

Exploring the Syntactic Abilities of RNNs with Multi-task Learning

Émile Enguehard¹ Yoav Goldberg² Tal Linzen^{3,4}

¹Département d’informatique, ENS, PSL Research University

²Computer Science Department, Bar Ilan University

³LSCP & IJN, CNRS, EHESS and ENS, PSL Research University

⁴Department of Cognitive Science, Johns Hopkins University

{emile.enguehard,tal.linzen}@ens.fr yoav.goldberg@gmail.com

Abstract

Recent work has explored the syntactic abilities of RNNs using the subject-verb agreement task, which diagnoses sensitivity to sentence structure. RNNs performed this task well in common cases, but faltered in complex sentences (Linzen et al., 2016). We test whether these errors are due to inherent limitations of the architecture or to the relatively indirect supervision provided by most agreement dependencies in a corpus. We trained a single RNN to perform both the agreement task and an additional task, either CCG supertagging or language modeling. Multi-task training led to significantly lower error rates, in particular on complex sentences, suggesting that RNNs have the ability to evolve more sophisticated syntactic representations than shown before. We also show that easily available agreement training data can improve performance on other syntactic tasks, in particular when only a limited amount of training data is available for those tasks. The multi-task paradigm can also be leveraged to inject grammatical knowledge into language models.

1 Introduction

Recurrent neural networks (RNNs) have seen rapid adoption in natural language processing applications. Since these models are not equipped with explicit linguistic representations such as dependency parses or logical forms, new methods are needed to characterize the linguistic generalizations that they capture. One such method is drawn from behavioral psychology: the network is tested on cases that are carefully selected to be

informative as to the generalizations that the network has acquired.

Linzen et al. (2016) have recently applied this methodology to evaluate how well a trained RNN captures sentence structure, using the **agreement prediction task** (Bock and Miller, 1991; Elman, 1991). The form of an English verb often depends on its subject. Identifying the subject of a given verb of requires sensitivity to sentence structure. Consequently, testing an RNN on its ability to choose the correct form of a verb in context can shed light on the sophistication of its syntactic representations (see Section 2.1 for details).

RNNs trained specifically to perform the agreement task can achieve very good average performance on a corpus, with accuracy close to 99%. However, error rates increase substantially on complex sentences (Linzen et al., 2016, 2017), suggesting that the syntactic knowledge acquired by the RNN is imperfect. Finally, when the RNN is trained as a language model rather than specifically on the agreement task, its sensitivity to subject-verb agreement, measured as the relative probability of the grammatical and ungrammatical forms of the verb, degrades dramatically.

Are the limitations that RNNs showed in previous work inherent to their architecture, or can these limitations be mitigated by stronger supervision? We address this question using multi-task learning, where the same model is encouraged to develop representations that are simultaneously useful for multiple tasks. To provide the RNN with an incentive to develop more sophisticated representations, we trained it to perform one of two tasks: the first is combinatory categorical grammar (CCG) supertagging (Bangalore and Joshi, 1999), a sequence labeling task likely to require robust syntactic representations; the second task is language modeling.

We also investigate the inverse question: can

tasks such as supertagging benefit from joint training with the agreement task? This question is of practical interest. Large training sets for the agreement task are much easier to create than training sets for supertagging, which are based on manually parsed sentences. If the training signal from the agreement prediction task proves to be beneficial for supertagging, this could lead to improved supertagging (and therefore parsing) performance in languages in which we only have a small amount of parsed training sentences.

We found that multi-task learning, either with LM or with CCG supertagging, improved the performance of the RNN on the agreement prediction task. The benefits of combined training with supertagging can be quite large: accuracy in challenging relative clause sentences increased from 50.6% to 76.2%. This suggests that RNNs are in principle capable of acquiring much better syntactic representations than those they learned from the corpus in Linzen et al. (2016).

In the other direction, joint training on the agreement prediction task did not improve overall language model perplexity, but made the model more syntax-aware: grammatically appropriate verb forms had higher probability than grammatically inappropriate ones. When a limited amount of CCG training data was available, joint training on agreement prediction led to improved supertagging accuracy. These findings suggest that multi-task training with auxiliary syntactic tasks such as agreement prediction can lead to improved performance on standard NLP tasks.

2 Background and Related Work

2.1 Agreement Prediction

English present-tense third-person verbs agree in number with their subject: singular subjects require singular verbs (*the boy smiles*) and plural subjects require plural verbs (*the boys smile*). Subjects in English are not overtly marked, and complex sentences often have multiple subjects corresponding to different verbs. Identifying the subject of a particular verb can therefore be non-trivial in sentences that have multiple nouns:

- (1) The only championship **banners** that are currently displayed within the building **are** for national or NCAA Championships.

