Bottom-Up Unranked Tree-to-Graph Transducers for Translation into Semantic Graphs

Johanna Björklund
Umeå University
Umeå, Sweden
johanna@cs.umu.se

Shay Cohen
University of Edinburgh
Edinburgh, UK
scohen@inf.ed.ac.uk

Frank Drewes
Umeå University
Umeå, Sweden
drewes@cs.umu.se

Giorgio Satta
University of Padova
Padova, Italy
satta@dei.unipd.it

Abstract

We present a formal model for translating unranked syntactic trees, such as dependency trees, into semantic graphs. The proposed tree-to-graph transducers can serve as a formalization of transition-based systems which recently have been shown to perform very well, yet hitherto lack a suitable formal basis. Our model features "extended" rules and a useful normal form, comes with an efficient translation algorithm, and can in a straightforward manner be equipped with weights.

1 Introduction

In dependency semantic parsing, one is given a natural language sentence and has to output a directed graph representing an associated, most-likely semantic analysis. Semantic parsing integrates tasks that have usually been addressed separately in statistical natural language processing, such as named entity recognition, word sense disambiguation, semantic role labeling, and coreference resolution. Semantic parsing is receiving considerable attention nowadays, as attested by the number of approaches being proposed for its solution (Oepen et al., 2014, 2015) and by the variety of existing semantic representations and available datasets (Kuhlmann and Oepen, 2016).

A successful approach to dependency semantic parsing by Wang et al. (2015b,a) first parses the input sentence into a dependency tree t, and then applies a transition-based algorithm that translates t into a dependency graph in Abstract Meaning Representation (AMR), a popular semantic formalism developed by Banarescu et al. (2013). In this work, we present a finite-state transducer for tree-to-graph translation that can serve as a mathematical model for (Wang et al., 2015b) and, more in general, for work in syntax/semantic interface.

Bottom-up tree transducers (Thatcher, 1973) have gained significant attention in the field of machine translation, where they are used to map syntactic phrase structure trees from source to target languages. This holds in particular for their "extended" version, which may process, in a single step, sections of the input consisting of several symbols; see (Maletti et al., 2009) and references therein. We propose a similar formalism for dependency semantic parsing, mapping syntactic dependency trees into directed graphs that represent the associated semantic interpretation.

When translating dependency trees into graphs in a bottom-up fashion, we face two problems. Firstly, bottom-up tree transducers process *ranked* trees, i.e., the number of children at each node is bounded by some constant. Thus, typically, these tree transducers use a single rule to process in one shot a node along with all of its (previously processed) children in the source tree. In contrast, in the case of dependency trees there is no global constant that limits the number of children a node may have, and processing all of the children by means of a single rule is problematic.

Secondly, in an output tree of a bottom-up tree transducer, nodes that are located near one another are translations of nodes in a source tree that are in close proximity as well. This condition is often referred to as *locality*. Locality does no longer hold true when translating trees into graphs. In fact, so-called reentrancy nodes in a graph have several parents, which are translations of nodes in the source tree whose distance from one another may not be bounded. Reentrancies thus require some form of nonlocal processing, generally not found in tree transducers.

The main contribution of this work is a finite-state tree-to-graph transducer that processes dependency trees in a bottom-up, left-to-right fashion. Our solution to the two problems mentioned above is rather simple. Each node and its children are processed by means of several translation rules, using some sort of independence assumption. Furthermore, in order to implement reentrancy, each translated subtree produces a graph that comes with a record of selected vertices, to be made accessible later in the translation process.

While our transducers use extended translation rules in the sense of (Maletti et al., 2009), they can be cast in a simple normal form, very useful for processing. We provide a polynomial time algorithm for translating an input dependency tree into a packed graph forest, from which each translation graph can efficiently be recovered.

Related work. Bottom-up tree-to-graph transducers were introduced by Engelfriet and Vogler (1994, 1998) who based their work on hyperedge replacement. Since the graph construction mechanism we use is equivalent to hyperedge replacement, our notion of tree-to-graph transducers is essentially an unranked and extended generalization of theirs, except for the fact that ours cannot create multiple copies of unbounded material in the input. This ability seems inappropriate for modeling natural language semantics.

When viewed as an accepting device, our t2g transducers turn into a new type of unranked tree automata that is as space-efficient as the stepwise automata of Carme et al. (2004). In addition, t2g transducers have the advantages of (i) doing away with the need for a binarization of the input tree, and (ii) allowing extended rules that make the representation yet more compact.

The system by Wang et al. (2015b) has inspired our t2g transducers. A technical comparison between their formalism and ours is made in Example 2. An alternative approach to the syntax-semantics interface exploits multi-component synchronous tree-adjoining grammars; see Nesson and Shieber (2006) and references therein. However, these models produce tree-like semantic representations, as opposed to general graphs.

A common approach in semantic parsing is to extend existing syntactic dependency parsers to produce graphs, realizing translation models from strings to graphs, as opposed to the tree-to-graph model investigated here. On this line, transition-based, greedy parsers have been adapted by Ballesteros and Al-Onaizan (2017), Damonte et al. (2017), Peng et al. (2018) and Vilares and Gómez-Rodríguez (2018). Despite the fact that the input is a bare string, these systems exploit features obtained from a precomputed run of a dependency parser, thus committing to some best parse tree, similarly to the pipeline model of Wang et al. (2015b). Dynamic programming parsers have also been adapted to produce graphs by Kuhlmann and Jonsson (2015) and Schluter (2015). Semantic translation from strings to graphs is further investigated by Jones et al. (2012) and Peng et al. (2015) using synchronous hyperedge replacement grammars, who provide unsupervised learning algorithms for grammar extraction.

2 Preliminaries

In this section we introduce some basic notation and terminology that we use throughout this paper.

