

# Between a Rock and a Hard Place – Uniform Parsing for Hyperedge Replacement DAG Grammars

Henrik Björklund, Frank Drewes, and Petter Ericson<sup>(✉)</sup>

Department of Computing Science, Umeå University, Umeå, Sweden  
{henrikb,drewes,pettter}@cs.umu.se

**Abstract.** Motivated by applications in natural language processing, we study the uniform membership problem for hyperedge-replacement grammars that generate directed acyclic graphs. Our major result is a low-degree polynomial-time algorithm that solves the uniform membership problem for a restricted type of such grammars. We motivate the necessity of the restrictions by two different NP-completeness results.

**Keywords:** Graph grammar · Hyperedge replacement · Abstract meaning representation · DAG grammar · Uniform membership problem · Parsing

## 1 Introduction

Hyperedge-replacement grammars (HRGs [5, 7]) are one of the most successful formal models for the generation of graph languages, because their properties resemble those of context-free grammars to a great extent. Unfortunately, polynomial parsing is an exception from this rule: graph languages generated by HRGs may be NP-complete. Thus, not only is the uniform membership problem intractable (unless  $P \neq NP$ ), but the non-uniform one is as well [1, 8].

Recently, Chiang et al. [4] advocated the use of hyperedge-replacement for describing meaning representations in natural language processing (NLP), and in particular the abstract meaning representations (AMRs) proposed by Banarescu et al. [2], and described a general recognition algorithm together with a detailed complexity analysis. Unsurprisingly, the running time of the algorithm is exponential even in the non-uniform case, one of the exponents being the maximum degree of nodes in the input graph. Unfortunately, this is one of the parameters one would ideally not wish to limit, since AMRs may have unbounded node degree. However, AMRs are directed acyclic graphs (DAGs), a fact that is not exploited in [4]. Another recent approach to HRG parsing is [6], where predictive top-down parsing in the style of SLL(1) parsers is proposed. This is a uniform approach yielding parsers of quadratic running time in the size of the input graph, but the generation of the parser from the grammar is not guaranteed to run in polynomial time. (For a list of earlier attempts to HRG parsing, see [6].)

In this paper, we study the complexity of the membership problem for DAG-generating HRGs. Since NLP applications usually involve a machine learning component in which the rules of a grammar are inferred from a corpus, and hence the resulting HRG cannot be assumed to be given beforehand, we are mainly interested in efficient algorithms for the uniform membership problem. We propose restricted DAG-generating HRGs and show, in Sect. 4, that their uniform membership problem is solvable in time  $\mathcal{O}(n^2 + nm)$ , where  $m$  and  $n$  are the sizes of the grammar and the input graph, resp. In linguistic applications, where grammars are usually much larger than the structures to be parsed, this is essentially equivalent to  $\mathcal{O}(nm)$ . To our knowledge, this is the first uniform polynomial-time parsing algorithm for a non-trivial subclass of HRGs. Naturally, the restrictions are rather strong, but we shall briefly argue in Sect. 5 that they are reasonable in the context of AMRs. We furthermore motivate the restrictions with two NP-completeness results for DAG-generating HRGs, in Sect. 6.

To save space, most proofs have been omitted. They are available in [3].

## 2 Preliminaries

The set of non-negative integers is denoted by  $\mathbb{N}$ . For  $n \in \mathbb{N}$ ,  $[n]$  denotes  $\{1, \dots, n\}$ . Given a set  $S$ , let  $S^{\circledast}$  be the set of non-repeating lists of elements of  $S$ . If  $sw \in S^{\circledast}$  with  $s \in S$ , we shall also denote  $sw$  by  $(s, w)$ . If  $\preceq$  is a (partial) ordering of  $S$ , we say that  $s_1 \cdots s_k \in S^{\circledast}$  *respects*  $\preceq$  if  $s_i \preceq s_j$  implies  $i \leq j$ .

**Hypergraphs and DAGs.** A *ranked alphabet* is a pair  $(\Sigma, \text{rank})$  consisting of a finite set  $\Sigma$  of symbols and a *ranking function*  $\text{rank} : \Sigma \rightarrow \mathbb{N}$  which assigns a *rank*  $\text{rank}(a)$  to every symbol  $a \in \Sigma$ . We usually identify  $(\Sigma, \text{rank})$  with  $\Sigma$  and keep ‘rank’ implicit.

