Inducing syntactic trees from BERT representations

Rudolf Rosa and **David Mareček**
Charles University, Faculty of Mathematics and Physics
Institute of Formal and Applied Linguistics
Malostranské náměstí 25, 118 00 Prague, Czech Republic
{rosa, marecek}@ufal.mff.cuni.cz

1 Introduction

One of the traditional linguistic criteria for recognizing dependency relations in a sentence is that a head of a syntactic construction determines its syntactic category and can often replace it without any damage of syntactic correctness (Kübler et al., 2009; Lopatková et al., 2005). Mareček and Žabokrtský (2012) use similar principle for unsupervised dependency parsing. They estimate "reducibility", a property describing how easily a word may be omitted from sentence without damaging it. Their hypothesis is that reducible words are more likely to occur as leaves in dependency trees. They simply declare a word reducible if the same sentence without this word appears elsewhere in the corpus. This very sparse method showed that reducibility of adjectives and adverbs is high, whereas reducibility of verbs is quite low.

With the advance of neural-network language models, e.g. BERT (Devlin et al., 2018) or ELMo (Peters et al., 2018), there are new ways how to estimate reducibilities of words. In this paper, we use English model of BERT and explore how a deletion of one word in a sentence changes representations of other words. Our hypothesis is that removing a reducible word (e.g. an adjective) does not affect representation of other words so much as removing e.g. the main verb, which makes the sentence ungrammatical and of "high surprise" for the language model.

We estimate reducibilities of individual words and also of longer phrases (word n-grams), study their syntax-related properties, and then also use them to induce full dependency trees. A significant difference between our work and most previous works (Hewitt and Manning, 2019; Belinkov, 2018) is that we estimate the reducibilities and dependency trees directly from the models, without any training on syntactically annotated data.

2 Reducibility scores from BERT

We use the pretrained English model (BERT-Large, uncased)[1], and sentences from the development part of the EWT English treebank from Universal Dependencies 2.3 (Nivre et al., 2018); we subselect sentences containing only words included in the vocabulary of the BERT model.

We compute the *reducibility* $r_{p,s}$ of each *phrase* p (any word or continuous sequence of words) in each sentence s as the average change of BERT representations of the words in the sentence when the phrase p is removed. By a BERT representation of a word we mean the state on the last layer of the BERT encoder on the position corresponding to the word.

$$r_{p,s} = \frac{1}{|s_{-p}|} \sum_{w \in s_{-p}} ||b_{w,s} - b_{w,s_{-p}}||_2 \quad (1)$$

where s_{-p} is the sentence s with phrase p removed, $b_{w,s}$ is the BERT representation of word w in sentence s, and the distance of the BERT representations is the Euclidean distance. The $b_{w,s_{-p}}$ representations are obtained by simply running BERT on the sentence s with phrase p deleted.

3 Linguistic properties of reducibilities

In Figure 1, we show reducibility scores of all the words in our testing data and average reducibility scores for individual part-of-speech tags. To visualise how the word reducibility correlates with being a leaf in the dependency tree, we color all the leaf instances by yellow and all the non-leaf by blue. It is apparent that in the right side of the graph, the blue instances prevail.

The absolute word reducibilities are different in each sentence, but we have found that the thresh-

---

[1] https://storage.googleapis.com/BERT_models/2018_10_18/uncased_L-24_H-1024_A-16.zip
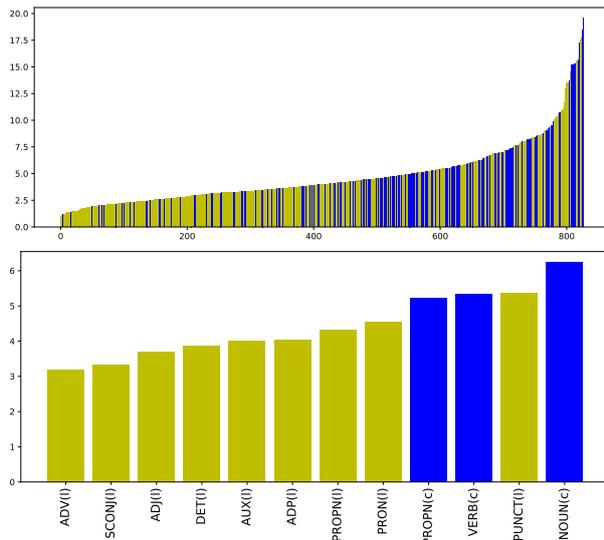
Figure 1: Distribution of reducibility scores across all the tested instances and averaged scores for POS tags. The leaf words are yellow, non-leaf words are blue.

