

Syntax-aware Neural Semantic Role Labeling with Supertags

Jungo Kasai^{♣*} Dan Friedman^{♣*} Robert Frank[◇]

Dragomir Radev[◇] Owen Rambow[♡]

[♣]University of Washington [♣]Google

[◇]Yale University [♡]Elemental Cognition, LLP

jkasai@cs.washington.edu danfriedman@google.com

{robert.frank, dragomir.radev}@yale.edu

owenr@elementalcognition.com

Abstract

We introduce a new syntax-aware model for dependency-based semantic role labeling that outperforms syntax-agnostic models for English and Spanish. We use a BiLSTM to tag the text with supertags extracted from dependency parses, and we feed these supertags, along with words and parts of speech, into a deep highway BiLSTM for semantic role labeling. Our model combines the strengths of earlier models that performed SRL on the basis of a full dependency parse with more recent models that use no syntactic information at all. Our local and non-ensemble model achieves state-of-the-art performance on the CoNLL 09 English and Spanish datasets. SRL models benefit from syntactic information, and we show that supertagging is a simple, powerful, and robust way to incorporate syntax into a neural SRL system.

1 Introduction

Semantic role labeling (SRL) is the task of identifying the semantic relationships between each predicate in a sentence and its arguments (Gildea and Jurafsky, 2002). While early research assumed that SRL models required syntactic information to perform well (Punyakanok et al., 2008), recent work has demonstrated that neural networks can achieve competitive and even state-of-the-art performance without any syntactic information at all (Zhou and Xu, 2015; Marcheggiani et al., 2017; He et al., 2017). These systems have the benefits of being simpler to implement and performing more robustly on foreign languages and out-of-domain data, cases where syntactic parsing is more difficult (Marcheggiani et al., 2017).

In this paper, we show that using supertags is an effective middle ground between using full syntactic parses and using no syntactic information

at all. A supertag is a linguistically rich description assigned to a lexical item. Supertags impose complex constraints on their local context, so supertagging can be thought of as “almost parsing” (Bangalore and Joshi, 1999). Supertagging has been shown to facilitate Tree-Adjoining Grammar (TAG) parsing (Bangalore et al., 2009; Friedman et al., 2017; Kasai et al., 2017, 2018) and Combinatory Categorical Grammar (CCG) parsing (Clark and Curran, 2007; Kummerfeld et al., 2010; Lewis et al., 2016; Xu, 2016).

We propose that supertags can serve as a rich source of syntactic information for downstream tasks without the need for full syntactic parsing. Following Ouchi et al. (2014), who used supertags to improve dependency parsing, we extract various forms of supertags from the dependency-annotated CoNLL 09 corpus. This contrasts with prior SRL work that uses TAG or CCG supertags (Chen and Rambow, 2003; Lewis et al., 2015). We train a bidirectional LSTM (BiLSTM) to predict supertags and feed the predicted supertag embedding, along with word and predicted part-of-speech embeddings, to another BiLSTM for semantic role labeling. Predicted supertags are represented by real-valued vectors, contrasting with approaches based on syntactic paths (Roth and Lapata, 2016; He et al., 2018) and syntactic edges (Marcheggiani and Titov, 2017; Strubell et al., 2018). This way of incorporating information alleviates the issue of error propagation from parsing.

Supertagging has many advantages as part of a natural language processing pipeline. First, as a straightforward sequence-labeling task, the supertagging architecture is much simpler than comparable systems for structured parsing. Second, it is simple to extract different forms of supertags from a dependency corpus to test different hypotheses about which kinds of syntactic information are most useful for downstream tasks. Our re-

*Work partially done at Yale University.

Token	Model 1	Model TAG
No	DEP/R	DEP/R
,	P/R	P/R
it	SBJ/R	-
was	ROOT+L_R	ROOT+SBJ/L_PRD/R
n't	ADV/L	ADV/L
black	NAME/R	NAME/R
Monday	PRD/L+L	-

Table 1: Supertags for the sentence “No, it wasn’t black Monday.”