Let  $\Sigma$  be a ranked alphabet. A (directed hyperedge-labeled) *hypergraph* over  $\Sigma$  is a tuple  $G = (V, E, \text{src}, \text{tar}, \text{lab})$  consisting of

- finite sets  $V$  and  $E \subseteq V \times V^{\circledast}$  of *nodes* and *hyperedges*, respectively
- *source* and *target mappings*  $\text{src} : E \rightarrow V$  and  $\text{tar} : E \rightarrow V^{\circledast}$  assigning to each hyperedge  $e$  its source  $\text{src}(e)$  and its sequence  $\text{tar}(e)$  of targets, and
- a *labeling*  $\text{lab} : E \rightarrow \Sigma$  such that  $\text{rank}(\text{lab}(e)) = |\text{tar}(e)|$  for every  $e \in E$ .

Below, we call hyperedges edges and hypergraphs graphs, for simplicity. Note that edges have only one source but several targets, similarly to the usual notion of term (hyper)graphs. The DAGs we shall consider below are, however, more general than term graphs in that nodes can have out-degree larger than one.

Continuing the formal definitions, a *path* in  $G$  is a (possibly empty) sequence  $e_1, e_2, \dots, e_k$  of edges such that for each  $i \in [k - 1]$  the source of  $e_{i+1}$  is a target of  $e_i$ . The *length* of a path is the number of edges it contains. A nonempty path is a *cycle* if the source of the first edge is a target of the last edge. If  $G$  does not contain any cycle then it is *acyclic* and is called a *DAG*. The *height* of a DAG  $G$  is the maximum length of any path in  $G$ . A node  $v$  is a *descendant* of a node  $u$  if  $u = v$  or there is a nonempty path  $e_1, \dots, e_k$  in  $G$  such that  $u = \text{src}(e_1)$  and  $v$  occurs in  $\text{tar}(e_k)$ . An edge  $e'$  is a *descendant edge* of an edge  $e$  if there is a path  $e_1, \dots, e_k$  in  $G$  such that  $e_1 = e$  and  $e_k = e'$ .

The *in-degree* of a node  $u \in V$  is the number of edges  $e$  such that  $u$  is a target of  $e$ . The *out-degree* of  $u$  is the number of edges  $e$  such that  $u$  is the source of  $e$ . A node with in-degree 0 is a *root* and a node with out-degree 0 is a *leaf*.

For a node  $u$  of a DAG  $G = (V, E, \text{src}, \text{tar}, \text{lab})$ , the *sub-DAG rooted at  $u$*  is the DAG  $G \downarrow_u$  induced by the descendants of  $u$ . Thus  $G \downarrow_u = (U, E', \text{src}', \text{tar}', \text{lab}')$  where  $U$  is the set of all descendants of  $u$ ,  $E' = \{e \in E \mid \text{src}(e) \in U\}$ , and  $\text{src}'$ ,  $\text{tar}'$ , and  $\text{lab}'$  are the restrictions of  $\text{src}$ ,  $\text{tar}$  and  $\text{lab}$  to  $E'$ . A leaf  $v$  of  $G \downarrow_u$  is *reentrant* if there exists an edge  $e \in E \setminus E'$  such that  $v$  occurs in  $\text{tar}(e)$ .

**DAG Grammars.** A *marked* graph is a tuple  $G = (V, E, \text{src}, \text{tar}, \text{lab}, X)$  where  $(V, E, \text{src}, \text{tar}, \text{lab})$  is a graph and  $X \in V^{\otimes}$  is nonempty. The sequence  $X$  is called the *marking* of  $G$ , and the nodes in  $X$  are referred to as *external nodes*. If  $X = (v, w)$  for some  $v \in V$  and  $w \in V^{\otimes}$  then we denote them by  $\text{root}(G)$  and  $\text{ext}(G)$ , resp. This is motivated by the form of our rules, which is defined next.

**Definition 1 (DAG grammar).** A DAG grammar is a system  $H = (\Sigma, N, S, P)$  where  $\Sigma$  and  $N$  are disjoint ranked alphabets of terminals and nonterminals, respectively,  $S$  is the starting nonterminal with  $\text{rank}(S) = 0$ , and  $P$  is a set of productions. Each production is of the form  $A \rightarrow F$  where  $A \in N$  and  $F$  is a marked DAG over  $\Sigma \cup N$  with  $|\text{ext}(F)| = \text{rank}(A)$  such that  $\text{root}(F)$  is the unique root of  $F$  and  $\text{ext}(F)$  contains only leaves of  $F$ .

Naturally, a terminal (nonterminal) edge is an edge labeled by a terminal (nonterminal, resp.). We may sometimes just call them terminals and nonterminals if there is no danger of confusion. By convention, we use capital letters to denote nonterminals, and lowercase letters for terminal symbols.

