

Full neural TAG parsing for French, English, and German LTAGs

Tatiana Bladier & Jakub Waszczuk

TreeGraSP Workshop

15 May 2019


Heinrich-Heine-Universität Düsseldorf

What is this talk about?

- ★ Data-driven full parsing for English, French, (and German) with Lexicalized Tree Adjoining Grammars (LTAG; Joshi and Schabes, 1997).
- ★ LTAG parsing = supertagging (almost parsing) + actual parsing (Bangalore and Joshi, 1999).
- ★ **Today we are presenting:**
 - Combination of the LTAG dependency parsing (Kasai et al. (2018)) and A* parsing algorithm (Waszczuk (2017)).
 - New English LTAG with sister-adjunction and regular adjunction.
 - Some issues we encountered: attachment ambiguities, extra nodes resulting during LTAG extraction.

Lexicalized Tree-Adjoining Grammar (LTAG)

- ★ LTAG: Linguistically motivated grammar formalism which supports an extended domain of locality (EDL).
- ★ An LTAG consists of a set of *elementary trees* which can be combined via *substitution* and *adjunction*.
- ★ Elementary trees are associated with lexical items and represent the span over which these items can specify their syntactic or semantic constraints.




LTAG Induction from Treebanks: Statistics

- The number of distinct LTAG supertags extracted from different treebanks is different, but is approximately around 4000.
- Almost the half of all supertags appear just once in an LTAG-annotated corpus.

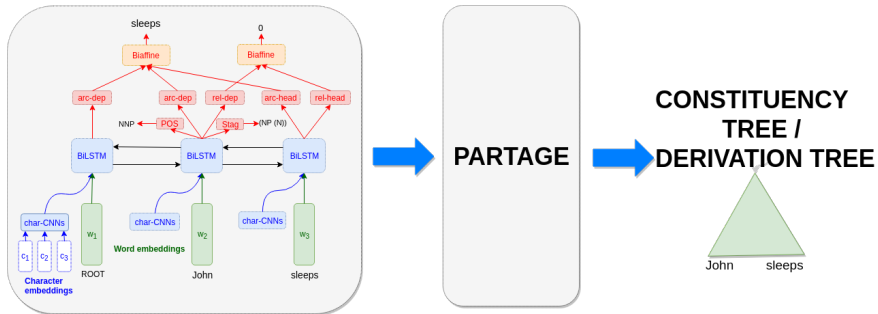
	French	German	English
	French Treebank	TiGer Treebank	Penn Treebank
Parameters	Bladier et al. (2018)	Kaeshammer (2012)	Kasai et al. (2017)
Distinct supertags	5103	3426	4727
Supertags occur. once	2611	1562	2165
POS tags	31	53	36
Sentences	21550	50000	44168
Avg. sentence length (in tokens)	29.81	17.51	23.90
Stag accuracy	82.43	88.51	89.32

LTAG induction statistics from different Treebanks (Bladier et al., 2018)

TAG parsing as dependency parsing (Kasai, 2018)



Neural dependency + Partage LTAG parsing model




A* parsing with k-best supertags and n-best arcs

- ★ A full LTAG derived tree can only be produced if the model predicts mutually compatible arcs and supertags → 1-best predicted supertag and 1-best dependency tree only allow for partial parsing.

		1-best output from Kasai (2018) parser		A* parser output	
		arcs	stags	arcs	stags
1	Nous	2	(CL ◊)	2	(CL ◊)
2	prions	0	(SENT (VN (CL) (V ◊)) (NP))	0	(SENT (VN (CL (V ◊)) (NP) (PP))
3	nos	4	(NP* (D ◊))	4	(NP* (D ◊))
4	lecteurs	2	(NP (N ◊))	2	(NP (N ◊))
5	de	2	(PP (P ◊) (VPinf))	2	(PP (P ◊) (VPinf))
6	bien	7	(VPinf* (ADV ◊))	7	(VPinf* (ADV ◊))
7	vouloir	5	(VPinf (VN (V ◊)) (VPinf))	5	(VPinf (VN (V ◊)) (VPinf))
8	nous	9	(CL ◊)	9	(CL ◊)
9	excuser	7	(VPinf (VN (CL) (V ◊)))	7	(VPinf (VN (CL) (V ◊)))