Position	Feature	0	1	2	TAG
Obligatory Parent	Direction	•	•	•	
	Relation	•	•	•	
Optional Parent	Direction	•	•	•	•
	Relation	•	•	•	•
Obligatory Dep.	Direction		•	•	•
	Relation			•	•
Optional Dep.	Direction		•	•	

Table 2: Supertag models for SRL. Models 1 and 2 are from Ouchi et al. (2014) and Model 0 is from Nguyen and Nguyen (2016).

sults show that supertags, by encoding just enough information, can improve SRL performance even compared to systems that incorporate complete dependency parses.

2 Our Models

2.1 Supertag Design

We experiment with four supertag models, two from Ouchi et al. (2014), one from Nguyen and Nguyen (2016), and one of our own design inspired by Tree Adjoining Grammar supertags (Bangalore and Joshi, 1999). Each model encodes a different set of attributes about the syntactic relationship between a word, its parent, and its dependents. Table 2 summarizes what information is expressed in each supertag model.

Model 0. A Model 0 supertag for a word w encodes the dependency relation and the relative position (direction) between w and its head, i.e. left (L), right (R), or no direction (ROOT) (Nguyen and Nguyen, 2016).

Model 1. A Model 1 supertag for w adds to the “parent information” from Model 0 the information of whether w possesses dependents to its left (L) or right (R) (Ouchi et al., 2014).

Model 2. A Model 2 supertag for w extends Model 1 by encoding the dependency relation between w and its obligatory dependents.¹ When w

¹Following Ouchi et al. (2014), we define obligatory dependents as those with relations ‘SBJ,’ ‘OBJ,’ ‘PRD,’ and ‘VC.’ For Spanish, we define obligatory syntactic arguments

lacks such obligatory children, we encode whether it possesses non-obligatory dependents to the left (L) or right (R) as in Model 1.

Model TAG. We propose Model TAG supertags that represent syntactic information analogously to TAG supertags (elementary trees) (Bangalore and Joshi, 1999). A Model TAG supertag encodes the dependency relation and the direction of the head of a word similarly to Model 0 if the dependency relation is non-obligatory (corresponding to *adjunction* nodes), and the information about obligatory dependents of verbs if any similarly to Model 2 (corresponding to *substitution* nodes).

2.2 Supertagger Model

Motivated by recent state-of-the-art supertaggers (TAG: Kasai et al. (2017, 2018); CCG: Lewis et al. (2016); Xu (2016)), we employ a bi-directional LSTM (BiLSTM) architecture for our supertagging. The input for each word is the concatenation of a dense vector representation of the word, a vector embedding of a predicted PTB-style POS tag (only for English),² and a vector output by character-level Convolutional Neural Networks (CNNs) for morphological information.

For POS tagging before English supertagging, we use the same hyperparameters as in Ma and Hovy (2016). For supertagging, we follow the hyperparameters chosen in Kasai et al. (2018) regardless of the supertag model that is employed. We initialize the word embeddings by the pre-trained 100 dimensional GloVe (Pennington et al., 2014) and the 300 dimensional FastText (Bojanowski et al., 2017) vectors for English and Spanish respectively.

2.3 Semantic Role Labeling

Our SRL model is most similar to the syntax-agnostic SRL model proposed by Marcheggiani et al. (2017). Our model differs in two ways: 1) we add randomly initialized 50 dimensional supertag embeddings to the input layer (Fig. 1), and 2) we use a modified LSTM with highway layers and regularization (0.5 dropout) as in He et al. (2017).

We use the same hyperparameters as in Marcheggiani et al. (2017) with randomly initialized 50 dimensional embeddings for supertags.³

as ‘dc,’ ‘subj,’ ‘cd,’ and ‘cpred.’

²For the English data, predicted PTB-style POS tags generally contribute to increases, approximately 0.2-0.4% in the dev set, whereas for Spanish adding predicted (coarse-grained) POS tags hurt the performance.

³We provide lists of hyperparameters in Appendix A.1.

Supertag	English				Spanish		
	# Stags	Dev	ID	OOD	# Stags	Dev	ID
Model 0	99	92.93	94.17	88.71	88	92.97	92.67
Model 1	298	91.07	92.50	86.51	220	90.63	90.37
Model 2	692	90.60	92.05	85.40	503	90.08	89.84
Model TAG	430	92.60	94.17	87.46	317	92.33	92.18

Table 3: Supertagging accuracies for English and Spanish. ID and OOD indicate the in-domain and out-of-domain evaluation data respectively. The # Stags columns show the number of supertags in the corresponding training set.