A derivation step of  $H$  is described as follows. Let  $G$  be a graph with an edge  $e$  such that  $\text{lab}(e) = A$  and let  $A \rightarrow F$  in  $P$  be a rule. Applying the rule involves replacing  $e$  with an unmarked copy of  $F$  in such a way that  $\text{src}(e)$  is identified with  $\text{root}(F)$  and for each  $i \in [|\text{tar}(e)|]$ , the  $i$ th node in  $\text{tar}(e)$  is identified with the  $i$ th node in  $\text{ext}(F)$ . Notice that  $|\text{tar}(e)| = |\text{ext}(F)|$  by definition. If the resulting graph is  $G'$ , we write  $G \Rightarrow_H G'$ . We write  $G \Rightarrow_H^* G'$  if  $G'$  can be derived from  $G$  in zero or more derivation steps. The *language*  $\mathcal{L}(H)$  of  $H$  are all graphs  $G$  over the terminal alphabet  $T$  such that  $S^\bullet \Rightarrow_H^* G$  where  $S^\bullet$  is the graph consisting of a single node and a single edge labeled by  $S$ .

The graphs produced by DAG grammars are connected, single-rooted, and as the name implies, acyclic. This can be proved in a straightforward manner by induction on the length of the derivation. In the following, we consider only graphs of height at least 1, as the (single) graph of height 0 requires simple but cumbersome special cases.

**Ordering the Leaves of a DAG.** Let  $G = (V, E, \text{src}, \text{tar}, \text{lab})$  be a DAG and let  $u$  and  $u'$  be leaves of  $G$ . We say that an edge  $e$  with  $\text{tar}(e) = w$  is a *common ancestor edge* of  $u$  and  $u'$  if there are  $t$  and  $t'$  in  $w$  such that  $u$  is a descendant of  $t$  and  $u'$  is a descendant of  $t'$ . If, in addition, there is no edge with its source in  $w$  that is a common ancestor edge of  $u$  and  $u'$ , we say that  $e$  is a *closest common ancestor edge* of  $u$  and  $u'$ .

Note that in a DAG, a pair of nodes can have more than one closest common ancestor edge.

**Definition 2.** Let  $G = (V, E, \text{src}, \text{tar}, \text{lab})$  be a DAG. Then  $\preceq_G$  is the partial order on the leaves of  $G$  defined by  $u \preceq_G u'$  if, for every closest common ancestor edge  $e$  of  $u$  and  $u'$ ,  $\text{tar}(e)$  can be written as  $wtw'$  such that  $t$  is an ancestor of  $u$  and all ancestors of  $u'$  in  $\text{tar}(e)$  are in  $w'$ .

### 3 Restricted DAG Grammars

DAG grammars are a special case of hyperedge-replacement grammars. We now define further restrictions that will allow polynomial time uniform parsing. Every rule  $A \rightarrow F$  of a *restricted DAG grammar* is required to satisfy the following conditions (in addition to the conditions formulated in Definition 1):

1. If a node  $v$  of  $F$  has in-degree larger than one, then  $v$  is a leaf
2. If  $F$  consists of exactly two edges  $e_1$  and  $e_2$ , both labeled by  $A$ , such that  $\text{src}(e_1) = \text{src}(e_2)$  and  $\text{tar}(e_1) = \text{tar}(e_2)$  we call  $A \rightarrow F$  a *clone rule*. Clone rules are the only rules in which a node can have out-degree larger than 1 and the only rules in which a nonterminal can have the root as its source.
3. For every nonterminal  $e$  in  $F$ , all nodes in  $\text{tar}(e)$  are leaves.
4. If a leaf of  $F$  has in-degree exactly one, then it is an external node or its unique incoming edge is terminal.
5. The leaves of  $F$  are totally ordered by  $\preceq_F$  and  $\text{ext}(F)$  respects  $\preceq_F$ .

We now demonstrate some properties of restricted DAG grammars.

**Lemma 3.** Let  $H = (\Sigma, N, S, P)$  be a restricted DAG grammar,  $G$  a DAG such that  $S^\bullet \Rightarrow_H^* G$ , and  $U$  the set of nodes of in-degree larger than 1 in  $G$ . Then  $U$  contains only leaves of  $G$  and  $\text{tar}(e) \in U^\otimes$  for every nonterminal  $e$  of  $G$ .


Note that the lemma implies that leaves with in-degree exactly one are only connected to terminal edges. The lemma is proven by induction on the length of derivations, starting with the observation that  $S^\bullet$  has the properties claimed. To simplify the presentation of our algorithm, we introduce a normal form.

**Definition 4.** A restricted DAG grammar  $H = (\Sigma, N, S, P)$  is on normal form if every rule  $A \rightarrow F$  in  $P$  has one of the following three forms.

- (a) The rule is a clone rule.
- (b)  $F$  has a single edge  $e$ , which is terminal.
- (c)  $F$  has height 2, the unique edge  $e$  with  $\text{src}(e) = \text{root}(F)$  is terminal, and all other edges are nonterminal.