Determining that the subject of the verb in bold-face is *banners* rather than the singular nouns

championship and *building* requires an understanding of the structure of the sentence.

In the agreement task, the learner is given the words leading up to a verb (a “preamble”), and is instructed to predict whether that verb will take the plural or singular form. This task is modeled after a standard psycholinguistic task, which is used to study syntactic representations in humans (Bock and Miller, 1991; Franck et al., 2002; Staub, 2009; Bock and Middleton, 2011).

Any English sentence with a third-person present-tense verb can be used as a training example for this task: all we need is a tagger that can identify such verbs and determine whether they are plural or singular. As such, large amounts of training data for this task can be obtained from a corpus.

The agreement task can often be solved using simple heuristics, such as copying the number of the most recent noun. It can therefore be useful to evaluate the model using sentences in which such a heuristic would fail because one or more nouns of the opposite number from the subject intervene between the subject and the verb; such nouns “attract” the agreement away from the grammatical subject. In general, the more such attractors there are the more difficult the task is for a sequence model that does not represent syntax (we focus on sentences in which **all** of the nouns between the subject and the verb are of the opposite number from the subject):

- (2) The **number** of men is not clear. (One attractor)
- (3) The **ratio** of men to women is not clear. (Two attractors)
- (4) The **ratio** of men to women and children is not clear. (Three attractors)

2.2 CCG Supertagging

Combinatory Categorical Grammar (CCG) is a syntactic formalism that relies on a large inventory of lexical categories (Steedman, 2000). These categories are known as *supertags*, and can be thought of as a fine-grained extension of the usual part-of-speech tags. For example, intransitive verbs (*smile*), transitive verbs (*build*) and raising verbs (*seem*) all have different tags: *SNP*, (*SNP*)/*NP* and (*SNP*)/(*SNP*), respectively.

CCG parsers typically rely on a supertagging step where each word in a sentence is associated

with an appropriate tag. In fact, supertagging is almost as difficult as finding the full CCG parse of the sentence: once the supertags are determined, only a small number of parses are possible. At the same time, supertagging is simple to set up as a machine learning problem, since at each word it amounts to a straightforward classification problem (Bangalore and Joshi, 1999). RNNs have shown excellent performance on this task, at least in English (Xu et al., 2015; Lewis et al., 2016; Vaswani et al., 2016).

In contrast with the agreement task, training data for supertagging needs to be obtained from parsed sentences which require expert annotation (Hockenmaier and Steedman, 2007); the amount of training data is therefore limited even in English, and much more sparse in other languages.

2.3 Language Modeling

The goal of a language model is to learn the distribution $\hat{p}(w_j|w_1, \dots, w_{j-1})$ of the j -th word in a sentence given the $j - 1$ words preceding it. We seek to minimize the mean negative log-likelihood of all sentences $s_i = w_{i,1} \dots w_{i,n_i}$ in our data:

$$L(\hat{p}) = -\frac{1}{Z} \sum_{i=1}^N \sum_{j=1}^{n_i} \log \hat{p}(w_{i,j}|w_{i,1:j-1}) \quad (1)$$

where $Z = \sum_{i=1}^N n_i$. Language modeling performance is often quantified using the perplexity $2^{L(\hat{p})}$. The effectiveness of RNNs in language modeling, in particular LSTMs, has been demonstrated in numerous studies (Mikolov et al., 2010; Sundermeyer et al., 2012; Jozefowicz et al., 2016).

2.4 Multitask Learning

The benefits of multi-task learning in neural networks are straightforward. Neural networks often require a large amount of training data to achieve good performance on a task. Even with a significant amount of training data, the signal may be too sparse for them to pick it up given their weak inductive biases. By training a network on a simple task for which large quantities of data are available, we can encourage it to evolve representations that would help its performance on the primary task (Caruana, 1998; Bakker and Heskes, 2003). This logic has been applied to various NLP tasks, with generally encouraging results (Collobert and Weston, 2008; Hashimoto et al., 2016; Sjøgaard

and Goldberg, 2016; Martínez Alonso and Plank, 2017; Bingel and Sjøgaard, 2017).

3 Methods

3.1 Datasets

We used two training datasets. The first is the corpus of approximately 1.5 million sentences from the English Wikipedia compiled by Linzen et al. (2016). All sentences had at most 50 words and contained at least one third-person present-tense agreement dependency. Following Linzen et al. (2016), we replaced rare words by their part-of-speech tags, using the Penn Treebank tag set (Marcus et al., 1993).¹

The second data set we used is the CCG-Bank (Hockenmaier and Steedman, 2007), a CCG version of the Penn Treebank. This corpus contained 48934 English sentences, 27299 of which include a present tense third-person verb agreement dependency. A negligible number of sentences longer than 90 words were removed. We applied the traditional split where Sections 2-21 are used for training and Section 23 for testing (41294 and 2407 sentences respectively).² Out of the 1363 different supertags that occur in the corpus, we only attempted to predict the 452 supertags that occurred at least ten times; we replaced the rest (0.2% of the tokens) by a dummy value.