For pre-trained word embeddings, we use the same word embeddings as the ones in Marcheggiani et al. (2017) for English and the 300-dimensional FastText vectors (Bojanowski et al., 2017) for Spanish. We use the predicates predicted by the mate-tools (Björkelund et al., 2009) (English) and Zhao et al. (2009) (Spanish) system in our models, again following Marcheggiani et al. (2017) to facilitate comparison. Our code is available online for easy replication of our results.⁴

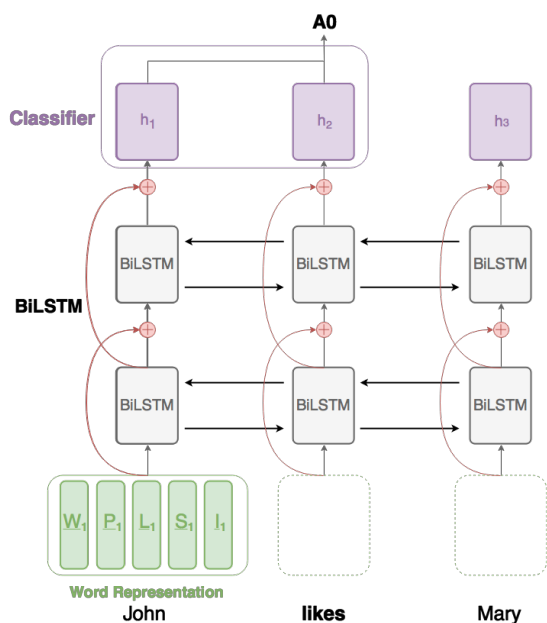


Figure 1: SRL architecture with a highway BiLSTM. W_1, P_1, L_1, S_1, I_1 indicate the word, POS, lemma, supertag, and predicate indicator embeddings for the first token, *John*. Here we only show two layers.

3 Results and Discussion

Table 3 provides our supertagging results for English and Spanish across the different types of supertag described above. Here we clearly see the general pattern that the more granular supertagging becomes, the less reliable it is, and finding the balance between granularity and predictability is critical. We present our SRL results in Tables 4-7 along with the results from a baseline

⁴https://github.com/jungokasai/stagging_srl.

BiLSTM model, which is our implementation of the syntax-agnostic model in Marcheggiani et al. (2017). We also present results for a BiLSTM model with dropout and highway connections but without supertags (BDH model), to distinguish the effects of supertags from the effects of better LSTM regularization. In every experiment we train the model five times, and present the mean score. Table 4 shows that Model 1 yields the best performance in the English dev set, and thus we only use Model 1 supertags for test evaluation. We primarily show results only with word type embeddings to conduct fair comparisons with prior work, but we also provide results with deep contextual word representations, ELMo (Peters et al., 2018), and compare our results with recent work that utilizes ELMo (He et al., 2018).⁵

English in-domain. Table 5 summarizes the results on the English in-domain test set. First, we were able to approximately replicate the results from Marcheggiani et al. (2017). Adding dropout and highway connections to our BiLSTM model improves performance by 0.5 points, to 88.1, and adding supertags improves results even further to 88.6. Our supertag model performs even better than the non-ensemble model in Marcheggiani and Titov (2017), in which the model is given the complete dependency parse of the sentence. This result suggests that supertags can be even more effective for SRL than a more complete representation of syntax. Furthermore, our supertag-based method with contextual representations achieves 90.2, a new state-of-the-art. Interestingly, the gain from supertagging decreases to 0.2 points (90.2 vs. 90.0) in the presence of contextual representations, suggesting that contextual representations encode some of the same syntactic information that supertags provide.

English out-of-domain. One of the advantages of using a syntax-agnostic SRL model is that such a model can perform relatively well on out-of-domain data, where the increased difficulty of syn-

⁵We used the pretrained ELMo available at <https://tfhub.dev/google/elmo/2>.