See Fig. 1 for examples of right-hand sides of the three types. In particular, right-hand sides  $F$  of the third type consist of nodes  $v, v_1, \dots, v_m, u_1, \dots, u_n$ , a terminal edge  $e$  and nonterminal edges  $e_1, \dots, e_k$  such that

- $v = \text{root}(F) = \text{src}(e)$  and  $v_1 \cdots v_m$  is a subsequence of  $\text{tar}(e)$ ,
- $\text{src}(e_i) \in \{v_1, \dots, v_m\}$  for all  $i \in [k]$ ,
- $\text{ext}(F)$  and  $\text{tar}(e_i)$ , for  $i \in [k]$ , are subsequences of  $u_1 \cdots u_n$ .



**Fig. 1.** Examples right-hand sides  $F$  of normal form rules of types (a), (b), and (c) for a nonterminal of rank 3. In illustrations such as these, boxes represent hyperedges  $e$ , where  $\text{src}(e)$  is indicated by a line and the nodes in  $\text{tar}(e)$  by arrows. Filled nodes represent the marking of  $F$ . Both  $\text{tar}(e)$  and  $\text{ext}(F)$  are drawn from left to right unless otherwise indicated by numbers

The proof of the following lemma follows the standard technique of dividing rules with large right-hand sides into several rules with smaller right-hand sides as in the proof of the Chomsky normal form of context-free grammars. The total size of the grammar does not change (except for a small constant factor).

**Lemma 5.** *Every restricted DAG grammar  $H$  can be transformed in linear time into a restricted DAG grammar  $H'$  on normal form such that  $\mathcal{L}(H) = \mathcal{L}(H')$ .*

One can now show that restrictions 1–5 imply that, in a DAG  $G$  generated by a restricted DAG grammar, the orders  $\preceq_{G \downarrow v}$  are consistent for all nodes  $v$ , that is, we have the following lemma the proof of which can be found in [3] (like the other proofs left out in this short version).

**Lemma 6.** *Let  $H$  be a restricted DAG grammar and  $G = (V, E, \text{src}, \text{tar}, \text{lab})$  a DAG generated by  $H$ . Then there is a total order  $\preceq$  on the leaves of  $G$  such that  $\preceq_G \subseteq \preceq$  and for every  $v \in V$  and every pair  $u, u'$  of reentrant nodes of  $G \downarrow_v$  we have  $u \preceq u' \Leftrightarrow u \preceq_{G \downarrow v} u'$ .*

If a DAG  $G$  has been derived by a restricted DAG grammar in normal form, it is uniquely determined which subgraphs of  $G$  have been produced by a nonterminal, and which leaves were connected to it at that point. In particular, given a non-leaf node  $v$  in  $G$ , consider the subgraph  $G \downarrow_v$ . Consider the earliest point in the derivation where there was a nonterminal  $e$  having  $v$  as its source. We say that  $e$  generated  $G \downarrow_v$ . From the structure of  $G$  and  $G \downarrow_v$ , we know that all reentrant nodes of  $G \downarrow_v$  are leaves and, by restriction 4, that  $e$  must have had exactly these reentrant leaves of  $G \downarrow_v$  as targets. By Lemma 6 and restriction 5, the order of these leaves in  $\text{tar}(e)$  coincides with the total order  $\preceq_{G \downarrow v}$ .

In other words, during the generation of  $G$  by a restricted DAG grammar,  $G \downarrow_v$  must be generated from a nonterminal  $e$  such that  $\text{src}(e) = v$  and  $\text{tar}(e)$  is uniquely determined by the condition that it consists of exactly the reentrant nodes of  $G \downarrow_v$  and respects  $\preceq_{G \downarrow v}$ . Therefore, we will from now on view  $G \downarrow_v$  as a *marked DAG*, where the marking is  $(v, \text{tar}(e))$ .