General Notation. The set of natural numbers (including zero) is denoted by $\mathbb{N}$, and $\mathbb{N}_+ = \mathbb{N} \setminus \{0\}$. For $n \in \mathbb{N}$ the set $\{1, \ldots, n\}$ is abbreviated to $[n]$. In particular, $[0] = \emptyset$. The set of all finite sequences of elements of a set S is written S^*, ε is the empty sequence, $S^+ = S^* \setminus \{\varepsilon\}$, and 2^S is the powerset of S. Given a sequence w, we write $[w]$ for the set of its elements. Concatenation of sequences s, s' is denoted by juxtaposition or, if required for notational clarity, as $s \cdot s'$.

Trees. Let Σ be an alphabet. The set T_Σ of (unranked) trees over Σ is the smallest set such that, for all $f \in \Sigma$ and $t_1, \ldots, t_n \in \mathrm{T}_\Sigma$ ($n \in \mathbb{N}$), we have $f(t_1, \ldots, t_n) \in \mathrm{T}_\Sigma$. In particular $f()$, which we abbreviate by f, is in T_Σ.

The nodes of a tree are identified by their Gorn addresses, which are sequences in $\mathbb{N}_+^*$: the root has the address ε, and if α is the address of a node in t_i then $i\alpha$ is the address of that node in $f(t_1, \ldots, t_n)$. The set of all nodes of t is $N(t)$ and the size of t is $|t| = |N(t)|$.

The label of node α in t is $t(\alpha)$. For $\Sigma' \subseteq \Sigma$, the set of all nodes $\alpha \in N(t)$ with $t(\alpha) \in \Sigma'$ is denoted by $N_{\Sigma'}(t)$. Throughout the paper, a subset $\{\alpha_1, \ldots, \alpha_k\}$ of the set of nodes of a tree t is denoted as $(\alpha_1, \ldots, \alpha_k)$ to stress that its nodes are listed in lexicographic order.

The following notion will play a crucial role in the definition of the translation step for our transducers in Section 3. Let $\square \notin \Sigma$ be a special symbol. A **context** is a tree $c \in \mathrm{T}_{\Sigma \cup \{\square\}}$ that contains exactly one occurrence of $\square$, and this occurrence is a leaf. Given such a context and a tree t, we let $c[t]$ denote the tree obtained from c by replacing $\square$ with t. Thus, $c[t] = t$ if $c = \square$, and otherwise $c[t] = f(s_1, \ldots, s_{i-1}, s_i[t], s_{i+1}, \ldots, s_n)$, where $c = f(s_1, \ldots, s_n)$ and $s_i \in \mathrm{T}_{\Sigma \cup \{\Box\}}$ is the context among $s_1, \ldots, s_n$. For contexts $c \neq \Box$, the notation $c[t]$ is straightforwardly extended to $c[t_1, \ldots, t_k]$ for trees $t_1, \ldots, t_k$ ($k \in \mathbb{N}$). It yields the tree obtained by inserting the sequence of subtrees $t_1, \ldots, t_k$ at the position marked by $\Box$. (This yields a tree since we only use it if $c \neq \Box$.) To be precise, if $c = f(s_1, \ldots, s_n)$ and $i \in [n]$ is the index such that $\Box$ occurs in s_i, then $c[t_1, \ldots, t_k]$ is equal to $f(s_1, \ldots, s_{i-1}, t_1, \ldots, t_k, s_{i+1}, \ldots, s_n)$ if we have $s_i = \Box$; otherwise, it is $f(s_1, \ldots, s_{i-1}, s_i[t_1, \ldots, t_n], s_{i+1}, \ldots, s_n)$.

Graphs. In our graphs, a designated group of vertices, called ports, is used to connect to some outer graph structures in the translation process.

For a given alphabet Δ, the set $\mathcal{G}_\Delta$ of graphs with labels in Δ consists of all quintuples $G = (V, E, lab, port)$ such that

1. V is a finite set of **vertices**,
2. $E \subseteq V \times \Delta \times V$ is the set of labeled **edges**,
3. $lab \colon V \to \Delta$ is a function labeling each vertex, and
4. $port \in V^*$ is a sequence of pairwise distinct vertices called **ports**.

The **size** of G is $|G| = |V| + |E|$. If $port = v_1 \cdots v_n$, then the p-th port v_p of G, $p \in [n]$, is denoted by $port(p)$ and $type(G) = |port|$ is the **type** of G. If the components of G are not explicitly named, they are denoted by V_G, E_G, lab_G, and $port_G$, respectively. For simplicity, we do not use separate sets for vertices and edge labels.

3 Bottom-Up Unranked Tree-to-Graph Transducers

Informally, our transducers process the input tree in a locally bottom-up, left-to-right manner. To apply a translation rule with a left-hand side s at a given node α, s must cover α together with $k \geq 0$ of its leftmost subtrees. Hence, these subtrees must have been processed earlier, to the extent necessary to make the part to be processed identical to s. Applying the rule then removes the subtrees and turns α into a state (or turns it from one state into another, if it already was a state due to an earlier step). Disregarding for the moment the partial output graphs involved, this is depicted schematically in Figure 1.

Note that, in particular, the number k of processed children can be zero, which means that single nodes can initially be turned into states by translation rules whose left-hand sides consist of just one node. More generally, rules in which the root of the left-hand side is an input symbol (with or without children) can be viewed as initializing the processing of the remaining children of that node by turning their parent into an “initial” state.

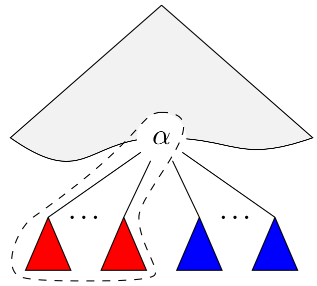

Figure 1: Rule application at node α is locally leftmost (any number, including zero, of the leftmost children of α are consumed) and bottom-up (the left-hand side covers those subtrees all the way down to the leaves). The result of applying a rule at α deletes the subtrees covered by the left-hand side and turns the label of α (a state or input symbol) into a state q.

