

A Transition-Based Directed Acyclic Graph Parser for UCCA

Daniel Hershcovich^{1,2}

Omri Abend²

Ari Rappoport²

¹The Edmond and Lily Safra Center for Brain Sciences

²School of Computer Science and Engineering

Hebrew University of Jerusalem

{danielh, oabend, arir}@cs.huji.ac.il

Abstract

We present the first parser for UCCA, a cross-linguistically applicable framework for semantic representation, which builds on extensive typological work and supports rapid annotation. UCCA poses a challenge for existing parsing techniques, as it exhibits reentrancy (resulting in DAG structures), discontinuous structures and non-terminal nodes corresponding to complex semantic units. To our knowledge, the conjunction of these formal properties is not supported by any existing parser. Our transition-based parser, which uses a novel transition set and features based on bidirectional LSTMs, has value not just for UCCA parsing: its ability to handle more general graph structures can inform the development of parsers for other semantic DAG structures, and in languages that frequently use discontinuous structures.

1 Introduction

Universal Conceptual Cognitive Annotation (UCCA, Abend and Rappoport, 2013) is a cross-linguistically applicable semantic representation scheme, building on the established Basic Linguistic Theory typological framework (Dixon, 2010a,b, 2012), and Cognitive Linguistics literature (Croft and Cruse, 2004). It has demonstrated applicability to multiple languages, including English, French, German and Czech, support for rapid annotation by non-experts (assisted by an accessible annotation interface (Abend et al., 2017)), and stability under translation (Sulem et al., 2015). It has also proven useful for machine translation evaluation (Birch et al., 2016). UCCA differs from syntactic schemes in terms of content and formal structure. It exhibits reentrancy,

discontinuous nodes and non-terminals, which no single existing parser supports. Lacking a parser, UCCA's applicability has been so far limited, a gap this work addresses.

We present the first UCCA parser, TUPA (Transition-based UCCA Parser), building on recent advances in discontinuous constituency and dependency graph parsing, and further introducing novel transitions and features for UCCA. Transition-based techniques are a natural starting point for UCCA parsing, given the conceptual similarity of UCCA's distinctions, centered around predicate-argument structures, to distinctions expressed by dependency schemes, and the achievements of transition-based methods in dependency parsing (Dyer et al., 2015; Andor et al., 2016; Kiperwasser and Goldberg, 2016). We are further motivated by the strength of transition-based methods in related tasks, including dependency graph parsing (Sagae and Tsujii, 2008; Ribeyre et al., 2014; Tokgöz and Eryiğit, 2015), constituency parsing (Sagae and Lavie, 2005; Zhang and Clark, 2009; Zhu et al., 2013; Maier, 2015; Maier and Lichte, 2016), AMR parsing (Wang et al., 2015a,b, 2016; Misra and Artzi, 2016; Goodman et al., 2016; Zhou et al., 2016; Damonte et al., 2017) and CCG parsing (Zhang and Clark, 2011; Ambati et al., 2015, 2016).

We evaluate TUPA on the English UCCA corpora, including in-domain and out-of-domain settings. To assess the ability of existing parsers to tackle the task, we develop a conversion procedure from UCCA to bilexical graphs and trees. Results show superior performance for TUPA, demonstrating the effectiveness of the presented approach.¹

The rest of the paper is structured as follows:

¹All parsing and conversion code, as well as trained parser models, are available at <https://github.com/danielhers/tupa>.

Section 2 describes UCCA in more detail. Section 3 introduces TUPA. Section 4 discusses the data and experimental setup. Section 5 presents the experimental results. Section 6 summarizes related work, and Section 7 concludes the paper.

2 The UCCA Scheme

UCCA graphs are labeled, directed acyclic graphs (DAGs), whose leaves correspond to the tokens of the text. A node (or *unit*) corresponds to a terminal or to several terminals (not necessarily contiguous) viewed as a single entity according to semantic or cognitive considerations. Edges bear a category, indicating the role of the sub-unit in the parent relation. Figure 1 presents a few examples.

UCCA is a multi-layered representation, where each layer corresponds to a “module” of semantic distinctions. UCCA’s *foundational layer*, targeted in this paper, covers the predicate-argument structure evoked by predicates of all grammatical categories (verbal, nominal, adjectival and others), the inter-relations between them, and other major linguistic phenomena such as coordination and multi-word expressions. The layer’s basic notion is the *scene*, describing a state, action, movement or some other relation that evolves in time. Each scene contains one main relation (marked as either a Process or a State), as well as one or more Participants. For example, the sentence “After graduation, John moved to Paris” (Figure 1a) contains two scenes, whose main relations are “graduation” and “moved”. “John” is a Participant in both scenes, while “Paris” only in the latter. Further categories account for inter-scene relations and the internal structure of complex arguments and relations (e.g. coordination, multi-word expressions and modification).

One incoming edge for each non-root node is marked as *primary*, and the rest (mostly used for implicit relations and arguments) as *remote* edges, a distinction made by the annotator. The primary edges thus form a tree structure, whereas the remote edges enable reentrancy, forming a DAG.

While parsing technology in general, and transition-based parsing in particular, is well-established for syntactic parsing, UCCA has several distinct properties that distinguish it from syntactic representations, mostly UCCA’s tendency to abstract away from syntactic detail that do not affect argument structure. For instance, consider the following examples where the concept of a scene

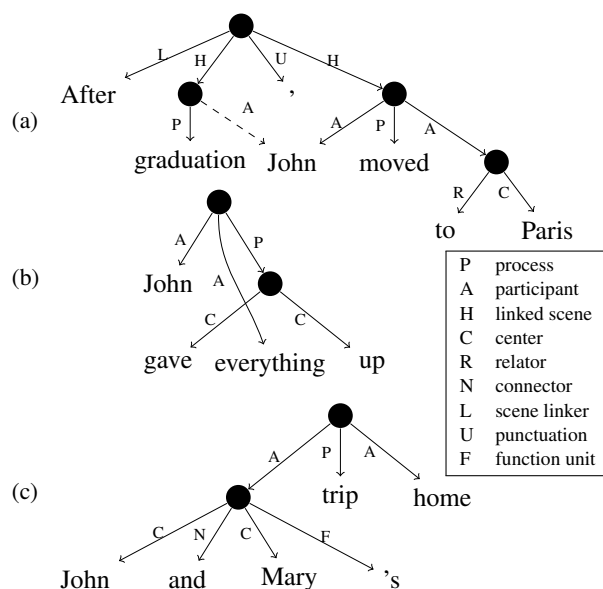


Figure 1: UCCA structures demonstrating three structural properties exhibited by the scheme. (a) includes a remote edge (dashed), resulting in “John” having two parents. (b) includes a discontinuous unit (“gave ... up”). (c) includes a coordination construction (“John and Mary”). Pre-terminal nodes are omitted for brevity. Right: legend of edge labels.

has a different rationale from the syntactic concept of a clause. First, non-verbal predicates in UCCA are represented like verbal ones, such as when they appear in copula clauses or noun phrases. Indeed, in Figure 1a, “graduation” and “moved” are considered separate events, despite appearing in the same clause. Second, in the same example, “John” is marked as a (remote) Participant in the graduation scene, despite not being overtly marked. Third, consider the possessive construction in Figure 1c. While in UCCA “trip” evokes a scene in which “John and Mary” is a Participant, a syntactic scheme would analyze this phrase similarly to “John and Mary’s shoes”.