Architecture	P	R	F ₁
BiLSTM	87.27	85.16	86.20
BiLSTM + DOut	86.49	86.11	86.30
BiLSTM + DOut + HWay	86.97	86.43	86.70
BDH + Model 0	87.47	86.46	86.96
BDH + Model 1	87.69	86.72	87.20
BDH + Model 2	87.54	86.09	86.81
BDH + Model TAG	87.78	86.07	86.92

Table 4: Results on the CoNLL 2009 dev set for English. BDH stands for BiLSTM + Dropout + Highway.

tactic parsing can cause errors in a syntax-based system (Marcheggiani et al., 2017). Unfortunately we were not able to replicate the out-of-domain results of Marcheggiani et al. (2017): our implementation of the BiLSTM achieves a score of 76.4, compared to their reported score of 77.7. However, we note that incorporating supertags into our own model improves performance, with our best model achieving a score of 77.6. Our supertag-based model also substantially outperforms the full dependency-based models (Roth and Lapata, 2016; Marcheggiani and Titov, 2017). This suggests that syntax with a certain degree of granularity is useful even across domains. Our supertag-based method alleviates the issue of error propagation from syntactic parsing. Finally, our model with contextual representations yields 80.8, an improvement of 1.5 F1 points over the previous state-of-the-art (He et al., 2018), which also uses ELMo.

Spanish. Table 7 shows the results on the Spanish test data. Our BiLSTM implementation yields lower performance than Marcheggiani et al. (2017): our model achieves a score of 79.1, compared to their reported score of 80.3. However, our BDH model yields a score of 80.8, already achieving state-of-the-art performance. Adding supertags to BDH improves the score further to 81.0. This suggests that while the gains are relatively small, the supertag-based approach still helps Spanish SRL. Supertags slightly improve performance when contextual representations are used (83.0 vs. 82.9). See appendices for details.

Following the analysis in Roth and Lapata (2016), we show plots of the BiLSTM, BDH (BiLSTM + Dropout + Highway), and Model 1 role labeling performance for sentences with varying number of words (in-domain: Fig. 2; out-of-domain: Fig. 3). Note first that BDH outperforms the baseline BiLSTM model in a relatively uniform manner across varying sentence lengths. The benefits of Model 1 supertags, in contrast, come more from longer sentences, especially in the out-

Non-ensemble System	P	R	F ₁
FitzGerald et al. (2015)	–	–	87.3
Roth and Lapata (2016)	90.0	85.5	87.7
Marcheggiani et al. (2017)	88.7	86.8	87.7
Marcheggiani and Titov (2017)	89.1	86.8	88.0
BiLSTM	88.5	86.7	87.6
BDH	88.3	87.8	88.1
BDH + Model 1	89.0	88.2	88.6
+ Contextual Representations			
He et al. (2018) (ELMo)	89.7	89.3	89.5
BDH + ELMo	90.3	89.7	90.0
BDH + Model 1 + ELMo	90.3	90.0	90.2
Ensemble System			
FitzGerald et al. (2015)	–	–	87.7
Roth and Lapata (2016)	90.3	85.7	87.9
Marcheggiani and Titov (2017)	90.5	87.7	89.1

Table 5: Results on the CoNLL 2009 in-domain test set for English. All standard deviations in $F_1 < 0.12$.

Non-ensemble System	P	R	F ₁
FitzGerald et al. (2015)	–	–	75.2
Roth and Lapata (2016)	76.9	73.8	75.3
Marcheggiani et al. (2017)	79.4	76.2	77.7
Marcheggiani and Titov (2017)	78.5	75.9	77.2
BiLSTM	77.2	75.6	76.4
BDH	77.8	76.6	77.2
BDH + Model 1	78.0	77.2	77.6
+ Contextual Representations			
He et al. (2018) (ELMo)	81.9	76.9	79.3
BDH + ELMo	81.1	80.4	80.8
BDH + Model 1 + ELMo	81.0	80.5	80.8
Ensemble System			
FitzGerald et al. (2015)	–	–	75.5
Roth and Lapata (2016)	79.7	73.6	76.5
Marcheggiani and Titov (2017)	80.8	77.1	78.9

Table 6: Results on the CoNLL 2009 out-of-domain test set for English. The standard deviation in F_1 ranges between 0.2 and 0.35.

System	P	R	F ₁
Zhao et al. (2009)	83.1	78.0	80.5
Roth and Lapata (2016)	83.2	77.4	80.2
Marcheggiani et al. (2017)	81.4	79.3	80.3
BiLSTM	79.8	78.4	79.1
BDH	82.0	79.7	80.8
BDH + Model 1	81.9	80.2	81.0
BDH + ELMo	83.1	82.8	82.9
BDH + Model 1 + ELMo	83.1	83.0	83.0

Table 7: Results on the CoNLL 2009 test set for Spanish. All standard deviations in $F_1 < 0.1$.