## 4 A Polynomial Time Algorithm

We present the parsing algorithm in pseudocode, after which we explain various subfunctions used therein. Intuitively, we work bottom-up on the graph in a manner resembling bottom-up finite-state tree automata, apart from where a node has out-degree greater than one. We assume that a total order  $\preceq$  on the leaves of the input DAG  $G$ , as ensured by Lemma 6, is computed in a preprocessing step before the algorithm is executed. At the same time, the sequence  $w_v$  of external nodes of each sub-DAG  $G\downarrow_v$  is computed. (Recall from the paragraph above that these are the reentrant leaves of  $G\downarrow_v$ , ordered according to  $\preceq_{G\downarrow_v}$ .) For a DAG  $G$  of size  $n$ , this can be done in time  $O(n^2)$  by a bottom-up process. To explain how, let us denote the set of all leaves of  $G\downarrow_v$  by  $U_v$  for every node  $v$  of  $G$ . We proceed as follows. For a leaf  $v$ , let  $\preceq_v = \{(v, v)\}$  and  $w_v = v$ . For every edge  $e$  with  $\text{tar}(e) = u_1 \dots u_k$  such that  $u_i$  has already been processed for all  $i \in [k]$ , first check if  $\preceq_0 = \bigcup_{i \in [k]} \preceq_{u_i}$  is a partial order. If so, define  $\preceq_e$  to be the unique extension of  $\preceq_0$  given as follows. Consider two nodes  $u, u' \in U_{\text{src}(e)}$  that are not ordered by  $\preceq_0$ . If  $i, j$  are the smallest indices such that  $u \in U_{u_i}$  and  $u' \in U_{u_j}$ , then  $u \preceq_e u'$  if  $i < j$ . Note that  $\preceq_e$  is uniquely determined and total. Moreover, let  $w_e$  be the unique sequence in  $U_{\text{src}(e)}^{\otimes}$  which respects  $\preceq_0$  and contains exactly the nodes in  $U_{\text{src}(e)}$  which are targets of edges of which  $e$  is not an ancestor edge. Similarly, if  $v$  is a node and all edges  $e_1, \dots, e_k$  having  $v$  as their source have already been processed, check if  $\bigcup_{i \in [k]} \preceq_{e_i}$  is a partial order. If so, define  $\preceq_e$  to be any total extension of this order. Moreover, check that  $w_{e_1} = \dots = w_{e_k}$ , and let  $w_v$  be exactly this sequence. The preprocessing may fail for some graphs, but as these may not be part of  $L(G)$  for any restricted DAG grammar  $G$ , we simply reject.

After this preprocessing, Algorithm 1 can be run. As the sequences  $w_u$  of external nodes for each sub-DAG  $G\downarrow_u$  were computed in the preprocessing step, we consider this information to be readily available in the pseudocode. This, together with the assumption that the DAG grammar  $H$  is in normal form allows for much simplification of the algorithm.

Walking through the algorithm step by step, we first extract the root node (line 2) and determine which kind of (sub-)graph we are dealing with (line 4): one with multiple outgoing edges from the root must have been produced by a cloning rule to be valid, meaning we can parse each constituent subgraph (line 5) recursively (line 6) and take the intersection of the resulting nonterminal edges (line 7). Each nonterminal that could have produced all the parsed subgraphs and has a cloning rule is entered into *returns* (line 8). The procedure `subgraphs_below` is used to partition the sub-DAG  $G\downarrow_v$  into one sub-DAG per edge having  $v$  as its source, by taking each such edge and all its descendant edges (and all their source and target nodes) as the subgraph. Note that the order among these subgraphs is undefined, though they are all guaranteed by the preprocessing to have the same sequence of external nodes  $w_v$ .

If, on the other hand, we have a single outgoing edge from the root node (line 9), we iterate through the subgraphs below the (unique) edge below the

**Algorithm 1.** Parsing of restricted graph grammars

---

```

1: function PARSES_TO(restricted DAG grammar  $H$  in normal form, DAG  $G$ )
2:    $v \leftarrow \text{root}(G)$ 
3:    $\text{returns} \leftarrow \emptyset$ 
4:   if  $\text{out\_degree}(v) > 1$  then
5:     for  $G_i \leftarrow \text{subgraphs\_below}(v)$  do
6:        $X_i \leftarrow \text{parses\_to}(G_i)$ 
7:        $X \leftarrow \bigcap_i X_i$ 
8:        $\text{returns} \leftarrow \{A \in X \mid \text{has\_clone\_rule}(A)\}$ 
9:   else
10:     $e \leftarrow \text{edge\_below}(v)$ 
11:     $\text{children} \leftarrow ()$ 
12:    for  $v' \leftarrow \text{targets}(e)$  do
13:      if  $\text{leaf}(v')$  then
14:         $\text{append}(\text{children}, \text{external\_node}(v'))$ 
15:      else
16:         $\text{append}(\text{children}, \text{parses\_to}(G \downarrow_{v'}))$ 
17:       $\text{returns} \leftarrow \{A \mid (A \rightarrow F) \in P \text{ and } \text{match}(F, e, \text{children})\}$ 
18:    return  $\text{returns}$ 

```

---

root node (line 12). Nodes are marked either with a set of nonterminals (that the subgraph below the nodes can parse to) (line 16), or, if the node is a leaf, with a Boolean indicating whether or not the node is reentrant in the currently processed subgraph  $G$  (line 14).