These examples demonstrate that a UCCA parser, and more generally semantic parsers, face an additional level of ambiguity compared to their syntactic counterparts (e.g., “after graduation” is formally very similar to “after 2pm”, which does not evoke a scene). Section 6 discusses UCCA in the context of other semantic schemes, such as AMR (Banarescu et al., 2013).

Alongside recent progress in dependency parsing into projective trees, there is increasing interest in parsing into representations with more general structural properties (see Section 6). One such property is *reentrancy*, namely the sharing of semantic units between predicates. For instance, in Figure 1a, “John” is an argument of both “gradu-

ation” and “moved”, yielding a DAG rather than a tree. A second property is *discontinuity*, as in Figure 1b, where “gave up” forms a discontinuous semantic unit. Discontinuities are pervasive, e.g., with multi-word expressions (Schneider et al., 2014). Finally, unlike most dependency schemes, UCCA uses *non-terminal nodes* to represent units comprising more than one word. The use of non-terminal nodes is motivated by constructions with no clear head, including co-ordination structures (e.g., “John and Mary” in Figure 1c), some multi-word expressions (e.g., “The Haves and the *Have Nots*”), and prepositional phrases (either the preposition or the head noun can serve as the constituent’s head). To our knowledge, no existing parser supports all structural properties required for UCCA parsing.

3 Transition-based UCCA Parsing

We now turn to presenting TUPA. Building on previous work on parsing reentrancies, discontinuities and non-terminal nodes, we define an extended set of transitions and features that supports the conjunction of these properties.

Transition-based parsers (Nivre, 2003) scan the text from start to end, and create the parse incrementally by applying a *transition* at each step to the parser’s state, defined using three data structures: a buffer B of tokens and nodes to be processed, a stack S of nodes currently being processed, and a graph $G = (V, E, \ell)$ of constructed nodes and edges, where V is the set of *nodes*, E is the set of *edges*, and $\ell : E \rightarrow L$ is the *label* function, L being the set of possible labels. Some states are marked as *terminal*, meaning that G is the final output. A classifier is used at each step to select the next transition based on features encoding the parser’s current state. During training, an oracle creates training instances for the classifier, based on gold-standard annotations.

Transition Set. Given a sequence of tokens w_1, \dots, w_n , we predict a UCCA graph G over the sequence. Parsing starts with a single node on the stack (an artificial root node), and the input tokens in the buffer. Figure 2 shows the transition set.

In addition to the standard SHIFT and REDUCE operations, we follow previous work in transition-based constituency parsing (Sagae and Lavie, 2005), adding the NODE transition for creating new non-terminal nodes. For every $X \in L$, NODE_X creates a new node on the buffer as a par-

ent of the first element on the stack, with an X -labeled edge. LEFT-EDGE_X and RIGHT-EDGE_X create a new primary X -labeled edge between the first two elements on the stack, where the parent is the left or the right node, respectively. As a UCCA node may only have one incoming primary edge, EDGE transitions are disallowed if the child node already has an incoming primary edge. LEFT-REMOTE_X and RIGHT-REMOTE_X do not have this restriction, and the created edge is additionally marked as *remote*. We distinguish between these two pairs of transitions to allow the parser to create remote edges without the possibility of producing invalid graphs. To support the prediction of multiple parents, node and edge transitions leave the stack unchanged, as in other work on transition-based dependency graph parsing (Sagae and Tsujii, 2008; Ribeyre et al., 2014; Tokgöz and Eryiğit, 2015). REDUCE pops the stack, to allow removing a node once all its edges have been created. To handle discontinuous nodes, SWAP pops the second node on the stack and adds it to the top of the buffer, as with the similarly named transition in previous work (Nivre, 2009; Maier, 2015). Finally, FINISH pops the root node and marks the state as terminal.

Classifier. The choice of classifier and feature representation has been shown to play an important role in transition-based parsing (Chen and Manning, 2014; Andor et al., 2016; Kiperwasser and Goldberg, 2016). To investigate the impact of the type of transition classifier in UCCA parsing, we experiment with three different models.

1. Starting with a simple and common choice (e.g., Maier and Lichte, 2016), $\text{TUPA}_{\text{Sparse}}$ uses a linear classifier with sparse features, trained with the averaged structured perceptron algorithm (Collins and Roark, 2004) and MIN-UPDATE (Goldberg and Elhadad, 2011): each feature requires a minimum number of updates in training to be included in the model.²
2. Changing the model to a feedforward neural network with dense embedding features, TUPA_{MLP} (“multi-layer perceptron”), uses an architecture similar to that of Chen and Manning (2014), but with two rectified linear layers

²We also experimented with a linear model using dense embedding features, trained with the averaged structured perceptron algorithm. It performed worse than the sparse perceptron model and was hence discarded.

Before Transition				Transition	After Transition				Terminal?	Condition
Stack	Buffer	Nodes	Edges		Stack	Buffer	Nodes	Edges		
S	$x \mid B$	V	E	SHIFT	$S \mid x$	B	V	E	–	
$S \mid x$	B	V	E	REDUCE	S	B	V	E	–	
$S \mid x$	B	V	E	NODE _X	$S \mid x$	$y \mid B$	$V \cup \{y\}$	$E \cup \{(y, x)_X\}$	–	$x \neq \text{root}$
$S \mid y, x$	B	V	E	LEFT-EDGE _X	$S \mid y, x$	B	V	$E \cup \{(x, y)_X\}$	–	$\left\{ \begin{array}{l} x \notin w_{1:n}, \\ y \neq \text{root}, \\ y \not\sim_G x \end{array} \right.$
$S \mid x, y$	B	V	E	RIGHT-EDGE _X	$S \mid x, y$	B	V	$E \cup \{(x, y)_X\}$	–	
$S \mid y, x$	B	V	E	LEFT-REMOTE _X	$S \mid y, x$	B	V	$E \cup \{(x, y)_X^*\}$	–	
$S \mid x, y$	B	V	E	RIGHT-REMOTE _X	$S \mid x, y$	B	V	$E \cup \{(x, y)_X^*\}$	–	
$S \mid x, y$	B	V	E	SWAP	$S \mid y$	$x \mid B$	V	E	–	$i(x) < i(y)$
[root]	\emptyset	V	E	FINISH	\emptyset	\emptyset	V	E	+	

Figure 2: The transition set of TUPA. We write the stack with its top to the right and the buffer with its head to the left. $(\cdot, \cdot)_X$ denotes a primary X -labeled edge, and $(\cdot, \cdot)_X^*$ a remote X -labeled edge. $i(x)$ is a running index for the created nodes. In addition to the specified conditions, the prospective child in an EDGE transition must not already have a primary parent.

instead of one layer with cube activation. The embeddings and classifier are trained jointly.