parser	UAS
left chain baseline	6.8
right chain baseline	29.5
algorithm D	31.1
algorithm D, red. punct.	33.1
algorithm R	37.0
algorithm R, red. punct.	40.6

Table 1: Parsing results, Unlabelled Attachment Score

old separating leafs and non-leaves in a given sentence is around $1.2\times$ the average word reducibility in that sentence. This allows us to separate leafs from non-leaves with an accuracy of 74.5%, compared to the baseline of 66.4% (assuming everything is a leaf).

The syntactic root of the sentence tends to be the least reducible word. It is so in 34% of the sentences; or in 46% if we ignore punctuation, which tends to be very irreducible. The random baseline here is 13%.

The dependency edge direction can be identified with a 70.6% accuracy, assuming that the parent should be less reducible than the child node. The right-chain baseline, assuming that parent is to the right of the child, has a 65.8% accuracy.

4 Building dependency trees

We now examine to which extent the reducibilities extracted from BERT can be used to build dependency trees. We propose two algorithms, and compare them to an uninformed baseline, which is the right chain, attaching each node to its right neighbor; the rightmost node becomes the root. For English, this is quite a strong baseline.

Algorithm D: In Algorithm D, we construct a projective dependency tree based on phrase reducibilities. We use the recursive headed brackets encoding, where each subtree is enclosed in one pair of brackets, containing subtrees (phrases in brackets) and just one head (word without brack-

ets), e.g.: *((subtree) head (subtree) (subtree))*. In each step, we greedily insert a new pair of brackets corresponding to the most reducible phrase such that the resulting structure still satisfies the following conditions: (a) brackets do not cross each other (b) each subtree has a head. Table 1 shows that the resulting structures only slightly surpass the uninformed baseline (this can be improved by explicitly setting a low reducibility for punctuation).

Algorithm R: Algorithm R directly builds upon the right-chain baseline, modifying it by introducing a constraint that the parent of each node must be less reducible than the child node; the least reducible node becomes the root. Each node is thus attached to the nearest subsequent more reducible node; or to the root if all subsequent nodes are less reducible. Table 1 shows that this outperforms the baseline by 7 (or 11) percentage points.

5 Conclusions

We examine to what extent reducibility, which underlies dependency syntactic structures, can be estimated from BERT representations. We devise a method based on measuring the differences in the representations when a word or phrase is removed from the sentence, and denoting this difference as the reducibility score of that word or phrase. We find that such scores partially correspond to the notion of reducibility in dependency trees, seeing a tendency of child nodes and leaf nodes to be more reducible than parent nodes and the root. We also show that these scores can be used in a simple parsing algorithm to construct dependency trees which are more accurate than an uninformed baseline.

Acknowledgements

References

Yonatan Belinkov. 2018. *On Internal Language Representations in Deep Learning:An Analysis of Machine Translation and Speech Recognition*. Ph.D. thesis, Massachusetts Institute of Technology.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805.

John Hewitt and Christopher D. Manning. 2019. Structural Probe for Finding Syntax in Word Representations. In *Proceedings of NAACL 2019*.

Sandra Kübler, Ryan T. McDonald, and Joakim Nivre. 2009. *Dependency Parsing*. Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers.