The `match` function used in line 17 deserves a closer description, as much of the complexity calculations depend on this function taking no more than time linear in the size of the right-hand side. Let  $\text{src}(e) = v$  and  $\text{tar}(e) = v_1 \cdots v_k$ . Each  $v_i$  has an entry in `emphchildren`. If  $v_i$  is a leaf it is a Boolean, otherwise a set of nonterminal labels. From  $G$  and `children`, we create a DAG  $G'$  as follows. Let  $T$  be the union of  $\{v, v_1, \dots, v_k\}$  and the set of leaves  $\ell$  of  $G$  such that  $\ell$  is reentrant to  $G$  (as indicated by `children`) or there is an  $i \in [k]$  with  $\ell$  being external in  $G \downarrow_{v_i}$ . Let  $T = \{v, v_1, \dots, v_k, t_1, \dots, t_p\}$ . Then  $G'$  has the set of nodes  $U = \{u, u_1, \dots, u_k, s_1, \dots, s_p\}$ . Let  $h$  be the bijective mapping with  $h(v) = u$  and  $h(v_i) = u_i$  for every  $i \in [k]$  and  $h(t_i) = (s_i)$  for every  $i \in [p]$ . We extend  $h$  to sequences in the obvious way. The root of  $G$  is  $u$  and there is a single edge  $d$  connected to it such that  $\text{lab}(d) = \text{lab}(e)$ ,  $\text{src}(d) = u$  and  $\text{tar}(d) = u_1 \cdots u_k$ . For every  $i \in [k]$  such that  $v_i$  is not a leaf,  $G'$  has an edge  $d_i$  with  $\text{src}(d_i) = u_i$  and  $\text{tar}(d_i) = h(w_i)$ , where  $w_i$  is the subsequence of leaves of  $G \downarrow_{v_i}$  that belong to  $T$ , ordered by  $\preceq$ . The edge is labeled by the *set* of nonterminals `children[i]`.

Once `match` has built  $G'$  it tests whether there is a way of selecting exactly one label for each nonterminal edge in  $G'$  such that the resulting graph is isomorphic to *rhs*. This can be done in linear time since the leaves of both  $G'$  and *rhs* are totally ordered and, furthermore, the ordering on  $v_1 \cdots v_k$  and  $u_1 \cdots u_k$  makes the matching unambiguous.

Let us now discuss the running time of Algorithm 1. Entering the `if` branch of `parses_to`, we simply recurse into each subgraph and continue parsing.

The actual computation in the `if`-clause is minor: an intersection of the  $l$  sets of nonterminals found. Each time we reach the `else` clause in `parses_to`, we consume one terminal edge of the input graph. We recurse once for each terminal edge below this (no backtracking), so the parsing itself enters the `else`-clause  $n$  times, where  $n$  is the number of terminal edges in the input graph. For each rule  $r = A \rightarrow F$ , we build and compare at most  $|F|$  nodes or edges in the `match` function. Thus, it takes  $\mathcal{O}(nm)$  operations to execute Algorithm 1 in order to parse a graph with  $n$  terminal hyperedges according to a restricted DAG grammar  $H$  in normal form of size  $m$ . If  $H$  is not in normal form, Lemma 5 can be used to normalize it in linear time. Since the process does not affect the size of  $H$  by more than a (small) linear factor, the time bound is not affected. Finally, a very generous estimation of the running time of the preprocessing stage yields a bound of  $\mathcal{O}(n^2)$ , because  $n$  edges (and at most as many nodes) have to be processed, each one taking no more than  $n$  steps. Altogether, we have shown the following theorem, the main result of this paper.

**Theorem 7.** *The uniform membership problem for restricted DAG grammars is solvable in time  $\mathcal{O}(n^2 + mn)$ , where  $n$  is the size of the input graph and  $m$  is the size of the grammar.*


## 5 Representing and Generating AMRs

An Abstract Meaning Representation (AMR) is an ordinary directed edge-labeled acyclic graph expressing the meaning of a sentence. An example expressing “*Anna’s cat is missing her*” is shown in Fig. 2. The root corresponds to the concept “missing”, which takes two arguments, the misser and the missed.

In this representation every node has a special “instance edge” that determines the concept represented by its source node (`miss`, `cat`, `anna`). The most important concepts are connected to (specific meanings of) verbs, which have a number of mandatory arguments `arg0`, `arg1` depending on the concept in question. While the representation shown is not directly compatible with the restrictions introduced in Sect. 3 a simple translation helps. Every concept with its  $k$  mandatory arguments is turned into a hyperedge of rank  $k + 1$ , the target nodes of which represent the instance (a leaf) and the roots of the arguments. The resulting hypergraph is shown in Fig. 2 on the right. Note that all shared nodes on the left (corresponding to cross-references) are turned into reentrant leaves. This is important because in a DAG generated by a restricted DAG grammar only leaves can have an in-degree greater than 1.