- Finally, **TUPA_{BiLSTM}** uses a bidirectional LSTM for feature representation, on top of the dense embedding features, an architecture similar to Kiperwasser and Goldberg (2016). The BiLSTM runs on the input tokens in forward and backward directions, yielding a vector representation that is then concatenated with dense features representing the parser state (e.g., existing edge labels and previous parser actions; see below). This representation is then fed into a feedforward network similar to TUPA_{MLP}. The feedforward layers, BiLSTM and embeddings are all trained jointly.

For all classifiers, inference is performed greedily, i.e., without beam search. Hyperparameters are tuned on the development set (see Section 4).

Features. TUPA_{Sparse} uses binary indicator features representing the words, POS tags, syntactic dependency labels and existing edge labels related to the top four stack elements and the next three buffer elements, in addition to their children and grandchildren in the graph. We also use bi- and trigram features based on these values (Zhang and Clark, 2009; Zhu et al., 2013), features related to discontinuous nodes (Maier, 2015, including separating punctuation and gap type), features representing existing edges and the number of parents and children, as well as the past actions taken by the parser. In addition, we use a novel, UCCA-specific feature: number of remote children.³

For TUPA_{MLP} and TUPA_{BiLSTM}, we replace all indicator features by a concatenation of the vector embeddings of all represented elements: words,

POS tags, syntactic dependency labels, edge labels, punctuation, gap type and parser actions. These embeddings are initialized randomly. We additionally use external word embeddings initialized with pre-trained word2vec vectors (Mikolov et al., 2013),⁴ updated during training. In addition to dropout between NN layers, we apply word dropout (Kiperwasser and Goldberg, 2016): with a certain probability, the embedding for a word is replaced with a zero vector. We do not apply word dropout to the external word embeddings.

Finally, for all classifiers we add a novel real-valued feature to the input vector, **ratio**, corresponding to the ratio between the number of terminals to number of nodes in the graph G . This feature serves as a regularizer for the creation of new nodes, and should be beneficial for other transition-based constituency parsers too.

Training. For training the transition classifiers, we use a dynamic oracle (Goldberg and Nivre, 2012), i.e., an oracle that outputs a set of optimal transitions: when applied to the current parser state, the gold standard graph is reachable from the resulting state. For example, the oracle would predict a NODE transition if the stack has on its top a parent in the gold graph that has not been created, but would predict a RIGHT-EDGE transition if the second stack element is a parent of the first element according to the gold graph and the edge between them has not been created. The transition predicted by the classifier is deemed correct and is applied to the parser state to reach the subsequent state, if the transition is included in the set of optimal transitions. Otherwise, a random optimal transition is applied, and for the perceptron-based parser, the classifier’s weights are updated

³See Appendix A for a full list of used feature templates.

⁴<https://goo.gl/6ovEhC>

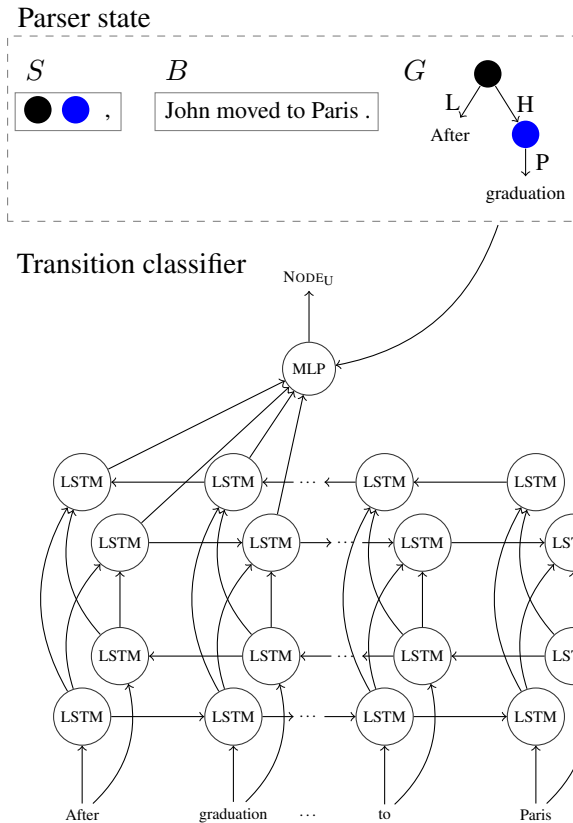


Figure 3: Illustration of the TUPA model. Top: parser state (stack, buffer and intermediate graph). Bottom: TUPA_{BILSTM} architecture. Vector representation for the input tokens is computed by two layers of bidirectional LSTMs. The vectors for specific tokens are concatenated with embedding and numeric features from the parser state (for existing edge labels, number of children, etc.), and fed into the MLP for selecting the next transition.

according to the perceptron update rule.

POS tags and syntactic dependency labels are extracted using spaCy (Honnibal and Johnson, 2015).⁵ We use the categorical cross-entropy objective function and optimize the NN classifiers with the Adam optimizer (Kingma and Ba, 2014).

4 Experimental Setup

Data. We conduct our experiments on the UCCA Wikipedia corpus (henceforth, *Wiki*), and use the English part of the UCCA *Twenty Thousand Leagues Under the Sea* English-French parallel corpus (henceforth, *20K Leagues*) as out-of-domain data.⁶ Table 1 presents some statistics for the two corpora. We use passages of indices up to 676 of the *Wiki* corpus as our training set, passages 688–808 as development set, and passages 942–1028 as in-domain test set. While

⁵<https://spacy.io>

⁶<http://cs.huji.ac.il/~oabend/ucca.html>

	Wiki			20K
	Train	Dev	Test	Leagues
# passages	300	34	33	154
# sentences	4268	454	503	506
# nodes	298,993	33,704	35,718	29,315
% terminal	42.96	43.54	42.87	42.09
% non-term.	58.33	57.60	58.35	60.01
% discont.	0.54	0.53	0.44	0.81
% reentrant	2.38	1.88	2.15	2.03
# edges	287,914	32,460	34,336	27,749
% primary	98.25	98.75	98.74	97.73
% remote	1.75	1.25	1.26	2.27
Average per non-terminal node				
# children	1.67	1.68	1.66	1.61

Table 1: Statistics of the *Wiki* and *20K Leagues* UCCA corpora. All counts exclude the root node, implicit nodes, and linkage nodes and edges.