Markéta Lopatková, Martin Plátek, and Vladislav Kuboň. 2005. Modeling syntax of free word-order languages: Dependency analysis by reduction. In *Lecture Notes in Artificial Intelligence, Proceedings of the 8th International Conference, TSD 2005*, volume 3658 of *Lecture Notes in Computer Science*, pages 140–147, Berlin / Heidelberg. Springer.

David Mareček and Zdeněk Žabokrtský. 2012. Exploiting reducibility in unsupervised dependency parsing. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, EMNLP-CoNLL '12, pages 297–307, Stroudsburg, PA, USA. Association for Computational Linguistics.

Joakim Nivre, Mitchell Abrams, Željko Agić, Lars Ahrenberg, Lene Antonsen, Katya Aplonova, Maria Jesus Aranzabe, Gashaw Arutie, Masayuki Asahara, Luma Ateyah, Mohammed Attia, Aitziber Atutxa, Liesbeth Augustinus, Elena Badmaeva, Miguel Ballesteros, Esha Banerjee, Sebastian Bank, Verginica Barbu Mititelu, Victoria Basmov, John Bauer, Sandra Bellato, Kepa Bengoetxea, Yevgeni Berzak, Irshad Ahmad Bhat, Riyaz Ahmad Bhat, Erica Biagetti, Eckhard Bick, Rogier Blokland, Victoria Bobicev, Carl Börstell, Cristina Bosco, Gosse Bouma, Sam Bowman, Adriane Boyd, Aljoscha Burchardt, Marie Candito, Bernard Caron, Gauthier Caron, Gülşen Cebiroğlu Eryiğit, Flavio Massimiliano Cecchini, Giuseppe G. A. Celano, Slavomír Čéplö, Savas Cetin, Fabricio Chalub, Jinho Choi, Yongseok Cho, Jayeol Chun, Silvie Cinková, Aurélie Collomb, Çağrı Çöltekin, Miriam Connor, Marine Courtin, Elizabeth Davidson, Marie-Catherine de Marneffe, Valeria de Paiva, Arantza Diaz de Ilarraza, Carly Dickerson, Peter Dirix, Kaja Dobrovoljc, Timothy Dozat, Kira Droganova, Puneet Dwivedi, Marhaba Eli, Ali Elkahky, Binyam Ephrem, Tomaž Erjavec, Aline Etienne, Richárd Farkas, Hector Fernandez Alcalde, Jennifer Foster, Cláudia Freitas, Katarína Gajdošová, Daniel Galbraith, Marcos Garcia, Moa Gärdenfors, Sebastian Garza, Kim Gerdes, Filip Ginter, Iakes Goenaga, Koldo Gojenola, Memduh Gökırmak, Yoav Goldberg, Xavier Gómez Guinovart, Berta Gonzáles Saavedra, Matias Grioni, Normunds Grūzītis, Bruno Guillaume, Céline Guillot-Barbance, Nizar Habash, Jan Hajič, Jan Hajič jr., Linh Hà Mỹ, Na-Rae Han, Kim Harris, Dag Haug, Barbora Hladká, Jaroslava Hlaváčová, Florinel Hociung, Petter Hohle, Jena Hwang, Radu Ion, Elena Irimia, Ọlájídé Ishola, Tomáš Jelínek, Anders Johannsen, Fredrik Jørgensen, Hüner Kaşıkara, Sylvain Kahane, Hiroshi Kanayama, Jenna Kanerva, Boris Katz, Tolga Kayadelen, Jessica Kenney, Václava Kettnerová, Jesse Kirchner, Kamil Kopacewicz, Natalia Kotsyba, Simon Krek, Sookyoung Kwak, Veronika Laippala, Lorenzo Lambertino, Lucia Lam, Tatiana Lando, Septina Dian Larasati, Alexei Lavrentiev, John Lee, Phng Lê H`ông, Alessandro Lenci, Saran Lertpradit, Herman