Figure 2: In-domain test results by sentence length.

Model	V/A0			V/A1			V/A2			V/AM		
	P	R	F	P	R	F	P	R	F	P	R	F
Mate-tools	91.2	87.4	89.3	91.0	90.8	90.9	82.8	76.9	79.7	79.3	74.4	76.8
Path-LSTM	90.8	89.2	90.0	91.0	91.9	91.4	84.3	76.9	80.4	82.2	72.4	77.0
BiLSTM	91.1	89.7	90.4	92.1	90.9	91.5	84.0	75.0	79.2	77.7	76.9	77.3
BDH	90.9	90.8	90.9	91.5	92.4	92.0	80.3	76.1	78.1	79.6	79.1	79.3
Model 0	92.3	92.2	92.3	93.4	92.7	93.0	81.9	77.8	79.8	79.1	79.5	79.3
Model 1	92.5	91.6	92.0	93.0	92.8	92.9	80.9	80.3	80.6	80.1	78.6	79.4
Model 2	91.9	90.1	91.0	92.5	92.4	92.4	79.2	77.8	78.5	79.9	78.2	79.1
TAG	91.7	89.9	90.8	92.5	93.3	92.9	82.1	77.3	79.6	80.4	78.3	79.3

Model	N/A0			N/A1			N/A2			N/AM		
	P	R	F	P	R	F	P	R	F	P	R	F
Mate-Tools	86.1	74.9	80.2	84.9	82.2	83.5	81.4	74.7	77.9	78.6	72.0	75.2
Path-LSTM	86.9	78.2	82.3	87.5	84.4	85.9	82.4	76.8	79.5	79.5	69.2	74.0
BiLSTM	85.1	79.5	82.2	85.8	83.4	84.6	81.0	76.4	78.7	72.8	71.4	72.1
BDH	83.9	80.1	82.0	84.8	86.1	85.5	80.6	77.0	78.8	71.3	77.9	74.5
Model 0	87.2	77.6	82.1	86.2	85.4	85.8	79.9	79.2	79.5	69.4	79.4	74.0
Model 1	84.1	80.7	82.4	85.2	86.0	85.6	79.6	79.6	79.6	75.2	76.3	75.8
Model 2	86.0	79.5	82.6	85.4	85.5	85.5	80.5	77.3	78.9	73.3	76.1	74.6
TAG	83.9	79.8	81.8	84.9	86.3	85.6	81.5	75.8	78.6	72.3	72.9	72.6

Table 8: English in-domain test results by predicate category and role label. The mate-tools (Björkelund et al., 2009) and Path-LSTM results are taken from Roth and Lapata (2016).




Figure 3: Out-of-domain results by sentence length.

of-domain test set. This implies that the supertag model is robust to the sentence length, probably because supertags encode relations between words that are linearly distant in the sentence, information that a simple BiLSTM is unlikely to recover.

Table 8 reports SRL results broken down by predicate category (V: Verb, Propbank; N: Noun, Nombank) and semantic role. We can observe that the various supertag models differ in their performance for different predicate-role pairs, suggesting that different kinds of linguistic information are relevant for identifying the different roles. Overall, Model 1 supertags achieve the most consistent improvements over BiLSTM and BiLSTM + Dropout + Highway (BDH) in V/A0, V/A1, V/A2, V/AM, N/A2, and N/AM. Moreover, Model 1 even improves on Path-LSTM (Roth and Lapata, 2016) by large margins in V/A0, V/A1, V/AM, and N/AM, even though the Path-LSTM model has the benefit of using the complete dependency path between each word and

its head. This shows that supertags can be even more effective for SRL than more granular syntactic information—even quite simple supertags, like Model 0, which encode only the dependency arc between a word and its head.

4 Conclusion and Future Work

We presented state-of-the-art SRL systems on the CoNLL 2009 English and Spanish data that make crucial use of dependency-based supertags. We showed that supertagging serves as an effective middle ground between syntax-agnostic approaches and full parse-based approaches for dependency-based semantic role labeling. Supertags give useful syntactic information for SRL and allow us to build an SRL system that does not depend on a complex architecture. We have also seen that the choice of the linguistic content of a supertag makes a significant difference in its utility for SRL. In this work, all models are developed independently for English and Spanish. However, sharing some part of SRL models could improve performance (Mulcaire et al., 2018, 2019). In future work, we will explore crosslingual transfer for supertagging and semantic role labeling.