UCCA edges can cross sentence boundaries, we adhere to the common practice in semantic parsing and train our parsers on individual sentences, discarding inter-relations between them (0.18% of the edges). We also discard linkage nodes and edges (as they often express inter-sentence relations and are thus mostly redundant when applied at the sentence level) as well as implicit nodes.⁷ In the out-of-domain experiments, we apply the same parsers (trained on the *Wiki* training set) to the *20K Leagues* corpus without parameter re-tuning.

Implementation. We use the DyNet package (Neubig et al., 2017) for implementing the NN classifiers. Unless otherwise noted, we use the default values provided by the package. See Appendix C for the hyperparameter values we found by tuning on the development set.

Evaluation. We define a simple measure for comparing UCCA structures $G_p = (V_p, E_p, \ell_p)$ and $G_g = (V_g, E_g, \ell_g)$, the predicted and gold-standard graphs, respectively, over the same sequence of terminals $W = \{w_1, \dots, w_n\}$. For an edge $e = (u, v)$ in either graph, u being the parent and v the child, its yield $y(e) \subseteq W$ is the set of terminals in W that are descendants of v . Define the set of *mutual edges* between G_p and G_g :

$$M(G_p, G_g) = \{(e_1, e_2) \in E_p \times E_g \mid y(e_1) = y(e_2) \wedge \ell_p(e_1) = \ell_g(e_2)\}$$

Labeled precision and recall are defined by dividing $|M(G_p, G_g)|$ by $|E_p|$ and $|E_g|$, respectively, and F-score by taking their harmonic mean.

⁷Appendix B further discusses linkage and implicit units.

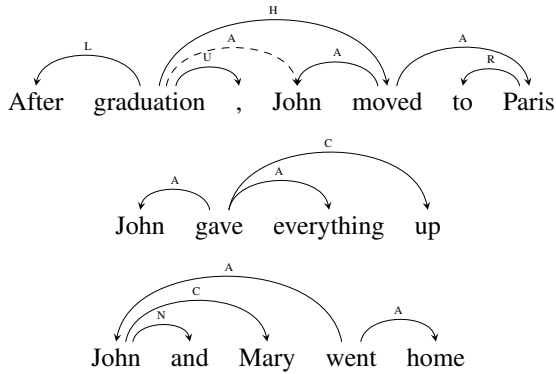


Figure 4: Bilingual graph approximation (dependency graph) for the sentences in Figure 1.

We report two variants of this measure: one where we consider only primary edges, and another for remote edges (see Section 2). Performance on remote edges is of pivotal importance in this investigation, which focuses on extending the class of graphs supported by statistical parsers.

We note that the measure collapses to the standard PARSEVAL constituency evaluation measure if G_p and G_g are trees. Punctuation is excluded from the evaluation, but not from the datasets.

Comparison to bilingual graph parsers. As no direct comparison with existing parsers is possible, we compare TUPA to bilingual dependency graph parsers, which support reentrancy and discontinuity but not non-terminal nodes.

To facilitate the comparison, we convert our training set into bilingual graphs (see examples in Figure 4), train each of the parsers, and evaluate them by applying them to the test set and then reconstructing UCCA graphs, which are compared with the gold standard. The conversion to bilingual graphs is done by heuristically selecting a head terminal for each non-terminal node, and attaching all terminal descendants to the head terminal. In the inverse conversion, we traverse the bilingual graph in topological order, creating non-terminal parents for all terminals, and attaching them to the previously-created non-terminals corresponding to the bilingual heads.⁸

In Section 5 we report the upper bounds on the achievable scores due to the error resulting from the removal of non-terminal nodes.

Comparison to tree parsers. For completeness, and as parsing technology is considerably more

⁸See Appendix D for a detailed description of the conversion procedures.




Figure 5: Tree approximation (constituency) for the sentence in Figure 1a (top), and bilingual tree approximation (dependency) for the same sentence (bottom). These are identical to the original graphs, apart from the removal of remote edges.

mature for tree (rather than graph) parsing, we also perform a *tree approximation* experiment, converting UCCA to (bilingual) trees and evaluating constituency and dependency tree parsers on them (see examples in Figure 5). Our approach is similar to the tree approximation approach used for dependency graph parsing (Agić et al., 2015; Fernández-González and Martins, 2015), where dependency graphs were converted into dependency trees and then parsed by dependency tree parsers. In our setting, the conversion to trees consists simply of removing remote edges from the graph, and then to bilingual trees by applying the same procedure as for bilingual graphs.

Baseline parsers. We evaluate two bilingual graph semantic dependency parsers: DAGParser (Ribeyre et al., 2014), the leading transition-based parser in SemEval 2014 (Oepen et al., 2014) and TurboParser (Almeida and Martins, 2015), a graph-based parser from SemEval 2015 (Oepen et al., 2015); UPARSE (Maier and Lichte, 2016), a transition-based constituency parser supporting discontinuous constituents; and two bilingual tree parsers: MaltParser (Nivre et al., 2007), and the stack LSTM-based parser of Dyer et al. (2015, henceforth “LSTM Parser”). Default settings are used in all cases.⁹ DAGParser and UPARSE use beam search by default, with a beam size of 5 and 4 respectively. The other parsers are greedy.

5 Results

Table 2 presents our main experimental results, as well as upper bounds for the baseline parsers, re-

⁹For MaltParser we use the ARCEAGER transition set and SVM classifier. Other configurations yielded lower scores.