An (unranked, linear, nondeleting) **bottom-up tree-to-graph transducer** (briefly t2g transducer) is a tuple $tg = (\Sigma, \Delta, Q, R, \mu, F)$ consisting of

1. finite input and output alphabets Σ and Δ;
2. a finite set Q of **states** disjoint with Σ, where every state $q \in Q$ has a type $type(q) \in \mathbb{N}$;
3. a finite set R of translation rules to be defined below;
4. a merging function $\mu \colon 2^\Delta \setminus \{\emptyset\} \to \Delta$; and
5. a set $F \subseteq Q$ of **final states**.

Note that the merging function is finite (because Δ is). It allows us to determine the label of a vertex obtained by merging vertices with different labels. While we do not place any restrictions on μ, it is reasonable to assume that, in linguistic settings, μ will be generated by a binary function in the sense that $\mu(\{\delta\}) = \delta$ and $\mu(\Delta' \cup \{\delta\}) = \mu(\{\mu(\Delta'), \delta\})$ for all $\delta \in \Delta$ and $\Delta' \in 2^\Delta \setminus \{\emptyset\}$. Thus, in this case μ can be efficiently represented by a table of size $|\Delta|^2$.

A **translation rule** $s \to \langle q, G\rangle$ consists of a left-hand side $s \in \mathrm{T}_{\Sigma \cup Q} \setminus Q$ and a right-hand side $\langle q, G\rangle$, where $q \in Q$ and G is a graph with $type(G) = type(q)$. G must fulfill the following additional condition: if $P = \{\alpha{:}p \mid \alpha \in N_Q(s),\ p \in [type(s(\alpha))]\}$ then $G \in \mathcal{G}_{\Delta \cup 2^P}$ and every $\alpha{:}p \in P$ occurs at most once in the labels of vertices in G. A vertex v carrying a label in 2^P is called a **docking vertex**. Intuitively, each

$\alpha{:}p \in lab_G(v)$ is a syntactic name (or formal parameter) referring to the pth port of the graph G_α associated with the node matched by α. During the application of the rule the pth port of G_α will be merged with v. This is formalized next.

A **configuration** of tg is a pair $\langle t, \Gamma \rangle$ with $t \in \mathrm{T}_{\Sigma \cup Q}$ such that $\Gamma\colon N_Q(t) \to \mathcal{G}_\Delta$, where $type(\Gamma(\alpha)) = type(t(\alpha))$ for every $\alpha \in N_Q(t)$. Given an input tree $t_0 \in \mathrm{T}_\Sigma$, the computation of a transducer starts with $\langle t_0, \Gamma_0 \rangle$ where Γ_0 is the function with the domain $N_Q(t_0) = \emptyset$. Suppose inductively that, after some computation steps, a configuration $\langle t, \Gamma \rangle$ has been reached. A translation rule $s \to \langle q, G \rangle$ can be applied to this configuration if t can be written as $t = c[f(t_1, \ldots, t_n)]$, such that $s = f(t_1, \ldots, t_k)$ for some $k \leq n$. If so, let α be the node in c such that $c(\alpha) = \square$. Then there is a computation step $\langle t, \Gamma \rangle \to_{tg} \langle \overline{t}, \overline{\Gamma} \rangle$ with $\overline{t} = c[q(t_{k+1}, \ldots, t_n)]$, where $\overline{\Gamma}$ is as follows:

1. For every node $\overline{\beta} \in N_Q(\overline{t}) \setminus \{\alpha\}$, if β is the corresponding node in t, then $\overline{\Gamma}(\overline{\beta}) = \Gamma(\beta)$.[1]
2. $\overline{\Gamma}(\alpha)$ is obtained as follows:
   First, take the disjoint union of G and all graphs $\Gamma(\alpha\beta)$, $\beta \in N_Q(s)$, the ports of the resulting graph being those of G.
   Second, if a docking vertex $v \in V_G$ has the label $lab_G(v) = \{\alpha_1{:}p_1, \ldots, \alpha_m{:}p_m\}$, merge v with all $v_i = port_{\Gamma(\alpha_i)}(p_i)$ for $i \in [m]$ and label the merged vertex by $\mu(\{lab_{\Gamma(\alpha_1)}(v_1), \ldots, lab_{\Gamma(\alpha_m)}(v_m)\})$.

Example 1 Consider the sentence "The president loves, respects, and fears himself." A slightly simplified Universal Dependencies parse tree of the sentence is shown leftmost in Figure 2. Here, we have removed the "and" node as well as the additional root node above the "loves" node. Further, the edge labels in the tree should be considered as intermediate nodes (since our trees, for simplicity, and in contrast to graphs, do not have edge labels). The figure shows how a t2g transducer may turn the tree into a semantic graph akin to AMR.

In Step 1 we assume for the sake of illustration that the model has seen the leftmost path of the tree ("The president loves") often enough to construct an individual translation rule for it, and that it has also learned that the president referred to is usually Trump. Thus, the translation rule

$$\text{loves}(\texttt{nsubj}(\text{president}(\texttt{det}(\text{The})))) \to \langle q_0, G_0 \rangle$$

turns node "loves" into the state q_0 and its first dependent vanishes. The pair $\langle q_0, \Gamma(\varepsilon) \rangle = \langle q_0, G_0 \rangle$ is illustrated by a dashed box with $\Gamma(\varepsilon)$ shown inside. The numbers next to the vertices indicate the ports. Thus, all three vertices are ports.

In Step 2, we apply a translation rule of the form $q_0(\texttt{conj}(\text{respects})) \to \langle q_{\text{conj}}, G \rangle$ to add two vertices and four edges to the graph. The graph G in the right-hand side is shown in Figure 3. The ports of G become the ports of $\Gamma(\varepsilon)$, and each of the vertices labeled $\varepsilon{:}p$ is merged with the pth port of G_0 (i.e., of the $\Gamma(\varepsilon)$ of the previous step).