NO PARSE



A* LTAG parsing: general information

Principles

- ★ Bottom-up, Earley-style chart parser for LTAGs
- ★ Grammar compression \implies A* parsing
- ★ Support for adjunction (restricted) and sister-adjunction (unrestricted)

Implementation

- ★ Code + user documentation: <https://github.com/kawu/partage>
- ★ Quite extensive set of unit tests

The A* algorithm (idea)




Figure: Goal: visit only the part of the chart (marked in green) corresponding to the best-score derivation

A* LTAG parsing: interface

Input

1	John	2:1.0	(NP (N <>)):1.0
2	eats	0:1.0	(SENT (NP) (VP (V <>))):0.6 (SENT (NP) (VP (V <>) (NP))):0.4
3	an	4:0.5 1:0.5	(NP* (D <>)):1.0
4	apple	2:0.5 0:0.5	(NP (N <>)):1.0

Output

1	John	2	(NP (N <>))
2	eats	0	(SENT (NP) (VP (V <>) (NP)))
3	an	4	(NP* (D <>))
4	apple	2	(NP (N <>))

(SENT (NP (N John)) (VP (V eats) (NP (D an) (N apple))))

A* LTAG parsing: scoring

Scoring

$$P(\text{derivation}) \propto \prod_{i=1}^n (P(\text{supertag}_i) \times P(\text{head}_i))$$

where n = sentence length

A* heuristic (idea)

Given a chart item, for each not-yet-scanned word assume:

- ★ the most probable supertag
- ★ the most probable dependency head

Parsing experiments with A* parsing algorithm (1)

- ★ Number of sentences in FTB-dev with no parse as a function of the number of n-best supertags (columns) and k-best arcs (rows).

		k-best stags (FTB, dev)					
		1	2	4	6	8	10
n-best arcs	1	968	696	509	448	410	385
	5	841	296	53	14	5	2
	10	838	276	38	7	2	0

Neural Partage parsing for English (Chen's Grammar)

- ★ 4727 unique supertags, 2165 occur once, only regular adjunction.
- ★ Only supertags occurring in the train set are considered
→ 7 % (176 out of 2413) of the test sentences have -missing- supertags

	1-best output Kasai (2018)	10-best output (our work)		10-best output (our work, mod. PTB)	
	test	test	test < 25	test	test < 25
Unlabeled Attachment Score (UAS)	93.37	93.52	94.14	93.52	94.14
Labeled Attachment Score (LAS)	-	-	-	-	-
Stagging accuracy	91.49	91.53	91.71	91.53	91.71
Sentences with 100% correct stags + arcs:	30.84	33.84	45.05	33.84	45.05
# sentences without parse	791 (35.36%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)
Exact match	-	0.0	0.0	26.77	36.93
POS accuracy	-	95.63	95.63	98.03	97.92
# sentences	2237	2237	1394	2234	1392
Labeled F-measure	-	62.28	64.04	90.45	91.53

Neural Partage parsing for English (our LTAG)

- ★ 4703 unique supertags, 2331 occur once.
- ★ Sister-adjunction and regular adjunction.




	1-best output Kasai (2018)	10-best output (our work)	
	test	test	test < 25
Unlabeled Attachment Score (UAS)	91.13	91.43	93.04
Labeled Attachment Score (LAS)	–	–	–
Staggering accuracy	87.98	88	89.36
Sentences with 100% correct stags + arcs:	23.51	27.57	40.58
# sentences without parse	1269 (52.52%)	0 (0%)	0 (0%)
Exact match	–	23.34	35.66
POS accuracy	–	97.45	97.74
# sentences	2416	2416	1370
Labeled F-measure	–	86.11	88.86

Neural Partage parsing for French (our LTAG)

- ★ 5103 unique supertags, 2611 occur once, only sister-adjunction.
- ★ F1 score is almost state of the art: 83.11 reported by Cross and Huang (2016) for FTB.