	Wiki (in-domain)						20K Leagues (out-of-domain)					
	Primary			Remote			Primary			Remote		
	LP	LR	LF	LP	LR	LF	LP	LR	LF	LP	LR	LF
TUPA _{Sparse}	64.5	63.7	64.1	19.8	13.4	16	59.6	59.9	59.8	22.2	7.7	11.5
TUPA _{MLP}	65.2	64.6	64.9	23.7	13.2	16.9	62.3	62.6	62.5	20.9	6.3	9.7
TUPA _{BILSTM}	74.4	72.7	73.5	47.4	51.6	49.4	68.7	68.5	68.6	38.6	18.8	25.3
Bilexical Approximation (Dependency DAG Parsers)												
Upper Bound	91			58.3			91.3			43.4		
DAGParser	61.8	55.8	58.6	9.5	0.5	1	56.4	50.6	53.4	–	0	0
TurboParser	57.7	46	51.2	77.8	1.8	3.7	50.3	37.7	43.1	100	0.4	0.8
Tree Approximation (Constituency Tree Parser)												
Upper Bound	100			–			100			–		
UPARSE	60.9	61.2	61.1	–	–	–	52.7	52.8	52.8	–	–	–
Bilexical Tree Approximation (Dependency Tree Parsers)												
Upper Bound	91			–			91.3			–		
MaltParser	62.8	57.7	60.2	–	–	–	57.8	53	55.3	–	–	–
LSTM Parser	73.2	66.9	69.9	–	–	–	66.1	61.1	63.5	–	–	–

Table 2: Experimental results, in percents, on the *Wiki* test set (left) and the *20K Leagues* set (right). Columns correspond to labeled precision, recall and F-score, for both primary and remote edges. F-score upper bounds are reported for the conversions. For the tree approximation experiments, only primary edges scores are reported, as they are unable to predict remote edges. TUPA_{BILSTM} obtains the highest F-scores in all metrics, surpassing the bilexical parsers, tree parsers and other classifiers.

flecting the error resulting from the conversion.¹⁰

DAGParser and UPARSE are most directly comparable to TUPA_{Sparse}, as they also use a perceptron classifier with sparse features. TUPA_{Sparse} considerably outperforms both, where DAGParser does not predict any remote edges in the out-of-domain setting. TurboParser fares worse in this comparison, despite somewhat better results on remote edges. The LSTM parser of Dyer et al. (2015) obtains the highest primary F-score among the baseline parsers, with a considerable margin.

Using a feedforward NN and embedding features, TUPA_{MLP} obtains higher scores than TUPA_{Sparse}, but is outperformed by the LSTM parser on primary edges. However, using better input encoding allowing virtual look-ahead and look-behind in the token representation, TUPA_{BILSTM} obtains substantially higher scores than TUPA_{MLP} and all other parsers, on both primary and remote edges, both in the in-domain and out-of-domain settings. Its performance in absolute terms, of 73.5% F-score on primary edges, is encouraging in light of UCCA’s inter-annotator agreement of 80–85% F-score on them (Abend and Rappoport, 2013).

The parsers resulting from tree approximation

¹⁰The low upper bound for remote edges is partly due to the removal of implicit nodes (not supported in bilexical representations), where the whole sub-graph headed by such nodes, often containing remote edges, must be discarded.

are unable to recover any remote edges, as these are removed in the conversion.¹¹ The bilexical DAG parsers are quite limited in this respect as well. While some of the DAG parsers’ difficulty can be attributed to the conversion upper bound of 58.3%, this in itself cannot account for their poor performance on remote edges, which is an order of magnitude lower than that of TUPA_{BILSTM}.

6 Related Work

While earlier work on anchored¹² semantic parsing has mostly concentrated on shallow semantic analysis, focusing on semantic role labeling of verbal argument structures, the focus has recently shifted to parsing of more elaborate representations that account for a wider range of phenomena (Abend and Rappoport, 2017).

Grammar-Based Parsing. Linguistically expressive grammars such as HPSG (Pollard and Sag, 1994), CCG (Steedman, 2000) and TAG (Joshi and Schabes, 1997) provide a theory of the syntax-semantics interface, and have been used as a basis for semantic parsers by defining com-

¹¹We also experimented with a simpler version of TUPA lacking REMOTE transitions, obtaining an increase of up to 2 labeled F-score points on primary edges, at the cost of not being able to predict remote edges.

¹²By *anchored* we mean that the semantic representation directly corresponds to the words and phrases of the text.

positional semantics on top of them (Flickinger, 2000; Bos, 2005, among others). Depending on the grammar and the implementation, such semantic parsers can support some or all of the structural properties UCCA exhibits. Nevertheless, this line of work differs from our approach in two important ways. First, the *representations* are different. UCCA does not attempt to model the syntax-semantics interface and is thus less coupled with syntax. Second, while grammar-based *parsers* explicitly model syntax, our approach directly models the relation between tokens and semantic structures, without explicit composition rules.

Broad-Coverage Semantic Parsing. Most closely related to this work is Broad-Coverage Semantic Dependency Parsing (SDP), addressed in two SemEval tasks (Oepen et al., 2014, 2015). Like UCCA parsing, SDP addresses a wide range of semantic phenomena, and supports discontinuous units and reentrancy. In SDP, however, bilocal dependencies are used, and a head must be selected for every relation—even in constructions that have no clear head, such as coordination (Ivanova et al., 2012). The use of non-terminal nodes is a simple way to avoid this liability. SDP also differs from UCCA in the type of distinctions it makes, which are more tightly coupled with syntactic considerations, where UCCA aims to capture purely semantic cross-linguistically applicable notions. For instance, the “poss” label in the DM target representation is used to annotate syntactic possessive constructions, regardless of whether they correspond to semantic ownership (e.g., “John’s dog”) or other semantic relations, such as marking an argument of a nominal predicate (e.g., “John’s kick”). UCCA reflects the difference between these constructions.

Recent interest in SDP has yielded numerous works on graph parsing (Ribeyre et al., 2014; Thomson et al., 2014; Almeida and Martins, 2015; Du et al., 2015), including tree approximation (Agić and Koller, 2014; Schluter et al., 2014) and joint syntactic/semantic parsing (Henderson et al., 2013; Swayamdipta et al., 2016).

Abstract Meaning Representation. Another line of work addresses parsing into AMRs (Flanigan et al., 2014; Vanderwende et al., 2015; Pust et al., 2015; Artzi et al., 2015), which, like UCCA, abstract away from syntactic distinctions and represent meaning directly, using OntoNotes predi-

cates (Weischedel et al., 2013). Events in AMR may also be evoked by non-verbal predicates, including possessive constructions.

Unlike in UCCA, the alignment between AMR concepts and the text is not explicitly marked. While sharing much of this work’s motivation, not anchoring the representation in the text complicates the parsing task, as it requires the alignment to be automatically (and imprecisely) detected. Indeed, despite considerable technical effort (Flanigan et al., 2014; Pourdamghani et al., 2014; Werling et al., 2015), concept identification is only about 80%–90% accurate. Furthermore, anchoring allows breaking down sentences into semantically meaningful sub-spans, which is useful for many applications (Fernández-González and Martins, 2015; Birch et al., 2016).