It might seem that we only need graphs with nodes of out-degree at most 1, and thus no cloning rules for their generation. However, a concept such as `miss` can typically also have optional so-called modifiers, such as in “*Anna’s cat is missing her heavily today*”, not illustrated in the figure. Such modifiers can typically occur in any number. We can add them to the structure by increasing the rank of `miss` by 1, thus providing the edge with another target  $v$ . The out-degree of this node  $v$  would be the number of modifiers of `miss`. Using the notation





**Fig. 2.** Example translation of AMR

of Sect. 4, each sub-DAG  $G \downarrow_e$  given by one of the outgoing edges  $e$  of  $v$  would represent one (perhaps complex) modifier. To generate these sub-DAGs  $G \downarrow_e$  a restricted DAG grammar would use a nonterminal edge that has  $v$  as its source and which can be cloned. The latter makes it possible to generate any number of modifiers all of which can refer to the same shared concepts.

## 6 NP-Hardness Results

In order to motivate the rather harsh restrictions we impose on our grammars, we present NP-hardness results for two different classes of grammars that are obtained by easing the restrictions in different ways.

**Theorem 8.** *The uniform membership problem for DAG grammars that conform to restrictions 1–4 is NP-complete.*

*Proof.* The problem is in NP since the restrictions guarantee that derivations are of linear length in the size of the input graph. It remains to prove NP-hardness.

Let us consider an instance  $\varphi$  of the satisfiability problem SAT, i.e., a set  $\{C_1, \dots, C_m\}$  of clauses  $C_i$ , each being a set of literals  $x_j$  or  $\neg x_j$ , where  $j \in [n]$  for some  $m, n \in \mathbb{N}$ . Recall that the question asked is whether there is an assignment of truth values to the variables  $x_j$  such that each clause contains a true literal. We have to show how to construct a DAG grammar  $H$  conforming to conditions 1–4 and an input graph  $G$  such that  $G \in L(H)$  if and only if  $\varphi$  is satisfiable.

We first give a construction that violates conditions 4 and 5. It uses nonterminals  $S, K, K_i, K_{ij}$  with  $i \in [m], j \in [n]$ . The terminal labels are  $c$ , all  $j \in [m]$ , and an “invisible” label. The labels  $K, K_i, K_{ij}, c$  are of rank  $2n$ ,  $S$  is of rank 0 and the remaining ones are of rank 1. Figure 3 depicts the rules of the grammar. Note that the rules are on normal form.

The first row of rules generates  $2n$  leaves which, intuitively, represent  $x_1, \neg x_1, \dots, x_n, \neg x_n$  and are targets of a  $K$ -labeled nonterminal. The nonterminal is




Fig. 3. Reduction of SAT to the uniform membership problem


cloned any number of times (with the intention to clone it  $m$  times, once for each clause). Afterwards it “guesses” which clause  $C_i$  ( $i \in \mathbb{N}$ ) it should check. The second row of rules lets every  $K_i$  “guess” which literal makes  $C_i$  true. If the literal is negative, it interchanges the corresponding targets, otherwise it keeps their order. The third row of rules, for all pairs  $(x_\ell, \neg x_\ell)$  that are not used to satisfy  $C_i$ , interchanges the corresponding targets or keeps their order. Finally, it replaces the nonterminal edge by a terminal one.

Now, consider the input DAG  $G$  in Fig. 4 (left). Suppose that  $G$  is indeed generated by  $H$ . Since the  $j$ th outgoing tentacles of all  $c$ -labeled edges point to the same node (representing either  $x_j$  or  $\neg x_j$ ), a consistent assignment is obtained that satisfies  $\varphi$ . Conversely, a consistent assignment obviously gives rise to a corresponding derivation of  $G$ , thus showing that the reduction is correct.

Finally, let us note that changing the initial rule to the one shown in the left part of Fig. 4 (using a new terminal  $\diamond$  of rank 2) makes  $H$  satisfy condition 4 as well. This change being made, the input graph is changed by including two copies of the original input, both sharing their leaves, and adding a new root with an outgoing  $\diamond$ -hyperedge targeting the roots of the two copies.  $\square$


If we also disregard restriction 2, the non-uniform membership problem also becomes NP-complete, even if we only consider graphs of height 1.