3.2 Model

The model in all of our experiments was a standard single-layer LSTM.³ The first layer was a vector embedding of word tokens into D -dimensional space. The second was a D -dimensional LSTM. The following layers depended on the task. For agreement, the output layers consisted of a linear layer with a one-dimensional output and a sigmoid activation; for language modeling, a linear layer with an N -dimensional output, where N is the size of the lexicon, and a softmax activation; and for supertagging, a linear layer with an S -dimensional

¹In the LM experiments, we restricted ourselves to 10000 words, amounting to 91.2% of the all occurrences. In the CCG supertagging experiments, we used those 12, 126 words that occurred more than 150 times, amounting to 92.2% of the total number of occurrences.

²For experiments using this corpus, we use 15784 words occurring at least four times, amounting to 95.9% of occurrences, and replace other words by their POS tags.

³Our code and data are available at <https://github.com/emengd/multitask-agreement>.

output, where S is the number of possible tags, followed by a softmax activation.

The language modeling loss is the mean negative log-likelihood of the data given in Equation (1); the loss for agreement is the mean binary cross-entropy of the classifier:

$$L_{\text{agr}} = -\frac{1}{|S|} \sum_{s \in S} \log(\hat{q}(\text{num}(s)|_{s:\text{vb}}))$$

where \hat{q} is the estimated distribution of verb numbers, S the set of sentences, $\text{num}(s)$ the correct verb number in s and $s:\text{vb}$ the sentence up to the verb. The loss for CCG supertagging is the mean cross-entropy of the classifiers:

$$L_{\text{ST}} = -\frac{1}{\sum_s |s|} \sum_{s \in S} \sum_{w_j \in s} \log(\hat{r}(\text{tag}(w_j)|_{s:w_j}))$$

where \hat{r} is the estimated distribution of CCG supertags, $\text{tag}(w_j)$ is the correct tag of word w_j in s , and $s:w_j$ is the sentence s up to and including w_j .

We had at most two tasks in any given experiment. We considered two separate setups for learning from those two tasks: joint training and pre-training.

Joint training: In this setup we had parallel output layers for each task. Both output layers received the shared LSTM representations as their input. We define the global loss L as follows:

$$L = \frac{1}{1+r} L_1 + \frac{r}{1+r} L_2 \quad (2)$$

where L_1 and L_2 are the losses associated with each task, and r is the weighting ratio of task 2 relative to task 1. This means that r is a hyperparameter that needs to be tuned. Note that sample averaging occurs before formula (2) is applied.

Pre-training: In this setup, we first trained the network on one of the tasks; we then used the weights learned by the network for the embedding layer and the LSTM layer as the initial weights of a new network which we then trained on the second task.

3.3 Training

All neural networks were implemented in Keras (Chollet, 2015) and Theano (Theano Development Team, 2016). We use the AdaGrad optimizer. We use batch training with batch sizes 128 for language modeling experiments and 256 for supertagging experiments on supertagging.

4 Agreement and Supertagging

For the supertagging experiments we used the full CCG corpus as well as 30% of the Wikipedia corpus for the agreement task (20% for training and 10% for testing). We trained the model for 20 epochs. The accuracy figures we report are averaged across three runs. We set the size of the network D to 500 hidden units.⁴ We ran a single pre-training experiment in each direction, as well as four joint training experiments, with the weight r of the agreement task set to 0.1, 1, 10 or 100.

We considered two baselines for the agreement task: the last noun baseline predicts the number of the verb based on the number of the most recent noun, and the majority baseline always predicts a singular verb (singular verbs are more common than plural ones in our corpus). Our baseline for supertagging was a majority baseline that predicts for each word its most common supertag.


The agreement task predicts the number of the verb based only on its left context (the preamble). We trained our supertagging model in the same setup. Since our model did not have access to the right context of a word when determining its supertag, we could not expect to compete with state-of-the-art taggers that use right-context lookahead (Xu et al., 2015) or even bidirectional RNNs that read the entire sentence from right to left (Vaswani et al., 2016; Lewis et al., 2016); we therefore did not compare our accuracy to these taggers.

4.1 Overall Results


Figure 1 shows the overall results of the experiment. Multi-task training with supertagging significantly improved overall accuracy on the agreement task (Figure 1a), either with pre-training or joint training: compared to the single-task setup, the agreement error rate decreased by up to 40% in relative terms (from 2.04% to 1.24%). Conversely, multi-task training with agreement did not improve supertagging accuracy, either in the pre-training or in the joint training regime; supertagging accuracy decreased the higher the weight of the agreement task (Figure 1b).