Acknowledgments

The authors thank Diego Marcheggiani for assistance in implementing SRL models and Diego Marcheggiani and the anonymous reviewers for their helpful feedback. This work was funded in part by the Funai Overseas Scholarship to JK.

References

- Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2015. [TensorFlow: Large-scale machine learning on heterogeneous systems](#). Software available from tensorflow.org.
- Srinivas Bangalore, Pierre Boullier, Alexis Nasr, Owen Rambow, and Benoît Sagot. 2009. [MICA: A probabilistic dependency parser based on tree insertion grammars \(application note\)](#). In *Proc. of NAACL-HLT*, pages 185–188, Boulder, Colorado. Association for Computational Linguistics.
- Srinivas Bangalore and Aravind K Joshi. 1999. Supertagging: An approach to almost parsing. *Computational Linguistics*, 25(2):237–265.
- Anders Björkelund, Love Hafdell, and Pierre Nugues. 2009. [Multilingual semantic role labeling](#). In *Proc. of CoNLL*, pages 43–48, Boulder, Colorado. Association for Computational Linguistics.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. [Enriching word vectors with subword information](#). *TACL*, 5:135–146.
- John Chen and Owen Rambow. 2003. [Use of deep linguistic features for the recognition and labeling of semantic arguments](#). In *Proc. of EMNLP*, pages 41–48. Association for Computational Linguistics.
- Stephen Clark and James R. Curran. 2007. Wide-coverage efficient statistical parsing with CCG and log-linear models. *Computational Linguistics*, 33:493–552.
- John C. Duchi, Elad Hazan, and Yoram Singer. 2011. [Adaptive subgradient methods for online learning and stochastic optimization](#). *JMLR*, 12:2121–2159.
- Nicholas FitzGerald, Oscar Täckström, Kuzman Ganchev, and Dipanjan Das. 2015. [Semantic role labeling with neural network factors](#). In *Proc. of EMNLP*, pages 960–970, Lisbon, Portugal. Association for Computational Linguistics.
- Dan Friedman, Jungo Kasai, R. Thomas McCoy, Robert Frank, Forrest Davis, and Owen Rambow. 2017. [Linguistically rich vector representations of supertags for TAG parsing](#). In *Proc. of TAG+*, pages 122–131, Umeå, Sweden. Association for Computational Linguistics.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson Liu, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer. 2018. [AllenNLP: A deep semantic natural language processing platform](#). arXiv:1803.07640.
- Daniel Gildea and Daniel Jurafsky. 2002. Automatic labeling of semantic roles. *Computational Linguistics*, 28(3):245–288.
- Filip Ginter, Jan Hajič, Juhani Luotolahti, Milan Straka, and Daniel Zeman. 2017. [CoNLL 2017 shared task - automatically annotated raw texts and word embeddings](#). LINDAT/CLARIN digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.
- Luheng He, Kenton Lee, Mike Lewis, and Luke Zettlemoyer. 2017. [Deep semantic role labeling: What works and what’s next](#). In *Proc. of ACL*, pages 473–483, Vancouver, Canada. Association for Computational Linguistics.
- Shexia He, Zuchao Li, Hai Zhao, and Hongxiao Bai. 2018. [Syntax for semantic role labeling, to be, or not to be](#). In *Proc. of ACL*, pages 2061–2071, Melbourne, Australia. Association for Computational Linguistics.
- Jungo Kasai, Robert Frank, R. Thomas McCoy, Owen Rambow, and Alexis Nasr. 2017. [TAG parsing with neural networks and vector representations of supertags](#). In *Proc. of EMNLP*, pages 1713–1723, Copenhagen, Denmark. Association for Computational Linguistics.
- Jungo Kasai, Robert Frank, Pauli Xu, William Merrill, and Owen Rambow. 2018. [End-to-end graph-based TAG parsing with neural networks](#). In *Proc. of NAACL-HLT*, pages 1181–1194, New Orleans, Louisiana. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Lei Ba. 2015. [ADAM: A Method for Stochastic Optimization](#). In *Proc. of ICLR*, San Diego, California.
- Jonathan K. Kummerfeld, Jessika Roesner, Tim Dawborn, James Haggerty, James R. Curran, and Stephen Clark. 2010. [Faster parsing by supertagger adaptation](#). In *Proc. of ACL*, pages 345–355, Uppsala, Sweden. Association for Computational Linguistics.
- Mike Lewis, Luheng He, and Luke Zettlemoyer. 2015. [Joint A* CCG parsing and semantic role labelling](#). In *Proc. of EMNLP*, pages 1444–1454, Lisbon, Portugal. Association for Computational Linguistics.
- Mike Lewis, Kenton Lee, and Luke Zettlemoyer. 2016. [LSTM CCG parsing](#). In *Proc. of NAACL-HLT*, pages 221–231, San Diego, California. Association for Computational Linguistics.