Step 3 processes $\text{fears}(\texttt{dobj}(\text{himself}))$, after which Step 4 finally combines the two graphs by applying $q_{\text{conj}}(\texttt{conj}(q_1)) \to \langle q_{\text{f}}, H \rangle$, where H contains a vertex with label $\{\varepsilon{:}2, \varepsilon{:}3, 11{:}2\}$ (see Figure 4). This merges the vertices labeled "Trump", "?", and "himself" into one. Here, we assume that the merging μ of vertex labels gives proper names precedence over pronouns, which in turn take precedence over "?".

Example 2 The transition-based system of Wang et al. (2015b) (and subsequent versions) translates dependency trees into AMRs by visiting nodes and dependency arcs of the input tree bottom-up and left-to-right. At each node or arc, it greedily applies one out of eight alternative actions, turning the tree into a graph. Six actions are local, meaning that they involve nodes at a close distance in the input tree. These include node or arc relabelling, reversing arc directions, deleting a node, and deleting an arc by merging its two nodes. Each of these actions can easily be captured by some individual translation rule of a t2g transducer.

Two remaining actions are nonlocal: one reattaches a node and the other creates a new arc to form a reentrancy. These actions are restricted to local actions for efficiency reasons (Wang et al., 2015b, Section 3.2), so reattachment takes place to the grandparent or great grandparent, and reentrancy involves sibling nodes only. While t2g can simulate reattachment, a major difference between the two models lies in the creation of reentrancies. Wang's system can repeatedly apply the reentrancy action, turning for instance n sibling nodes into a clique, for any n. This is not possible in our model, since translation rules can create reentrancies only by accessing a fixed number

[1] Here, the node corresponding to $\overline{\beta}$ is defined in the obvious way, to take care of the change of Gorn addresses that results from the deletion of $t_1, \ldots, t_k$: if $\overline{\beta} = \alpha i \gamma$ ($i \in \mathbb{N}_+$), then its corresponding node in t is $\alpha(k+i)\gamma$. If α is not a proper prefix of $\overline{\beta}$, then the corresponding node is $\overline{\beta}$ itself.

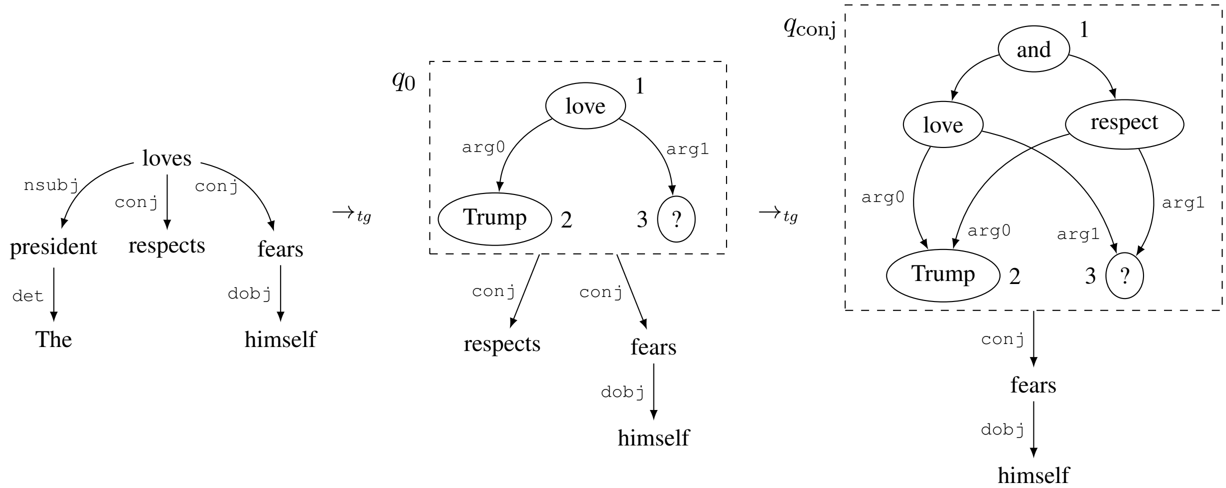

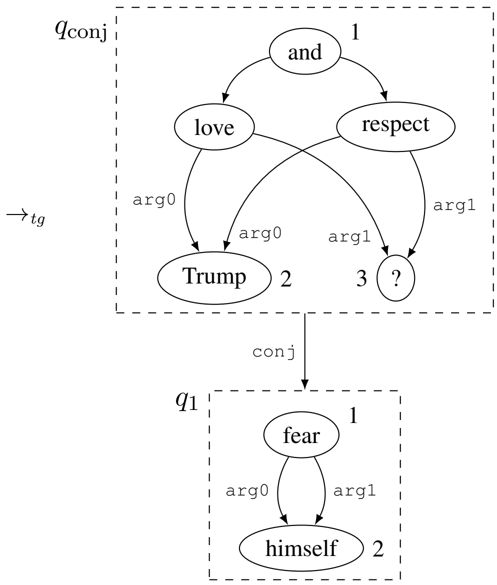

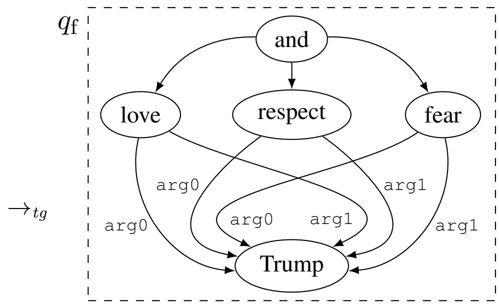

Figure 2: Computation of a t2g transducer applied to the Universal Dependencies parse tree of "The president loves, respects, and fears himself."