Several transition-based AMR parsers have been proposed: CAMR assumes syntactically parsed input, processing dependency trees into AMR (Wang et al., 2015a,b, 2016; Goodman et al., 2016). In contrast, the parsers of Damonte et al. (2017) and Zhou et al. (2016) do not require syntactic pre-processing. Damonte et al. (2017) perform concept identification using a simple heuristic selecting the most frequent graph for each token, and Zhou et al. (2016) perform concept identification and parsing jointly. UCCA parsing does not require separately aligning the input tokens to the graph. TUPA creates non-terminal units as part of the parsing process.

Furthermore, existing transition-based AMR parsers are not general DAG parsers. They are only able to predict a subset of reentrancies and discontinuities, as they may remove nodes before their parents have been predicted (Damonte et al., 2017). They are thus limited to a sub-class of AMRs in particular, and specifically cannot produce arbitrary DAG parses. TUPA’s transition set, on the other hand, allows general DAG parsing.¹³

7 Conclusion

We present TUPA, the first parser for UCCA. Evaluated in in-domain and out-of-domain settings, we show that coupled with a NN classifier and BiLSTM feature extractor, it accurately predicts UCCA graphs from text, outperforming a variety of strong baselines by a margin.

Despite the recent diversity of semantic pars-

¹³See Appendix E for a proof sketch for the completeness of TUPA’s transition set.

ing work, the effectiveness of different approaches for structurally and semantically different schemes is not well-understood (Kuhlmann and Oepen, 2016). Our contribution to this literature is a general parser that supports multiple parents, discontinuous units and non-terminal nodes.

Future work will evaluate TUPA in a multi-lingual setting, assessing UCCA’s cross-linguistic applicability. We will also apply the TUPA transition scheme to different target representations, including AMR and SDP, exploring the limits of its generality. In addition, we will explore different conversion procedures (Kong et al., 2015) to compare different representations, suggesting ways for a data-driven design of semantic annotation.

A parser for UCCA will enable using the framework for new tasks, in addition to existing applications such as machine translation evaluation (Birch et al., 2016). We believe UCCA’s merits in providing a cross-linguistically applicable, broad-coverage annotation will support ongoing efforts to incorporate deeper semantic structures into various applications, such as sentence simplification (Narayan and Gardent, 2014) and summarization (Liu et al., 2015).

Acknowledgments

This work was supported by the HUJI Cyber Security Research Center in conjunction with the Israel National Cyber Bureau in the Prime Minister’s Office, and by the Intel Collaborative Research Institute for Computational Intelligence (ICRI-CI). The first author was supported by a fellowship from the Edmond and Lily Safra Center for Brain Sciences. We thank Wolfgang Maier, Nathan Schneider, Elior Sulem and the anonymous reviewers for their helpful comments.

References

Omri Abend and Ari Rappoport. 2013. [Universal Conceptual Cognitive Annotation \(UCCA\)](#). In *Proc. of ACL*. pages 228–238. <http://aclweb.org/anthology/P13-1023>.

Omri Abend and Ari Rappoport. 2017. The state of the art in semantic representation. In *Proc. of ACL*. To appear.

Omri Abend, Shai Yerushalmi, and Ari Rappoport. 2017. UCCAApp: Web-application for syntactic and semantic phrase-based annotation. In *Proc. of ACL: System Demonstration Papers*. To appear.

Željko Agić and Alexander Koller. 2014. [Potsdam: Semantic dependency parsing by bidirectional graph-tree transformations and syntactic parsing](#). In *Proc. of SemEval*. pages 465–470. <http://aclweb.org/anthology/S14-2081>.

Željko Agić, Alexander Koller, and Stephan Oepen. 2015. [Semantic dependency graph parsing using tree approximations](#). In *Proc. of IWCS*. pages 217–227. <http://aclweb.org/anthology/W15-0126>.

Mariana S. C. Almeida and André F. T. Martins. 2015. [Lisbon: Evaluating TurboSemanticParser on multiple languages and out-of-domain data](#). In *Proc. of SemEval*. pages 970–973. <http://aclweb.org/anthology/S15-2162>.

Bharat Ram Ambati, Tejaswini Deoskar, Mark Johnson, and Mark Steedman. 2015. [An incremental algorithm for transition-based CCG parsing](#). In *Proc. of NAACL*. pages 53–63. <http://aclweb.org/anthology/N15-1006>.

Bharat Ram Ambati, Tejaswini Deoskar, and Mark Steedman. 2016. [Shift-reduce CCG parsing using neural network models](#). In *Proc. of NAACL-HLT*. pages 447–453. <http://aclweb.org/anthology/N16-1052>.

Daniel Andor, Chris Alberti, David Weiss, Aliaksei Severyn, Alessandro Presta, Kuzman Ganchev, Slav Petrov, and Michael Collins. 2016. [Globally normalized transition-based neural networks](#). In *Proc. of ACL*. pages 2442–2452. <http://aclweb.org/anthology/P16-1231>.

Yoav Artzi, Kenton Lee, and Luke Zettlemoyer. 2015. [Broad-coverage CCG semantic parsing with AMR](#). In *Proc. of EMNLP*. pages 1699–1710. <http://aclweb.org/anthology/D15-1198>.

Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Martha Palmer, and Nathan Schneider. 2013. [Abstract Meaning Representation for semantics](#). In *Proc. of the Linguistic Annotation Workshop*. <http://aclweb.org/anthology/W13-2322>.

Alexandra Birch, Omri Abend, Ondřej Bojar, and Barry Haddow. 2016. [HUME: Human UCCA-based evaluation of machine translation](#). In *Proc. of EMNLP*. pages 1264–1274. <http://aclweb.org/anthology/D16-1134>.

Johan Bos. 2005. [Towards wide-coverage semantic interpretation](#). In *Proc. of IWCS*. volume 6, pages 42–53. <http://www.let.rug.nl/bos/pubs/Bos2005IWCS.pdf>.

Danqi Chen and Christopher Manning. 2014. [A fast and accurate dependency parser using neural networks](#). In *Proc. of EMNLP*. pages 740–750. <http://aclweb.org/anthology/D14-1082>.