- Xuezhe Ma and Eduard Hovy. 2016. [End-to-end sequence labeling via bi-directional LSTM-CNNs-CRF](#). In *Proc. of ACL*, pages 1064–1074, Berlin, Germany. Association for Computational Linguistics.
- Diego Marcheggiani, Anton Frolov, and Ivan Titov. 2017. [A simple and accurate syntax-agnostic neural model for dependency-based semantic role labeling](#). In *Proc. of CoNLL*, pages 411–420, Vancouver, Canada. Association for Computational Linguistics.
- Diego Marcheggiani and Ivan Titov. 2017. [Encoding sentences with graph convolutional networks for semantic role labeling](#). In *Proc. of EMNLP*, pages 1507–1516, Copenhagen, Denmark. Association for Computational Linguistics.
- Phoebe Mulcaire, Jungo Kasai, and Noah A. Smith. 2019. [Polyglot contextual representations improve crosslingual transfer](#). In *Proc. of NAACL-HLT*, page to appear.
- Phoebe Mulcaire, Swabha Swayamdipta, and Noah A. Smith. 2018. [Polyglot semantic role labeling](#). In *Proc. of ACL*, pages 667–672, Melbourne, Australia. Association for Computational Linguistics.
- Kiet Van Nguyen and Ngan Luu-Thuy Nguyen. 2016. Vietnamese transition-based dependency parsing with supertag features. In *Proc. of KSE*, pages 175–180, Hanoi, Vietnam. Springer International Publishing.
- Hiroki Ouchi, Kevin Duh, and Yuji Matsumoto. 2014. [Improving dependency parsers with supertags](#). In *Proc. of EACL*, pages 154–158, Gothenburg, Sweden. Association for Computational Linguistics.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proc. of EMNLP*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proc. of NAACL-HLT*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Vasin Punyakanok, Dan Roth, and Wen-tau Yih. 2008. The importance of syntactic parsing and inference in semantic role labeling. *Computational Linguistics*, 34(2):257–287.
- Michael Roth and Mirella Lapata. 2016. [Neural semantic role labeling with dependency path embeddings](#). In *Proc. of ACL*, pages 1192–1202, Berlin, Germany. Association for Computational Linguistics.
- Emma Strubell, Patrick Verga, Daniel Andor, David Weiss, and Andrew McCallum. 2018. [Linguistically-informed self-attention for semantic role labeling](#). In *Proc. of EMNLP*, pages 5027–5038, Brussels, Belgium. Association for Computational Linguistics.
- Wenduan Xu. 2016. [LSTM shift-reduce CCG parsing](#). In *Proc. of EMNLP*, pages 1754–1764, Austin, Texas. Association for Computational Linguistics.
- Hai Zhao, Wenliang Chen, Jun’ichi Kazama, Kiyotaka Uchimoto, and Kentaro Torisawa. 2009. [Multilingual dependency learning: Exploiting rich features for tagging syntactic and semantic dependencies](#). In *Proc. of CoNLL*, pages 61–66, Boulder, Colorado. Association for Computational Linguistics.
- Jie Zhou and Wei Xu. 2015. [End-to-end learning of semantic role labeling using recurrent neural networks](#). In *Proc. of ACL-IJCNLP*, pages 1127–1137, Beijing, China. Association for Computational Linguistics.

A Appendices

A.1 Hyperparameters

All of our models are implemented in TensorFlow (Abadi et al., 2015).