Leung, Cheuk Ying Li, Josie Li, Keying Li, KyungTae Lim, Nikola Ljubešić, Olga Loginova, Olga Lyashevskaya, Teresa Lynn, Vivien Macketanz, Aibek Makazhanov, Michael Mandl, Christopher Manning, Ruli Manurung, Cătălina Mărănduc, David Mareček, Katrin Marheinecke, Héctor Martínez Alonso, André Martins, Jan Mašek, Yuji Matsumoto, Ryan McDonald, Gustavo Mendonça, Niko Miekka, Margarita Misirpashayeva, Anna Missilä, Cătălin Mititelu, Yusuke Miyao, Simonetta Montemagni, Amir More, Laura Moreno Romero, Keiko Sophie Mori, Shinsuke Mori, Bjartur Mortensen, Bohdan Moskalevskyi, Kadri Muischnek, Yugo Murawaki, Kaili Müürisep, Pinkey Nainwani, Juan Ignacio Navarro Horñiacek, Anna Nedoluzhko, Gunta Nešpore-Bērzkalne, Lng Nguyễn Thị, Huỳên Nguyễn Thị Minh, Vitaly Nikolaev, Rattima Nitisaroj, Hanna Nurmi, Stina Ojala, Adédayo Olúòkun, Mai Omura, Petya Osenova, Robert Östling, Lilja Øvrelid, Niko Partanen, Elena Pascual, Marco Passarotti, Agnieszka Patejuk, Guilherme Paulino-Passos, Siyao Peng, Cenel-Augusto Perez, Guy Perrier, Slav Petrov, Jussi Piitulainen, Emily Pitler, Barbara Plank, Thierry Poibeau, Martin Popel, Lauma Pretkalniņa, Sophie Prévost, Prokopis Prokopidis, Adam Przepiórkowski, Tiina Puolakainen, Sampo Pyysalo, Andriela Rääbis, Alexandre Rademaker, Loganathan Ramasamy, Taraka Rama, Carlos Ramisch, Vinit Ravishankar, Livy Real, Siva Reddy, Georg Rehm, Michael Rießler, Larissa Rinaldi, Laura Rituma, Luisa Rocha, Mykhailo Romanenko, Rudolf Rosa, Davide Rovati, Valentin Roca, Olga Rudina, Jack Rueter, Shoval Sadde, Benoît Sagot, Shadi Saleh, Tanja Samardžić, Stephanie Samson, Manuela Sanguinetti, Baiba Saulīte, Yanin Sawanakunanon, Nathan Schneider, Sebastian Schuster, Djamé Seddah, Wolfgang Seeker, Mojgan Seraji, Mo Shen, Atsuko Shimada, Muh Shohibussirri, Dmitry Sichinava, Natalia Silveira, Maria Simi, Radu Simionescu, Katalin Simkó, Mária Šimková, Kiril Simov, Aaron Smith, Isabela Soares-Bastos, Carolyn Spadine, Antonio Stella, Milan Straka, Jana Strnadová, Alane Suhr,

Umut Sulubacak, Zsolt Szántó, Dima Taji, Yuta Takahashi, Takaaki Tanaka, Isabelle Tellier, Trond Trosterud, Anna Trukhina, Reut Tsarfaty, Francis Tyers, Sumire Uematsu, Zdeňka Urešová, Larraitz Uria, Hans Uszkoreit, Sowmya Vajjala, Daniel van Niekerk, Gertjan van Noord, Viktor Varga, Eric Villemonte de la Clergerie, Veronika Vincze, Lars Wallin, Jing Xian Wang, Jonathan North Washington, Seyi Williams, Mats Wirén, Tsegay Woldemariam, Tak-sum Wong, Chunxiao Yan, Marat M. Yavrumyan, Zhuoran Yu, Zdeněk Žabokrtský, Amir Zeldes, Daniel Zeman, Manying Zhang, and Hanzhi Zhu. 2018. Universal dependencies 2.3. LINDAT/CLARIN digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.

Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proc. of NAACL*.