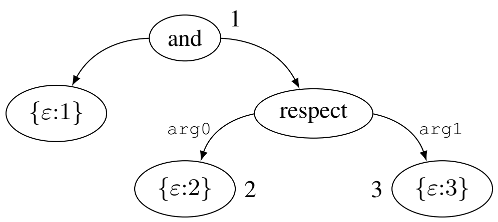

Figure 3: The graph G used in the translation rule $q_0(\texttt{conj}(\text{respects})) \to \langle q_{\text{conj}}, G\rangle$ of Step 2

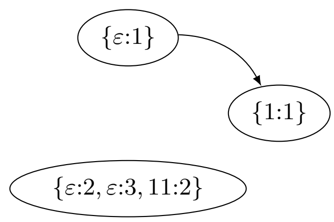

Figure 4: The graph H used in the translation rule $q_{\text{conj}}(\texttt{conj}(q_1)) \to \langle q_{\text{f}}, H\rangle$ of Step 4

of vertices "remembered" as ports. On the other hand, t2g transducers do implement nonlocality, by percolating port nodes at any distance in the underlying derivation tree. This results in reentrancies that extend further than to sibling nodes. Thus, in terms of translation power the two formalisms seem very close to each other in practice, but are formally incomparable.

Readers who are familiar with the concept of hyperedge replacement may have noticed that, except for the role of the merging function, the process described in item 2 in the definition of $\rightarrow_{tg}$ is just hyperedge replacement (where the replaced hyperedges are kept implicit).

A configuration $\langle t, \Gamma\rangle$ is **final** if $t \in F$, i.e., if the first component has been reduced to a single state, which is final. For an input tree t_0, the set of all output graphs computed by tg is denoted by $tg(t_0)$. It is the set of all graphs $\Gamma(\varepsilon)$ such that $\langle t_0, \Gamma_0\rangle \rightarrow^*_{tg} \langle t, \Gamma\rangle$ for some final configuration $\langle t, \Gamma\rangle$. The **transduction** computed by tg is the set $\{(t, g) \in \mathrm{T}_\Sigma \times \mathcal{G}_\Delta \mid g \in tg(t)\}$. The **domain language** of tg is $\{t \in \mathrm{T}_\Sigma \mid tg(t) \neq \emptyset\}$.

4 Derivation Trees

In this section, we describe how a computation of a t2g transducer tg can be represented by means of a tree over the alphabet R, the set of translation rules of tg. We call these trees **derivation trees** of tg. Derivation trees will be used in Section 6 to design efficient translation algorithms. We also show that the derivation trees of tg form a regular tree language (and, in fact, even a local one).

Consider a computation γ of tg that is of form $\langle t_0, \Gamma_0\rangle \rightarrow^+_{tg} \langle q, \Gamma\rangle$ with $t_0 \in \mathrm{T}_\Sigma$ and $q \in Q$. If γ consists of a single step, then a translation rule of the form $r_0 \colon t_0 \rightarrow \langle q, G\rangle$ has been used. In this case the derivation tree associated with γ, written $d(\gamma)$, is simply r_0.

If γ consists of more than one step, assume that at the last step of γ we have used a translation rule r_0 of the form $s \rightarrow \langle q, G\rangle$. We can then write γ as $\langle t_0, \Gamma_0\rangle \rightarrow^+_{tg} \langle s, \Gamma'\rangle \rightarrow_{r_0} \langle q, \Gamma\rangle$. Let γ' denote the first part $\langle t_0, \Gamma_0\rangle \rightarrow^+_{tg} \langle s, \Gamma'\rangle$ of the computation.[2] We define $rk(r_0) = |N_Q(s)|$, called the **rank** of r_0, and we let $N_Q(s) = (\alpha_1, \dots, \alpha_{rk(r_0)})$. In the derivation tree $d(\gamma)$, r_0 has $rk(r_0)$ direct subtrees, the i-th subtree corresponding to the subderivation that ended in the state at α_i. Accordingly, below we split the input tree t_0 into smaller pieces on the basis of the node addresses α_i.

In order to describe this thoroughly, we need to determine a correspondence between nodes in s and nodes in t_0. Intuitively, s is a segment at the top of t_0 that extends to the right. To see this, observe that the root ε of s corresponds to the root of t_0. Some of the children of ε in t_0 may have been consumed by γ' and thus are no longer present in s. However, this happens strictly from left to right. Therefore, if a child of ε in t_0 is still present in s, then all of its siblings to the right are still present, too. The same pattern continues recursively at the children of these nodes in s. The situation is illustrated schematically in Figure 5.

Formally, for a node $\alpha \in N(s)$ we define the

[2]Recall that $N_Q(s)$ is the set of all nodes in s labeled by states in Q, and that we write $N_Q(s) = (\alpha_1, \dots, \alpha_{rk(r_0)})$ to indicate that $\alpha_1, \dots, \alpha_{rk(r_0)}$ are listed in lexicographic order.

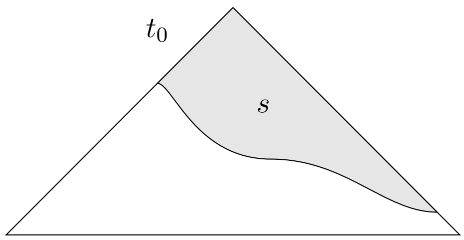

Figure 5: Schematic illustration of the part s of an input tree t_0 which is left after some computation steps. (Note that this is only a structural illustration; some of the node labels in s are not the same as in t_0 anymore, but have been replaced with states.)

pre-image $\overline{\alpha} \in N(t_0)$ of α inductively over the structure of s, as follows:

1. $\overline{\varepsilon} = \varepsilon$.
2. Assume that $\alpha \cdot 1, \dots, \alpha \cdot k \in N(s)$ are the children of a node α in s. Then it should be clear that the children of $\overline{\alpha}$ in t_0 are $\overline{\alpha} \cdot 1, \dots, \overline{\alpha} \cdot n$ for some $n \geq k$. We let $\overline{\alpha \cdot i} = \overline{\alpha} \cdot (i + n - k)$ for all $i \in [k]$.

For a set $N \subseteq N(s)$ of node addresses, we let $\overline{N} = \{\overline{\alpha} \mid \alpha \in N\}$.

We are now ready to split the tree t_0 into the subtrees that, via the computation γ', gave rise to the states at $\alpha_1, \dots, \alpha_{rk(r_0)}$. We do this by defining the sets $N_i \subseteq N(t_0)$ of their nodes, $i \in [rk(r_0)]$. For each $i \in [rk(r_0)]$, N_i is the set of all nodes $\beta \in N(t_0)$ such that $\overline{\alpha_i}$ is the first node in $\overline{N(s)}$ that appears on the path from β to the root of t_0.

