frequency of constant G in the trees. The entropies are between 2 and 3 nats, but are consistently lower for EM than the neural method, by 0.031 to 0.079 nats. Considering that the neural method achieves higher parsing accuracies, using the most common graph constants and edges possible evidently is not always optimal for performance. The syntactic regularities exploited by the neural method may contribute to its improved performance.

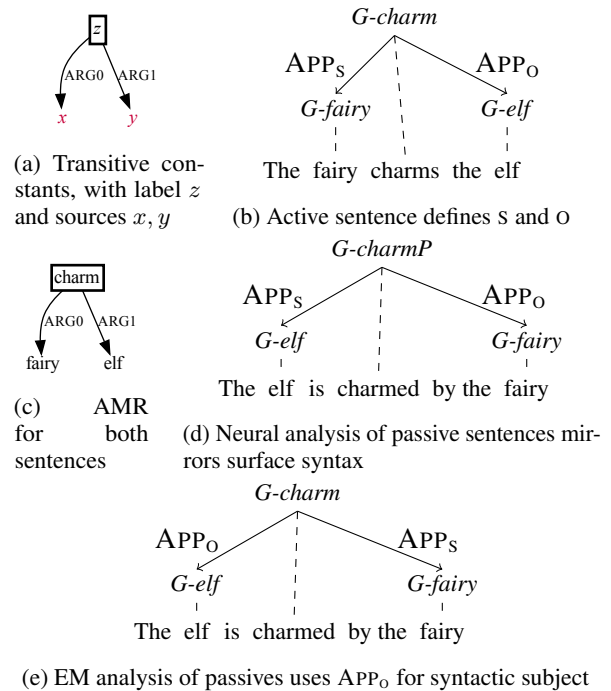


Figure 8: AMR examples of active and passive. See Fig. 2 for graph constants.

9 Conclusion

In this work, we presented a method to obtain the compositional structures for AM dependency parsing that relies much less on graphbank-specific heuristics written by experts. Our neural model learns linguistically meaningful argument slot names, as shown by our manual evaluation; in this regard, our model learns to do the job of the linguist. High parsing performance across graphbanks shows that the learned compositional structures are also well-suited for practical applications, promising easier adaptation of AM dependency parsing to new graphbanks.

Acknowledgments

We would like to thank the anonymous reviewers as well as Lucia Donatelli, Pia Weißenhorn and Matthias Lindemann for their thoughtful comments. This research was in part funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation), project KO 2916/2-2, and by the Dutch Research Council (NWO) as part of the project *Learning Meaning from Structure* (VI.Veni.194.057).

References

- Tatiana Anikina, Alexander Koller, and Michael Roth. 2020. [Predicting coreference in Abstract Meaning Representations](#). In *Proceedings of the Third Workshop on Computational Models of Reference, Anaphora and Coreference*, pages 33–38, Barcelona, Spain (online). Association for Computational Linguistics.
- Yoav Artzi, Kenton Lee, and Luke Zettlemoyer. 2015. [Broad-coverage CCG Semantic Parsing with AMR](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*.
- J. K. Baker. 1979. [Trainable grammars for speech recognition](#). *The Journal of the Acoustical Society of America*, 65(S1):S132–S132.
- 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*.
- Michele Bevilacqua, Rexhina Blloshmi, and Roberto Navigli. 2021. [One SPRING to rule them both: Symmetric AMR semantic parsing and generation without a complex pipeline](#). In *Proceedings of AAAI*.
- 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.
- Hubert Comon, Max Dauchet, Rémi Gilleron, Florent Jacquemard, Denis Lugiez, Sophie Tison, Marc Tommasi, and Christof Löding. 2007. *Tree Automata techniques and applications*. published online - <http://tata.gforge.inria.fr/>.
- Ann Copestake and Dan Flickinger. 2000. [An open-source grammar development environment and broad-coverage english grammar using HPSG](#). In *Proceedings of the Second conference on Language Resources and Evaluation (LREC)*.
- Ann Copestake, Alex Lascarides, and Dan Flickinger. 2001. [An algebra for semantic construction in constraint-based grammars](#). In *Proceedings of the 39th ACL*.
- Bruno Courcelle and Joost Engelfriet. 2012. *Graph Structure and Monadic Second-Order Logic, a Language Theoretic Approach*. Cambridge University Press.
- 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, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jason Eisner. 2016. [Inside-outside and forward-backward algorithms are just backprop \(tutorial paper\)](#). In *Proceedings of the Workshop on Structured Prediction for NLP*, pages 1–17.
- Federico Fancellu, Sorcha Gilroy, Adam Lopez, and Mirella Lapata. 2019. [Semantic graph parsing with recurrent neural network DAG grammars](#). 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 2769–2778, Hong Kong, China. Association for Computational Linguistics.
- Daniel Fernández-González and Carlos Gómez-Rodríguez. 2020. [Transition-based semantic dependency parsing with pointer networks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7035–7046, Online. Association for Computational Linguistics.
- 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.
- 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 *IWCS 2017 - 12th International Conference on Computational Semantics - Long papers*.
- Jonas Groschwitz, Matthias Lindemann, Meaghan Fowlie, Mark Johnson, and Alexander Koller. 2018. [AMR Dependency Parsing with a Typed Semantic Algebra](#). In *Proceedings of ACL*.
- Serhii Havrylov, Germán Kruszewski, and Armand Joulin. 2019. [Cooperative learning of disjoint syntax and semantics](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1118–1128, Minneapolis, Minnesota. Association for Computational Linguistics.
- Han He and Jinho Choi. 2020. [Establishing strong baselines for the new decade: Sequence tagging, syntactic and semantic parsing with BERT](#). In *The Thirty-Third International Flairs Conference*.