- Michael Collins and Brian Roark. 2004. Incremental parsing with the perceptron algorithm. In *Proc. of ACL*. pages 111–118. <http://aclweb.org/anthology/P04-1015>.
- William Croft and D Alan Cruse. 2004. *Cognitive linguistics*. Cambridge University Press.
- Marco Damonte, Shay B. Cohen, and Giorgio Satta. 2017. An incremental parser for abstract meaning representation. In *Proceedings of EACL*. <http://homepages.inf.ed.ac.uk/scohen/eacl17amr.pdf>.
- Robert M. W. Dixon. 2010a. *Basic Linguistic Theory: Grammatical Topics*, volume 2. Oxford University Press.
- Robert M. W. Dixon. 2010b. *Basic Linguistic Theory: Methodology*, volume 1. Oxford University Press.
- Robert M. W. Dixon. 2012. *Basic Linguistic Theory: Further Grammatical Topics*, volume 3. Oxford University Press.
- Yantao Du, Fan Zhang, Xun Zhang, Weiwei Sun, and Xiaojun Wan. 2015. Peking: Building semantic dependency graphs with a hybrid parser. In *Proc. of SemEval*. pages 927–931. <http://aclweb.org/anthology/S15-2154>.
- Chris Dyer, Miguel Ballesteros, Wang Ling, Austin Matthews, and Noah A. Smith. 2015. Transition-based dependency parsing with stack long short-term memory. In *Proc. of ACL*. pages 334–343. <http://aclweb.org/anthology/P15-1033>.
- Daniel Fernández-González and André FT Martins. 2015. Parsing as reduction. In *Proc. of ACL*. pages 1523–1533. <http://aclweb.org/anthology/P15-1147>.
- Jeffrey Flanigan, Sam Thomson, Jaime Carbonell, Chris Dyer, and Noah A. Smith. 2014. A discriminative graph-based parser for the abstract meaning representation. In *Proc. of ACL*. pages 1426–1436. <http://aclweb.org/anthology/P14-1134>.
- Daniel Flickinger. 2000. On building a more efficient grammar by exploiting types. In *Collaborative Language Engineering*, CLSI, Stanford, CA, volume 6, pages 15–28.
- Yoav Goldberg and Michael Elhadad. 2011. Learning sparser perceptron models. Technical report. <http://www.cs.bgu.ac.il/~yoav/publications>.
- Yoav Goldberg and Joakim Nivre. 2012. A dynamic oracle for arc-eager dependency parsing. In *Proc. of COLING*. pages 959–976. <http://aclweb.org/anthology/C12-1059>.
- James Goodman, Andreas Vlachos, and Jason Naradowsky. 2016. Noise reduction and targeted exploration in imitation learning for Abstract Meaning Representation parsing. In *Proc. of ACL*. pages 1–11. <http://aclweb.org/anthology/P16-1001>.
- James Henderson, Paola Merlo, Ivan Titov, and Gabriele Musillo. 2013. Multilingual joint parsing of syntactic and semantic dependencies with a latent variable model. *Computational Linguistics* 39(4):949–998. <http://cognet.mit.edu/node/27348>.
- Matthew Honnibal and Mark Johnson. 2015. An improved non-monotonic transition system for dependency parsing. In *Proc. of EMNLP*. pages 1373–1378. <http://aclweb.org/anthology/D15-1162>.
- Angelina Ivanova, Stephan Oepen, Lilja Øvrelid, and Dan Flickinger. 2012. Who did what to whom? A contrastive study of syntacto-semantic dependencies. In *Proc. of LAW*. pages 2–11. <http://aclweb.org/anthology/W12-3602>.
- Aravind Joshi and Yves Schabes. 1997. Tree-Adjoining Grammars. In Grzegorz Rozenberg and Arto Salomaa, editors, *Handbook of Formal Languages*, Springer, Berlin, volume 3, pages 69–124.
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *CoRR* abs/1412.6980. <http://arxiv.org/abs/1412.6980>.
- Eliyahu Kiperwasser and Yoav Goldberg. 2016. Simple and accurate dependency parsing using bidirectional LSTM feature representations. *TACL* 4:313–327. <https://transacl.org/ojs/index.php/tacl/article/view/885>.
- Lingpeng Kong, Alexander M. Rush, and Noah A. Smith. 2015. Transforming dependencies into phrase structures. In *Proc. of NAACL HLT*. <https://aclweb.org/anthology/N15-1080>.
- Marco Kuhlmann and Stephan Oepen. 2016. Towards a catalogue of linguistic graph banks. *Computational Linguistics* <https://mn.uio.no/ifi/english/people/aca/oe/cl.pdf>.
- Fei Liu, Jeffrey Flanigan, Sam Thomson, Norman Sadeh, and Noah A. Smith. 2015. Toward abstractive summarization using semantic representations. In *Proc. of NAACL*. pages 1077–1086. <http://aclweb.org/anthology/N15-1114>.
- Wolfgang Maier. 2015. Discontinuous incremental shift-reduce parsing. In *Proc. of ACL*. pages 1202–1212. <http://aclweb.org/anthology/P15-1116>.
- Wolfgang Maier and Timm Lichte. 2016. Discontinuous parsing with continuous trees. In *Proc. of Workshop on Discontinuous Structures in NLP*. pages 47–57. <http://aclweb.org/anthology/W16-0906>.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *CoRR* abs/1301.3781. <https://arxiv.org/pdf/1301.3781>.
- Dipendra K Misra and Yoav Artzi. 2016. Neural shift-reduce CCG semantic parsing. In *Proc. of EMNLP*. pages 1775–1786. <http://aclweb.org/anthology/D16-1183>.