**Theorem 9.** *There is a DAG grammar  $H$  that conforms to restrictions 1, 3, and 4, such that all graphs in  $\mathcal{L}(H)$  have height 1 and  $\mathcal{L}(H)$  is NP-complete.*



**Fig. 4.** Input graph in the proof of Theorem 8 (left) and modified starting rule (right)

The proof is by reduction from the membership problem for *context-free grammars with disconnecting* (CFGDs), using a result from [8]. A CFGD is an ordinary context-free grammar in Chomsky normal form, with additional rules  $A \rightarrow \diamond$ , where  $\diamond$  is a special symbol that cuts the string into two. Thus, an element in the generated language is a finite multiset of strings rather than a single string. As shown in [8], CFGDs can generate NP-complete languages. We represent a multiset  $\{s_1, \dots, s_k\}$  of strings  $s_i$  as a graph consisting of  $k$  DAGs of height 1 sharing their roots. If  $s_i = a_1 \cdots a_m$  then the DAG representing it consists of the root  $v$ , leaves  $u_0, \dots, u_m$ , and  $a_i$ -hyperedges  $e_i$  with  $\text{src}(e_i) = v$  and  $\text{tar}(e_i) = u_{i-1}u_i$ . Moreover, there are two “unlabeled” terminal edges from  $v$  to  $u_0$  and  $u_n$ , resp. Now, every CFGD can be turned into an equivalent DAG grammar using the schemata in Fig. 5.  $\square$



**Fig. 5.** Rules of a DAG grammar equivalent to a CFGD with initial nonterminal  $S_0$ , from left to right: initial rule,  $A \rightarrow BC$ ,  $A \rightarrow a$ ,  $A \rightarrow \diamond$ .

## 7 Future Work

A number of interesting questions remain open. Is it the case that lifting any one of our five restrictions, while keeping the others, leads to NP-hardness? It seems that the algorithm we propose leads to a fixed-parameter tractable algorithm, with the size of right-hand sides in the grammar as the parameter, when we lift restriction 5 (enforcing that the marking respects  $\leq_F$ ). Is this actually the

case and are there other interesting parameterizations that give tractability for some less restricted classes of grammars? Another open question is whether the algorithm for checking the structure of the input graph and computing the ordering on the leaves can be optimized to run in linear or  $\mathcal{O}(n \log n)$  time.

From a practical point of view, one should study in detail how well suited restricted DAG grammars are for describing linguistic structures such as AMRs. Which phenomena can be modeled in an appropriate manner and which cannot? Are there important aspects in AMRs that can be modeled by general DAG-generating HRGs but not by restricted DAG grammars? If so, can the restrictions be weakened appropriately without sacrificing polynomial parsability?

**Acknowledgements.** We gratefully acknowledge the support from the Swedish Research Council grant 621-2011-6080 and the EU FP7 MICO project. We are furthermore grateful to the anonymous referees for various helpful comments.

## References

1. Aalbersberg, I.J., Ehrenfeucht, A., Rozenberg, G.: On the membership problem for regular DNLC grammars. *Discrete Appl. Math.* **13**, 79–85 (1986)
2. Banarescu, L., Bonial, C., Cai, S., Georgescu, M., Griffitt, K., Hermjakob, U., Knight, K., Koehn, P., Palmer, M., Schneider, N.: Abstract meaning representation for sembanking. In: *Proceedings of 7th Linguistic Annotation Workshop, ACL 2013* (2013)
3. Björklund, H., Drewes, F., Ericson, P.: Between a rock and a hard place - Parsing for hyperedge replacement DAG grammars. Technical report UMINF 15.13, Umeå University (2015). <http://www8.cs.umu.se/research/uminf/index.cgi>
4. Chiang, D., Andreas, J., Bauer, D., Hermann, K.M., Jones, B., Knight, K.: Parsing graphs with hyperedge replacement grammars. In: *Proceedings of 51st Annual Meeting of the Association for Computational Linguistics (ACL 2013)*, vol. 1: Long Papers, pp. 924–932 (2013)
5. Drewes, F., Habel, A., Kreowski, H.J.: Hyperedge replacement graph grammars. In: Rozenberg, G. (ed.) *Handbook of Graph Grammars and Computing by Graph Transformation*, vol. 1: Foundations, chap. 2, pp. 95–162. World Scientific (1997)
6. Drewes, F., Hoffmann, B., Minas, M.: Predictive top-down parsing for hyperedge replacement grammars. In: Parisi-Presicce, F., Westfechtel, B. (eds.) *ICGT 2015*. LNCS, vol. 9151, pp. 19–34. Springer, Heidelberg (2015)
7. Habel, A.: *Hyperedge Replacement: Grammars and Languages*. LNCS, vol. 643. Springer, Heidelberg (1992)
8. Lange, K.J., Welzl, E.: String grammars with disconnecting or a basic root of the difficulty in graph grammar parsing. *Discrete Appl. Math.* **16**, 17–30 (1987)