- Jonathan Herzig and Jonathan Berant. 2020. [Span-based semantic parsing for compositional generalization](#).
- Nikita Kitaev and Dan Klein. 2018. [Constituency parsing with a self-attentive encoder](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2676–2686, Melbourne, Australia. Association for Computational Linguistics.
- Matthias Lindemann, Jonas Groschwitz, and Alexander Koller. 2019. [Compositional semantic parsing across graphbanks](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4576–4585, Florence, Italy. Association for Computational Linguistics.
- Matthias Lindemann, Jonas Groschwitz, and Alexander Koller. 2020. [Fast semantic parsing with well-typedness guarantees](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3929–3951, Online. Association for Computational Linguistics.
- Jean Maillard, Stephen Clark, and Dani Yogatama. 2019. [Jointly learning sentence embeddings and syntax with unsupervised tree-LSTMs](#). *Natural Language Engineering*, 25(4).
- Stephan Open, 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)*.
- Xiaochang Peng, Linfeng Song, and Daniel Gildea. 2015. [A synchronous hyperedge replacement grammar based approach for AMR parsing](#). In *Proceedings of the 19th Conference on Computational Language Learning*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, undefinukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 6000–6010, Red Hook, NY, USA. Curran Associates Inc.

Learning compositional structures for semantic graph parsing

Supplementary Materials

Jonas Groschwitz

Saarland University

jonasg@coli.uni-saarland.de

Meaghan Fowlie

Utrecht University

m.fowlie@uu.nl

Alexander Koller

Saarland University

koller@coli.uni-saarland.de

A Sampling Method for hand analysis

To sample trees, we compute for each sentence in the training set the best tree from its tree automaton, according to the neural weights of the best performing epoch. This ensures the AM trees evaluate to the correct graph. We then sample trees from this set for hand-analysis.

To get relevant sentences, we sampled 5-to-15-word sentences with graph constants from the following six categories:

Transitive verbs: graph constants with a labeled root and two arguments with edges labelled as in Table 1:

Sembank	subject	object
AMR	ARG0	ARG1
DM	ARG1	ARG2
PAS	verb_ARG1	verb_ARG2
PSD	ACT_arg	PAT_arg

Table 1: Transitive verbs

As explained in the main text, we define the *active* constants as those with the most common source allocation, and the *passive* constants as those with the active source allocation flipped. We sampled both active and passive source allocations.

Verbs with one argument: Graph constants just like the transitive ones but lacking one of the arguments. There are four of these, given both source allocations.

Generally these graph constants are used for more than just verbs; for each of the six categories we sampled until we had ten relevant sentences. We visualised the AM trees and categorised the phenomena, for example active or passive verbs, nominalised verbs, imperatives, relative clauses, gerund modifiers, and so forth.

To answer the question of whether the parser used consistent constants for active and passive transitive sentences, we sampled until we had ten sentences with active or passive main verbs. For the single-argument verbs, we also looked at nominalised verbs, modifiers, and so forth. (Sampling and visualisation scripts will be available together with the rest of our code on GitHub.)

B An algorithm to obtain all AM dep-trees for a graph

Let G be a graph partitioned into blobs. Let \mathcal{U}_G be the set of unrolled graphs for G that can be obtained by Algorithm 2 by varying the queue order.

Let further \mathcal{M}_G be the set of results of Algorithm 3 below for every input AM dep-tree $T = C_U$ for $U \in \mathcal{U}_G$ and every choice of set M as specified in the algorithm. Algorithm 3 switches the order of two nodes m and k , making k the head of the subtree previously headed by m . This change of head is only possible when the incoming edge of m is labeled MOD (for APP, the change of head changes the evaluation result). It also requires a MOD edge between m and k ; an APP edge with this type of swap would lead to a non-well-typed graph.