	1-best output Kasai (2018)	10-best output (our work)	
	test	test	test < 25
Unlabeled Attachment Score (UAS)	85.18	85.78	89.01
Labeled Attachment Score (LAS)	–	–	–
Stagging accuracy	82.43	81.83	84.20
Sentences with 100% correct stags + arcs:	9.39	13.37	29.13
# sentences without parse	981 (79.43%)	0(0%)	0 (0%)
Exact match	–	20.02	39.17
POS accuracy	–	96.94	97.30
# sentences	1234	1234	508
Labeled F-measure	–	82.84	87.35


Attachment ordering ambiguity



Problem with Chen's and Xia's LTAGs: extra nodes

- ★ Existing LTAG induction algorithms (Chen et al. (2006), Xia (1999)) add extra nodes to LTAG elementary trees.
 - Modified version of original treebanks.
- ★ Labeled F1 evalb score on the actual PTB gives (obviously!) bad results (around 60%).
- ★ We had to extract our own LTAG for English
 - Sister-adjunction and regular adjunction.
- ★ We have the same problem with German LTAG extracted by Miriam Kaeshammer Kaeshammer (2012):
 - The supertags produce a modified version of the TiGer treebank.
 - The supertags used by Westburg (2016) cannot be put together at all.

Bottom-up approach (Chen et al., 2006)




Issues with German LTAGs

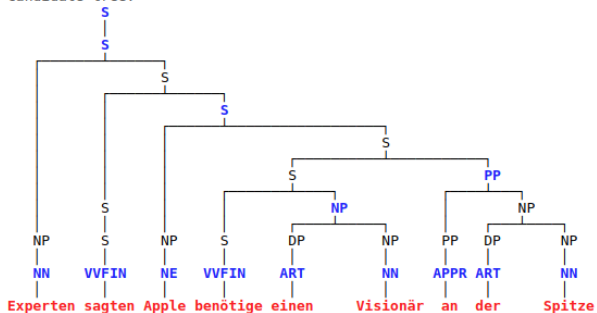
	token	head	stag Westburg (2016)	stag Kaeshammer (2012)
1	Experten	2	(NP (NN ◇))	(NP (NN ◇))
2	sagten	5	(S (NP)(VVFIN ◇)(S*))	(S (NP) (S (S (VVFIN ◇)) (S*)))
3	,	2	(PONCT* ◇)	(S (S*) (PONCT ◇))
4	Apple	5	(NP (NE ◇))	(NP (NE ◇))
5	benötige	0	(S (NP)(VVFIN ◇)(NP))	(S (NP) (S (S (VVFIN ◇)) (NP)))
6	einen	7	(DP (ART ◇))	(DP (ART ◇))
7	Visionär	5	(NP (DP)(NN ◇))	(NP (DP) (NP (NN ◇)))
8	an	5	(PP* (APPR ◇)(NP))	(S (S*) (PP (PP (APPR ◇)) (NP)))
9	der	10	(DP (ART ◇))	(DP (ART ◇))
10	Spitze	8	(NP (DP)(NN ◇))	(NP (DP) (NP (NN ◇)))
11	.	2	(PONCT* ◇)	(S (S*) (PONCT ◇))

Example from German LTAG

Sentence: Experten sagten Apple benötige einen Visionär an der Spitze
 Gold tree:



Candidate tree:



A* parsing with golden supertags and arcs (English)

- ★ English, our LTAG.
- ★ 4703 unique supertags, 2331 occur once.
- ★ Sister-adjunction and regular adjunction.

	gold supertags and arcs (our work)	
	test	test < 25
Unlabeled Attachment Score (UAS)	100	100
Labeled Attachment Score (LAS)	-	-
Stagging accuracy	100	100
Sentences with 100% correct stags + arcs:	100	100
# sentences without parse	0 (0%)	0 (0%)
Exact match	50.29	66.46
POS accuracy	100.00	100.00
Labeled Recall	95.94	97.10
Labeled Precision	95.05	96.38
Labeled F-measure	95.49	96.74
# sentences	2416	1464