Supertagging We follow the hyperparameters chosen in Kasai et al. (2018). Specifically, we list the hyperparameters in Table 9 for completeness and easy replication.

d_w (English word embeddings)	100
d_w (Spanish word embeddings)	300
d_{pos} (POS embeddings)	100
Char-CNN window size	3
Char-CNN # filters	30
Char-CNN character embedding size	30
d_h (LSTM hidden states)	512
k (BiLSTM depth)	4
LSTM dropout rate	0.5
Recurrent dropout rate	0.5
Batch Size	100
Adam (Kingma and Ba, 2015) lr rate	0.01
Adam β_1	0.9
Adam β_2	0.999

Table 9: Supertagging Hyperparameters.

SRL We follow the hyperparameters of Marcheggiani et al. (2017) and add highway connections (He et al., 2017) and LSTM dropout. Concretely, we use the hyperparameters shown in Table 10.

Contextual Representations For English, we use the pretrained ELMo model available at <https://tfhub.dev/google/elmo/2>. For Spanish, we use a multilingual fork (Mulcaire




Figure 4: V / A0 case where BDH assigns A0 to *President* (blue arc) while Model 1 correctly assigns A0 to *Wollaeger* (red arc). The predicted Model 1 supertags for *President* and *Wollaeger* are NAME/R and SBJ/R+L respectively.




Figure 5: V/A1 where BDH assigns A1 to *money* (blue arc) while Model 1 correctly assigns A1 to *that* (red arc). The predicted Model 1 supertags for *that* and *money* are OBJ/L+R and SBJ/R+R respectively.




Figure 6: V/A2 case where BDH assigns AM-DIR to *loose* (blue arc) while Model 1 correctly assign A2 (red arc). The predicted supertag for *loose* is PRD/L (predicative complement). Notice that the “PRT” (particle) or “DIR” (adverbial of direction) feature is not predicted that could have misled the labeling. Interestingly, the gold parse and gold POS tag for *loose* treat it as an adverbial modifier to *turned*.




Figure 7: N/A2 case where BDH assigns A3 for the predicate *buy* to *at* (blue arc) while Model 1 correctly assigns A2 for the predicate *prices* (red arc). The predicted Model 1 supertag for *at* was NMOD/L+R, correctly resolving the PP attachment ambiguity.

d_w (English word embeddings)	100
d_w (Spanish word embeddings)	300
d_{pos} (POS embeddings)	16
d_l (lemma embeddings)	100
d_s (supertag embeddings)	50
d_h (LSTM hidden states)	512
d_r (role representation)	128
d_l' (output lemma representation)	128
k (BiLSTM depth)	4
α (word dropout)	.25
LSTM dropout rate	0.5
Batch Size	100
Adam lr rate	0.01
Adam β_1	0.9
Adam β_2	0.999

Table 10: SRL Hyperparameters

Character CNNs	
Char embedding size	16
(# Window Size, # Filters)	(1, 32), (2, 32), (3, 68), (4, 128), (5, 256), (6, 512), (7, 1024)
Activation	Relu
Word-level LSTM	
LSTM size	2048
# LSTM layers	2
LSTM projection size	256
Use skip connections	Yes
Inter-layer dropout rate	0.1
Training	
Batch size	128
Unroll steps (Window Size)	20
# Negative samples	64
# Epochs	10
Adagrad (Duchi et al., 2011) lr rate	0.2
Adagrad initial accumulator value	1.0

Table 11: Spanish Language Model Hyperparameters.

et al., 2019)⁶ of the AllenNLP library (Gardner et al., 2018), and train a language model on the pre-segmented Spanish data provided by Ginter et al. (2017).⁷ The original AllenNLP library uses a byte representation. On the other hand, the multilingual fork assigns a unique character id to each unicode character, thereby avoiding the need to recognize mutibyte representations. We follow the hyperparameters chosen in Mulcaire et al. (2019) (Table 11), and randomly sample 50 million tokens from the Spanish data for training.

A.2 Supplementary Analysis

We show examples from the dev set in Figures 4-7 where a model without supertags mislabels (dashed blue arcs) and Model 1 (red arcs) correctly labels. In all those cases, it is clear that the pre-

dicted supertags are playing a crucial role in guiding role labeling.

⁶<https://github.com/pmulcaire/rosita/>

⁷<https://lindat.mff.cuni.cz/repository/xmlui/handle/11234/1-1989>