Thus, N_i consists of $\overline{\alpha_i}$ and those of its descendants which are not in s anymore, i.e., which have already been "consumed" by the computation γ' in the process of producing node α_i in s. For each $i \in [rk(r_0)]$, define tree t_i as the portion of tree t_0 that is induced by the nodes in N_i.

For every $i \in [rk(r_0)]$, consider the translation rules of γ' that are applied to the nodes in N_i. Clearly, the restriction of γ' to them yields a new computation γ_i of the form $\langle t_i, \Gamma_0\rangle \rightarrow^+_{tg} \langle s(\alpha_i), \Gamma_i\rangle$ whose length is at most the length of γ' and thus less than the length of γ. Let $d(\gamma_i)$ be the derivation tree associated with γ_i. Then we define the derivation tree $d(\gamma)$ to be $r_0(d(\gamma_1), \dots, d(\gamma_{rk(r_0)}))$.

The inductive procedure above associates a unique derivation tree $d(\gamma)$ to each computation γ of tg. Observe that each node of $d(\gamma)$ has a label $r \in R$ and a number of children $rk(r)$. This means that the set of derivation trees of tg is defined over a finite, ranked alphabet. More precisely, it is the set $D(tg)$ recognized by a bottom-up finite-state tree automaton A whose set of states is Q, with F the accepting states. The rules of the automaton A are all $r(q_1, \dots, q_k) \to q$ such that:

1. $r\colon t \to \langle q, G\rangle$ is a translation rule of tg,
2. $N_Q(t) = (\alpha_1, \dots, \alpha_k)$, and
3. $t(\alpha_i) = q_i$ for all $i \in [k]$.

Given a derivation tree $dt \in D(tg)$, such that $dt = r(dt_1, \dots, dt_k)$, we can compute its input tree $in(dt)$ and its output graph $out(dt)$ recursively, as follows. Suppose the root r of dt is the translation rule $r\colon t \to \langle q, G\rangle$ with $N_Q(t) = (\alpha_1, \dots, \alpha_{rk(r)})$.

1. If $t_i = in(dt_i)$ for all $i \in [k]$, then $in(dt)$ is obtained from t and $t_1, \dots, t_k$ by fusing each node α_i with the root of t_i and making $t_i(\varepsilon)$ the label of the fused node. The subtrees of t_i are added to the left of the leftmost subtree of α_i in t. (If α_i is a leaf, t_i just replaces α_i.)
2. If the graph $G_i = out(dt_i)$ for every index $i \in [k]$, then the output graph $out(dt)$ is obtained from the disjoint union of G and $G_1, \dots, G_k$ by merging each docking vertex $v \in V_G$ with ports in $G_1, \dots, G_k$, as follows: if $lab_G(v) = \{\alpha_{i_1}{:}p_1, \dots, \alpha_{i_m}{:}p_m\}$, then v is merged with all $v_j = port_{G_{i_j}}(p_j)$, $j \in [m]$, and the resulting vertex is labeled by $\mu(\{lab_{G_{i_1}}(v_1), \dots, lab_{G_{i_m}}(v_m)\})$.

Note that the definition of $out(dt)$ simply reiterates the way in which computations are defined to construct output graphs. As a consequence, it is a straightforward task to show that $dt = d(\gamma)$ for a computation γ that consumes $in(dt)$ and yields the output graph $out(dt)$.

5 Arc-Factored Normal Form

A translation rule is in **arc-factored normal form** if its left-hand side is in $\Sigma \cup Q(Q)$. A t2g transducer is in arc-factored normal form if each of its translation rules is in arc-factored normal form.

We now show that every t2g transducer tg can effectively be transformed into a t2g transducer in arc-factored normal form which computes the same transduction. First, introduce a new state q_f of type 0 for every $f \in \Sigma$ that occurs in some left-hand side of a translation rule, add the rule $f \to \langle q_f, \emptyset\rangle$ (where $\emptyset$ denotes the empty graph), and replace f by q_f in the left-hand sides of all original rules. Clearly, the computed transduction remains the same and all rules which violate the condition of the arc-factored normal form have left-hand sides in T_Q.

Now, we split rules with large left-hand sides into smaller ones. As long as the transducer is not in arc-factored normal form, select any translation rule $s \to \langle q, G\rangle$ such that $|s| > 2$. Then s has the form $c[q_1(q_2, t_1, \dots, t_n)]$ for some context c, states q_1, q_2, and trees $t_1, \dots, t_n$ ($n \geq 0$). If $k = type(q_1)$ and $\ell = type(q_2)$, we decompose the translation rule into two rules, namely $q_1(q_2) \to \langle q_{1;2}, H\rangle$ and $c[q_{1;2}(t_1, \dots, t_n)] \to \langle q, G'\rangle$, where $q_{1;2}$ is a fresh state with $type(q_{1;2}) = k + \ell$.

The intermediate graph H consists of $k + \ell$ isolated vertices $u_1, \dots, u_k, v_1, \dots, v_\ell$ with $port_H = u_1 \cdots u_k v_1 \cdots v_\ell$ and, for all $i \in [k]$ and $j \in [\ell]$, $lab_H(u_i) = \varepsilon{:}i$ and $lab_H(v_j) = 1{:}j$. The effect of this translation rule is to take the disjoint union of the graphs associated with the two nodes, concatenating the port sequences.

The graph G' is obtained from G by appropriately renaming the references of the form $\alpha{\cdot}1{:}p$ where α is the address of $\Box$ in c: for every $p \in [\ell]$, if $\alpha{\cdot}1{:}p$ occurs in a label of a vertex in G, then it is replaced by $\alpha{:}(\ell + p)$. Moreover, in every label each port reference of the form $\alpha{\cdot}i{:}p$ for $i > 1$ is replaced by $\alpha{\cdot}(i - 1){:}p$.