A* parsing with golden supertags and arcs (French)

- ★ French, our LTAG.
- ★ 5103 unique supertags, 2611 occur once.
- ★ Sister-adjunction only.

	gold supertags and arcs (our work)	
	test	test < 25
Unlabeled Attachment Score (UAS)	100	100
Labeled Attachment Score (LAS)	-	-
Stagging accuracy	100	100
Sentences with 100% correct stags + arcs:	100	100
# sentences without parse	0 (0%)	0 (0%)
Exact match	90.52	95.87
POS accuracy	100.00	100.00
Labeled Recall	99.50	99.61
Labeled Precision	99.50	99.61
Labeled F-measure	99.50	99.61
# sentences	1234	508

Neural Partage: Web Interface

<https://tagstar.phil.hhu.de/parser/>


Neural ParTAGe

Enter a sentence in English, German, or French. The sentence will be parsed and the most probable parse tree will be shown ([show technical details](#)).


Alternatively, upload a text file with sentences which should be parsed.

Sentence:

Enter a sentence in English, German, or French.



(show/hide supertags)



Neural Partage: Web Interface

`https://tagstar.phil.hhu.de/parser/`


Neural ParTAGe

Enter a sentence in English, German, or French. The sentence will be parsed and the most probable parse tree will be shown ([show technical details](#)).

Alternatively, upload a text file with sentences which should be parsed.

Sentence:

Enter a sentence in English, German, or French.



[\(show/hide supertags\)](#)

[Imprint](#) [Privacy](#)

Conclusion and future work

- ★ Partage makes full neural parsing with LTAGs possible.
- ★ LTAG grammar design is crucial (attachment ambiguities, number of supertags)
 - bottleneck problem.

Our future work:

- improve LTAGs for English and French;
- improve the website;
- German LTAG.

Thank you!

THANK YOU VERY MUCH FOR YOUR ATTENTION!

References I

- Bangalore, S. and Joshi, A. K. (1999). Supertagging: An approach to almost parsing. *Computational linguistics*, 25(2):237–265.
- Bladier, T., van Cranenburgh, A., Samih, Y., and Kallmeyer, L. (2018). German and french neural supertagging experiments for Itag parsing (to appear). *ACL 2018 Student Research Workshop, Melbourne, Australia*.
- Chen, J., Bangalore, S., and Vijay-Shanker, K. (2006). Automated extraction of tree-adjointing grammars from treebanks. *Natural Language Engineering*, 12(3):251–299.
- Cross, J. and Huang, L. (2016). Span-based constituency parsing with a structure-label system and provably optimal dynamic oracles. *arXiv preprint arXiv:1612.06475*.
- Joshi, A. K. and Schabes, Y. (1997). Tree-adjointing grammars. In *Handbook of formal languages*, pages 69–123. Springer.
- Kaeshammer, M. (2012). *A German treebank and lexicon for tree-adjointing grammars*. PhD thesis, Master’s thesis, Universitat des Saarlandes, Saarlandes, Germany.
- Kasai, J., Frank, R., McCoy, T., Rambow, O., and Nasr, A. (2017). Tag parsing with neural networks and vector representations of supertags. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1713–1723.
- Kasai, J., Frank, R., Xu, P., Merrill, W., and Rambow, O. (2018). End-to-end graph-based TAG parsing with neural networks. *CoRR*, abs/1804.06610.
- Waszczuk, J. (2017). *Leveraging MWEs in practical TAG parsing: towards the best of the two world*. PhD thesis.
- Westburg, A. (2016). Supertagging for german - an implementation based on tree adjoining grammars. Master’s thesis, Heinrich Heine University of Düsseldorf.
- Xia, F. (1999). Extracting tree adjoining grammars from bracketed corpora. In *Proceedings of the 5th Natural Language Processing Pacific Rim Symposium (NLPRS-99)*, pages 398–403.