Finally, let \mathcal{R}_G be the set of results of Algorithm 4 for every input AM dep-tree $T \in \mathcal{M}_G$ and any valid choice of R and RT (valid as described in the algorithm). Algorithm 4 is like Algorithm 1 for reentrancy resolution, but can have resolution targets RT(n) that are higher in the tree than the lowest common ancestor of n and the REF- n nodes. Further, Algorithm 4 uses the same methodology to also move nodes that do not need resolution to become descendants of a ‘resolution target’ higher in the tree (i.e. R here can now also contain nodes for which no REF node exists).

Then the following Theorem 1 holds:

Algorithm 3: Modify-edge swapping

1 Input: an AM dep-tree T and a set M of pairs of consecutive edges in T of the form $\langle n \xrightarrow{\text{MOD}_n} m, m \xrightarrow{\text{MOD}_m} k \rangle$ such that no edge appears in multiple pairs.

2 for $\langle n \xrightarrow{\text{MOD}_n} m, m \xrightarrow{\text{MOD}_m} k \rangle \in M$:

3 Replace $n \xrightarrow{\text{MOD}_n} m$ in T with $n \xrightarrow{\text{MOD}_n} k$;

4 Replace $m \xrightarrow{\text{MOD}_m} k$ in T with $k \xrightarrow{\text{APP}_m} m$;

5 Add $\beta(m)$ (which always includes n) to the request at m in $\tau(k)$;

6 return T

Algorithm 4: Extended reentrancy resolution

1 Input: an AM dep-tree T ; a set $R \supseteq \{n \in N_G \mid \exists \text{ REF-}n \text{ node in } T\}$; and a map RT that assigns to each node $n \in R$ a resolution target $\text{RT}(n)$, that is at least as high as the lowest common ancestor of n and all REF- n nodes (if they exist), and that satisfies the conditions of Theorem 1.

2 while $R \neq \emptyset$:

3 Pick a $y \in R$ s.t. there is no $x \in R$, $x \neq y$, with y on an x -resolution path;

4 for $p \in y$ -resolution paths:

5 for $n \xrightarrow{\text{APP}} m \in p$:

6 if m is y or labeled REF- y :

7 Add $\beta(y)$ to the request at y in $\tau(n)$;

8 else:

9 Add $y[\beta(y)]$ to the request at m in $\tau(n)$;

10 Move the subtree of T rooted at y up to be an APP $_y$ daughter of $\text{RT}(y)$, unless $\text{RT}(y) = y$;

11 Delete all REF- y nodes from T ;

12 $R \leftarrow R - \{y\}$

13 return T

Theorem 1. Let G be a graph partitioned into blobs, and let \mathcal{T}_G be the set of all well-typed AM dep-trees with placeholder sources, using that blob partition, that evaluate to G . Then if $\mathcal{T}_G = \emptyset$, all AM dep-trees in \mathcal{R}_G are either not well-typed or do not evaluate to G . If however $\mathcal{T}_G \neq \emptyset$, then $\mathcal{R}_G = \mathcal{T}_G$.

# sources	DM	PAS	PSD	AMR
2	92.2	91.9	75.6	74.3
3	94.5	94.8	82.7	76.5
4	94.4	94.7	83.4	75.9
6	92.3	93.6	80.1	73.4

Table 2: Development set accuracies of the neural method for different numbers of source names.

C Additional Details

- AMR F-scores are *Smatch scores* (Cai and Knight, 2013)
- DM, PAS and PSD: We compute labeled F-score with the evaluation toolkit that was developed for the SDP shared task: <https://github.com/semantic-dependency-parsing/toolkit>
- We use the standard train/dev/test split for all corpora
- AMR corpus available through <https://amr.isi.edu/download.html> (requires LDC license)
- SDP corpora available through <https://catalog.ldc.upenn.edu/LDC2016T10> (requires LDC license)

Number of source names. We experimented with different numbers of source names in the joint neural method (Table 2). Mostly, three source names were most effective, except for PSD, where four were most effective. Two source names are not enough to model many common phenomena (e.g. ditransitive verbs, coordination of verbs); graphs containing these phenomena cannot be decomposed with two sources and are removed from the training set, reducing parsing accuracy. The higher performance of PSD with four sources may stem from PSD using flat coordination structures which require more source names; although this is also true for AMR where four source names are not beneficial. The drop with six source names may come from the fact that the latent space grows rapidly with more sources, making it harder to learn consistent source assignments.

Hyperparameters. See Table 3.

Activation function	tanh
Optimizer	Adam
Learning rate	0.001
Epochs	100
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 3: Common hyperparameters used in all experiments (the random trees, random weights and EM weights baselines use 40 epochs since they converge faster). For a complete description of the neural architecture, see [Lindemann et al. \(2019\)](#) and its supplementary materials.

References

- 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.
- Matthias Lindemann, Jonas Groschwitz, and Alexander Koller. 2019. [Compositional semantic parsing across graphbanks](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4576–4585, Florence, Italy. Association for Computational Linguistics.