- Shashi Narayan and Claire Gardent. 2014. [Hybrid simplification using deep semantics and machine translation](#). In *Proc. of ACL*. pages 435–445. <http://aclweb.org/anthology/P14-1041>.
- Graham Neubig, Chris Dyer, Yoav Goldberg, Austin Matthews, Waleed Ammar, Antonios Anastasopoulos, Miguel Ballesteros, David Chiang, Daniel Clothiaux, Trevor Cohn, Kevin Duh, Manaal Faruqui, Cynthia Gan, Dan Garrette, Yangfeng Ji, Lingpeng Kong, Adhiguna Kuncoro, Gaurav Kumar, Chaitanya Malaviya, Paul Michel, Yusuke Oda, Matthew Richardson, Naomi Saphra, Swabha Swayamdipta, and Pengcheng Yin. 2017. [DyNet: The dynamic neural network toolkit](#). *arXiv preprint arXiv:1701.03980* <https://arxiv.org/abs/1701.03980>.
- Joakim Nivre. 2003. [An efficient algorithm for projective dependency parsing](#). In *Proc. of IWPT*. pages 149–160. <http://aclweb.org/anthology/W06-2933>.
- Joakim Nivre. 2009. [Non-projective dependency parsing in expected linear time](#). In *Proc. of ACL*. pages 351–359. <http://aclweb.org/anthology/P09-1040>.
- Joakim Nivre, Johan Hall, Jens Nilsson, Atanas Chanev, Gülsen Eryigit, Sandra Kübler, Svetoslav Marinov, and Erwin Marsi. 2007. [MaltParser: A language-independent system for data-driven dependency parsing](#). *Natural Language Engineering* 13(02):95–135.
- 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 *Proc. of SemEval*. pages 915–926. <http://aclweb.org/anthology/S15-2153>.
- Stephan Open, Marco Kuhlmann, Yusuke Miyao, Daniel Zeman, Dan Flickinger, Jan Hajič, Angelina Ivanova, and Yi Zhang. 2014. [SemEval 2014 task 8: Broad-coverage semantic dependency parsing](#). In *Proc. of SemEval*. pages 63–72. <http://aclweb.org/anthology/S14-2008>.
- Carl Pollard and Ivan Sag. 1994. *Head Driven Phrase Structure Grammar*. CSLI Publications, Stanford, CA.
- Nima Pourdamghani, Yang Gao, Ulf Hermjakob, and Kevin Knight. 2014. [Aligning English strings with abstract meaning representation graphs](#). In *Proc. of EMNLP*. pages 425–429. <http://aclweb.org/anthology/D14-1048>.
- Michael Pust, Ulf Hermjakob, Kevin Knight, Daniel Marcu, and Jonathan May. 2015. [Parsing English into abstract meaning representation using syntax-based machine translation](#). In *Proc. of EMNLP*. pages 1143–1154. <http://aclweb.org/anthology/D15-1136>.
- Corentin Ribeyre, Eric Villemonte de la Clergerie, and Djamel Seddah. 2014. [Alpage: Transition-based semantic graph parsing with syntactic features](#). In *Proc. of SemEval*. pages 97–103. <http://aclweb.org/anthology/S14-2012>.
- Kenji Sagae and Alon Lavie. 2005. [A classifier-based parser with linear run-time complexity](#). In *Proc. of IWPT*. pages 125–132. <http://aclweb.org/anthology/W05-1513>.
- Kenji Sagae and Jun’ichi Tsujii. 2008. [Shift-reduce dependency DAG parsing](#). In *Proc. of COLING*. pages 753–760. <http://aclweb.org/anthology/C08-1095>.
- Natalie Schluter, Anders Søgaard, Jakob Elming, Dirk Hovy, Barbara Plank, Héctor Martínez Alonso, Anders Johanssen, and Sigrid Klerke. 2014. [Copenhagen-Malmö: Tree approximations of semantic parsing problems](#). In *Proc. of SemEval*. pages 213–217. <http://aclweb.org/anthology/S14-2034>.
- Nathan Schneider, Emily Danchik, Chris Dyer, and Noah A Smith. 2014. [Discriminative lexical semantic segmentation with gaps: running the MWE gamut](#). *TACL* 2:193–206. <http://aclweb.org/anthology/Q14-1016.pdf>.
- Mark Steedman. 2000. *The Syntactic Process*. MIT Press, Cambridge, MA.
- Elior Sulem, Omri Abend, and Ari Rappoport. 2015. [Conceptual annotations preserve structure across translations: A French-English case study](#). In *Proc. of S2MT*. pages 11–22. <http://aclweb.org/anthology/W15-3502>.
- Swabha Swayamdipta, Miguel Ballesteros, Chris Dyer, and Noah A. Smith. 2016. [Greedy, joint syntactic-semantic parsing with stack LSTMs](#). In *Proc. of CoNLL*. pages 187–197. <http://aclweb.org/anthology/K16-1019>.
- Sam Thomson, Brendan O’Connor, Jeffrey Flanagan, David Bamman, Jesse Dodge, Swabha Swayamdipta, Nathan Schneider, Chris Dyer, and Noah A. Smith. 2014. [CMU: Arcfactored, discriminative semantic dependency parsing](#). In *Proc. of SemEval*. pages 176–180. <http://aclweb.org/anthology/S14-2027>.
- Alper Tokgöz and Gülsen Eryigit. 2015. [Transition-based dependency DAG parsing using dynamic oracles](#). In *Proc. of ACL Student Research Workshop*. pages 22–27. <http://aclweb.org/anthology/P15-3004>.
- Lucy Vanderwende, Arul Menezes, and Chris Quirk. 2015. [An AMR parser for English, French, German, Spanish and Japanese and a new AMR-annotated corpus](#). In *Proc. of NAACL*. pages 26–30. <http://aclweb.org/anthology/N15-3006>.

- Chuan Wang, Sameer Pradhan, Xiaoman Pan, Heng Ji, and Nianwen Xue. 2016. CAMR at SemEval-2016 task 8: An extended transition-based amr parser. In *Proc. of SemEval*. pages 1173–1178. <http://aclweb.org/anthology/S16-1181>.
- Chuan Wang, Nianwen Xue, and Sameer Pradhan. 2015a. Boosting transition-based AMR parsing with refined actions and auxiliary analyzers. In *Proc. of ACL*. pages 857–862. <http://aclweb.org/anthology/P15-2141>.
- Chuan Wang, Nianwen Xue, and Sameer Pradhan. 2015b. A transition-based algorithm for AMR parsing. In *Proc. of NAACL*. pages 366–375. <http://aclweb.org/anthology/N15-1040>.
- Ralph Weischedel, Martha Palmer, Mitchell Marcus, Eduard Hovy, Sameer Pradhan, Lance Ramshaw, Nianwen Xue, Ann Taylor, Jeff Kaufman, Michelle Franchini, et al. 2013. OntoNotes release 5.0 LDC2013T19. *Linguistic Data Consortium, Philadelphia, PA* <https://catalog.ldc.upenn.edu/LDC2013T19>.
- Keenon Werling, Gabor Angeli, and Christopher D. Manning. 2015. Robust subgraph generation improves abstract meaning representation parsing. In *Proc. of ACL*. pages 982–991. <http://aclweb.org/anthology/P15-1095>.
- Yue Zhang and Stephen Clark. 2009. Transition-based parsing of the Chinese treebank using a global discriminative model. In *Proc. of IWPT*. Association for Computational Linguistics, pages 162–171. <http://aclweb.org/anthology/W09-3825>.
- Yue Zhang and Stephen Clark. 2011. Shift-reduce CCG parsing. In *Proc. of ACL*. pages 683–692. <http://aclweb.org/anthology/P11-1069>.
- Junsheng Zhou, Feiyu Xu, Hans Uszkoreit, Weiguang Qu, Ran Li, and Yanhui Gu. 2016. AMR parsing with an incremental joint model. In *Proc. of EMNLP*. pages 680–689. <http://aclweb.org/anthology/D16-1065>.
- Muhua Zhu, Yue Zhang, Wenliang Chen, Min Zhang, and Jingbo Zhu. 2013. Fast and accurate shift-reduce constituent parsing. In *Proc. of ACL*. pages 434–443. <http://aclweb.org/anthology/P13-1043>.