It should be clear that the two translation rules, executed one after the other, have precisely the same effect as the original one. This completes the proof of the arc-factored normal form.

Note that the size increase implied by the preceding construction is modest. Let us define the size of $tg = (\Sigma, \Delta, Q, R, \mu, F)$ to be the sum of the sizes of its rules. The size of a rule is the size of the left-hand side plus the size of the graph in the right-hand side. By the construction above, each rule will be decomposed into as many rules as there are arcs in the original left-hand side, and the size of graphs in the right-hand sides of intermediate translation rules is at most twice the largest type τ of states in Q. Hence, the total size of the new rules replacing $s \to \langle q, G\rangle$ is $O(|s| \cdot \tau + |G|)$. More sophisticated constructions can result in a smaller transducer. A rather simple optimization is to drop all ports from the discrete graph H which do not occur in G, and to identify those referenced in the label of the same docking vertex. We do not further pursue this here.

6 Translation into a Packed Forest

In this section we investigate the translation problem for our t2g transducers, defined as follows. Given a t2g transducer $tg = (\Sigma, \Delta, Q, R, \mu, F)$ and an unranked tree t, we have to provide a suitable representation for the set of all graphs that are translations of t under tg. We describe the construction only for t2g transducers tg in arc-factored normal form because it is both simpler and significantly more efficient.

We solve the translation problem in two steps. The first step annotates every occurrence of a symbol in t with its address in the tree, yielding $\hat{t}$, and constructs a new t2g transducer $tg_t = (\Sigma', \Delta, Q', R', \mu, F')$. The domain language of tg_t is $\{\hat{t}\}$ and the output graphs of tg_t are the graphs that are translations of t by tg. In the second step, we construct a finite state tree automaton M_t representing the set $D(tg_t)$ of all derivation trees of tg_t. We provide below a more precise description of these steps, but without an overly detailed formalization.

6.1 Grounding

For the first step in our translation algorithm, let $N(\hat{t}) = N(t)$ and $\hat{t}(\alpha) = t(\alpha)^\alpha$ for all $\alpha \in N(t)$. Thus, in $\hat{t}$ each occurrence of a symbol f is annotated with its address α. Next, we restrict the domain language of tg to the set $\{\hat{t}\}$, in such a way that the translation process of tg is "preserved". We call this construction the **grounding** of tg to t. For this, let $k_\alpha = \min\{i \in \mathbb{N}_+ \mid \alpha i \notin N(t)\}$ for every $\alpha \in N(t)$, i.e., k_α is the number of children of α plus one.

1. The input alphabet Σ' consists of all symbols appearing in $\hat{t}$.
2. The set Q' consists of all $\langle q, \alpha, i\rangle$ such that $q \in Q$, $\alpha \in N(t)$, and $i \in [k_\alpha]$. Intuitively, α records the position in the tree and i is the number of the next child to be consumed.
3. The set F' is $\{\langle q, \varepsilon, k_\varepsilon\rangle \mid q \in F\}$.
4. For every translation rule $f \to \langle q, G\rangle$ of tg and every $\alpha \in N(t)$ with $t(\alpha) = f$, we include $f^\alpha \to \langle\langle q, \alpha, 1\rangle, G\rangle$ in R'.
5. For every translation rule $q_1(q_2) \to \langle q, G\rangle$ of tg and every $\alpha i \in N(t)$ ($i \in \mathbb{N}_+$), we let $\langle q_1, \alpha, i\rangle(\langle q_2, \alpha i, k_{\alpha i}\rangle) \to \langle\langle q, \alpha, i+1\rangle, G\rangle$ be a translation rule in R'.

Note that the grounding algorithm above bears close similarity with the notion of parsing by intersection, which makes use of the construction proposed by Bar-Hillel et al. (1964) for producing a context-free grammar that generates the intersection of a language generated by a context-free grammar and a language recognized by a finite (string) automaton. It should thus be clear that $tg_t(\hat{s}) = \emptyset$ for all $s \in \mathrm{T}_\Sigma \setminus \{t\}$, and $tg_t(\hat{t}) = tg(t)$.

The size of the t2g transducer tg_t is the product of the sizes of tg and t. However, in practice many of its translation rules may be useless. It is possible to avoid this by interleaving the construction of the translation rules of tg_t with a simulation of the process of parsing by tg on input t. This has the advantage of pruning the search space, so that useless translation rules are filtered out. We do not further pursue this issue here.

6.2 Graph Forest

In the second step of our translation algorithm, we construct a suitable representation of all the graphs that are obtained in any translation of t based on tg. Using the t2g transducer tg_t from the previous step, we can apply the construction outlined in Section 4 and produce a finite state tree automaton M_t representing the set $D(tg_t)$ of all derivation trees of tg_t.

Together with the interpretation of generated derivation trees dt as $out(dt)$ this yields a compact representation of the set $tg(t)$ of graphs t translates into. We therefore call M_t a **graph forest** for the translation of t under tg.

One can now use standard algorithms to, e.g., generate the graphs of the form $out(dt)$. Further, if the rules of tg are equipped with weights from a suitable weight structure, these weights carry over to the rules of M_t in the obvious way. We can thus use n-best algorithms (Huang and Chiang, 2005) to efficiently extract the n "best" translations of t, e.g., those with the lowest weight.

7 Transducer with a Continuous Memory Register

We assume that the transduction now keeps a general state "processor" that is being updated during the transduction. Each state that is represented in the derivation during the transduction will be replaced by a vector (possibly a typed vector). The states of the machine (the set Q) are represented by vectors in $\mathbb{R}^\ell$ for some ℓ.

In addition, we assume that each node in the dependency tree is associated with a "context vector," for example, created by a TreeLSTM or a contextual word embedding network. This means that each node is accompanied by a vector $v \in \mathbb{R}^m$ for some m.

For our continuous-state transducer, we also assume the existence of a state update function $F\colon \mathbb{R}^m \times \mathbb{R}^\ell \to \mathbb{R}^\ell$. This function consumes a node in the tree, an embedding for graph and an existing state vector, and outputs a new state vector.

Now, the transducer works in a similar way to the case of a discrete state. It traverses the dependency tree, and whenever it needs to apply a rule, it takes some combination of the tree node vectors and the state vector from the current derivation (corresponding to the states in the left-hand side of a rule). It applies F multiple times perhaps, until it gets the final state. It then uses this state vector in a classification to find the best rule to apply at this point. Note that depending on the number of rules in the grammar, this could be quite an expensive classification problem.

Abstractly, the main idea is that there is a register that keeps being updated as the derivation is being generated. It is essentially used to lineraize the states in the derivation into a sequence of states that are being fed to F and it assists at each step to decide what rule to use next.

8 Conclusion

We have developed a novel finite-state transducer that implements nonlocal processing to translate unranked dependency trees into general graphs for semantic representation of natural language. The model can be used as a formalization of transition-based approaches to semantic parsing such as the one of Wang et al. (2015b). The next step in this project is the development of effective algorithms for unsupervised extraction of t2g translation rules from semantic graph corpora.

References

Miguel Ballesteros and Yaser Al-Onaizan. 2017. AMR parsing using stack-LSTMs. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1269–1275. Association for Computational Linguistics.

Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. Abstract meaning representation for sembanking. In *Proc. 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 178–186.

Yehoshua Bar-Hillel, Micha Perles, and Eliyahu Shamir. 1964. On formal properties of simple phrase structure grammars. In Y. Bar-Hillel, editor, *Language and Information: Selected Essays on their Theory and Application*, chapter 9, pages 116–150. Addison-Wesley.

Julien Carme, Joachim Niehren, and Marc Tommasi. 2004. Querying unranked trees with stepwise tree automata. In *Rewriting Techniques and Applications*, pages 105–118, Berlin, Heidelberg. Springer Berlin Heidelberg.

Marco Damonte, B. Shay Cohen, and Giorgio Satta. 2017. An incremental parser for abstract meaning representation. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 536–546. Association for Computational Linguistics.

Joost Engelfriet and Heiko Vogler. 1994. The translation power of top-down tree-to-graph transducers. *Journal of Computer and System Sciences*, 49:258–305.

Joost Engelfriet and Heiko Vogler. 1998. The equivalence of bottom-up and top-down tree-to-graph transducers. *Journal of Computer and System Sciences*, 56(3):332–356.

Liang Huang and David Chiang. 2005. Better k-best parsing. In *Proc. Conf. on Parsing Technology 2005*, pages 53–64. Association for Computational Linguistics.

Bevan Jones, Jacob Andreas, Daniel Bauer, Karl Moritz Hermann, and Kevin Knight. 2012. Semantics-based machine translation with hyperedge replacement grammars. In *Proceedings of COLING 2012*, pages 1359–1376. The COLING 2012 Organizing Committee.

Marco Kuhlmann and Peter Jonsson. 2015. Parsing to noncrossing dependency graphs. *Transactions of the Association for Computational Linguistics*, 3:559–570.

Marco Kuhlmann and Stephan Oepen. 2016. Towards a catalogue of linguistic graph banks. *Computational Linguistics*, 42(4):819–827.

Andreas Maletti, Jonathan Graehl, Mark Hopkins, and Kevin Knight. 2009. The power of extended top-down tree transducers. *SIAM J. Comput.*, 39(2):410–430.

Rebecca Nesson and Stuart M. Shieber. 2006. Simpler TAG semantics through synchronization. In *Proceedings of the 11th Conference on Formal Grammar*, Malaga, Spain.

Stephan Oepen, Marco Kuhlmann, Yusuke Miyao, Daniel Zeman, Silvie Cinkova, Dan Flickinger, Jan Hajic, and Zdenka Uresova. 2015. Semeval 2015 task 18: Broad-coverage semantic dependency parsing. In *Proceedings of the 9th Intl. Workshop on Semantic Evaluation (SemEval 2015)*, pages 915–926.

Stephan Oepen, Marco Kuhlmann, Yusuke Miyao, Daniel Zeman, Dan Flickinger, Jan Hajic, Angelina Ivanova, and Yi Zhang. 2014. Semeval 2014 task 8: Broad-coverage semantic dependency parsing. In *Proc. 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 63–72.

Xiaochang Peng, Daniel Gildea, and Giorgio Satta. 2018. AMR parsing with cache transition systems. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 4897–4904. AAAI Press.

Xiaochang Peng, Linfeng Song, and Daniel Gildea. 2015. A synchronous hyperedge replacement grammar based approach for amr parsing. In *Proceedings of the Nineteenth Conference on Computational Natural Language Learning*, pages 32–41. Association for Computational Linguistics.

Natalie Schluter. 2015. The complexity of finding the maximum spanning DAG and other restrictions for DAG parsing of natural language. In *Proceedings of the Fourth Joint Conference on Lexical and Computational Semantics, *SEM 2015, June 4-5, 2015, Denver, Colorado, USA.*, pages 259–268.

James W. Thatcher. 1973. Tree automata: An informal survey. In A. V. Aho, editor, *Currents in the Theory of Computing*, pages 143–172. Prentice Hall, New York.

David Vilares and Carlos Gómez-Rodríguez. 2018. A transition-based algorithm for unrestricted AMR parsing. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 142–149. Association for Computational Linguistics.

Chuan Wang, Nianwen Xue, and Sameer Pradhan. 2015a. Boosting transition-based AMR parsing with refined actions and auxiliary analyzers. In *Proc. 53rd Annual Meeting of the Association for Computational Linguistics and the 7th Intl. Joint Conf. on Natural Language Processing (Short Papers)*, pages 857–862.

Chuan Wang, Nianwen Xue, and Sameer Pradhan. 2015b. A transition-based algorithm for AMR parsing. In *Proc. 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, (NAACL HLT 2015), Denver, Colorado, USA, 2015*, pages 366–375.