Comparing the two multi-task learning regimes, the pre-training setup performed about as well as the joint training setup with the optimal r . In the following supertagging experiments we dispensed with the joint training setup, which is time con-

⁴In initial experiments $D = 50$ yielded supertagging results inferior to a majority choice baseline.



(a) Agreement



(b) Supertagging


Figure 1: Overall results of supertagging + agreement multi-task training.

suming since it requires trying multiple values of r , and focused only on the pre-training setup.


4.2 Effect of Corpus Size

To further investigate the relative contribution of the two supervision signals, we conducted a series of follow-up experiments in the pre-training setup, using subsets of varying size of both corpora. We also included POS tagging as an auxiliary task to determine to what extent the full parse of the sentence (approximated by supertags) is crucial to the improvements we have seen in the agreement task. Since POS tags contain less syntactic information than CCG supertags, we expect them to be less helpful as an auxiliary task. Penn Treebank POS tags distinguish singular and plural nouns and verbs, but CCG supertags do not; to put the two tasks on equal footing we removed number information from the POS tags. We trained for 15 epochs and averaged our results over 5 runs.

The results for the agreement task are shown in Figure 2a (baseline values are always calculated over the full corpora). The figure confirms



(a) Agreement



(b) Supertagging

Figure 2: The effect of corpus size on agreement and supertagging accuracy in multi-task settings.

the beneficial effect of supertagging pre-training (note that the scale starts at 0.8, not 0.9 as in Figure 1a). This effect was amplified when we used less training data for the agreement task. Pre-training on POS tagging yielded a similar though slightly weaker effect. This suggests that much of the improvement in syntactic representations due to pre-training on supertagging can also be gained from pre-training on POS tagging.

Finally, Figure 2b shows that pre-training on the agreement task improved supertagging accuracy when we only used 10% of the CCG corpus (increase in accuracy from 73.4% to 76.3%); however, even with agreement pre-training supertagging accuracy is lower than when the model is trained on the full CCG corpus (where accuracy was 83.1%).

In summary, the data for each task can be used to supplement the data for the other, but there is a large imbalance in the amount of information provided by each task. This is not surprising given that the CCG supertagging data is much richer than the agreement data for any individual sentence. Still, we showed that the syntactic sig-




Figure 3: Agreement accuracy as a function of the number of attractors intervening between the subject and the verb, for two different subsets of the agreement corpus (90% and 1% of the corpus).

nal from the agreement prediction task can help improve parsing performance when CCG training data is sparse; this weak but widely available source of syntactic supervision may therefore have a practical use in languages with smaller treebanks than English.

4.3 Attraction Errors

Most sentences are syntactically simple and do not pose particular challenges to the models: the accuracy of the last noun baseline in Figure 1a was close to 95%. To investigate the behavior of the model on more difficult sentences, we next break down our test sentences by the number of agreement attractors (see Section 2.1).

Our results, shown in Figure 3, confirm that attractors make the agreement task more difficult, and that pre-training helps overcome this difficulty. This effect is amplified when we only use a small subset of the agreement corpus. In this scenario, the accuracy of the single-task model on sentences with four attractors is only 20.4%. Pre-training makes it possible to overcome this difficulty to a significant extent (though not entirely), increasing the accuracy to 40.1% in the case of POS tagging and 51.2% in the case of supertagging. This suggests that a network that has developed sophisticated syntactic representations can transfer its knowledge to a new syntactic task using only a moderate amount of data.

4.4 Relative Clauses

In Linzen et al. (2016), attraction errors were particularly severe when the attractor was inside a rel-




Figure 4: Accuracy on sentences from Bock and Cutting (1992). Error bars indicate standard deviation across runs.

ative clause. To gain a more precise understanding of the errors and the extent to which pre-training can mitigate them, we turn to two sets of carefully constructed sentences from the psycholinguistic literature (Linzen et al., 2017). Bock and Cutting (1992) compared preambles with prepositional phrase modifiers to closely matched relative clause modifiers:

- (5) PREPOSITIONAL: The demo tape(s) from the popular rock singer(s)...
- (6) RELATIVE: The demo tape(s) that promoted the popular rock singer(s)...

They constructed 24 such sentence pairs. Each of the sentences in each pair has four versions, with all possible combinations of the number of the subject and the attractor. We refer to them as SS for singular-singular (*tape, singer*), SP for singular-plural (*tape, singers*), and likewise PS and PP. We replaced out-of-vocabulary words with their POS, and further streamlined the materials by always using *that* as the relativizer.

We retrained the single-task and pre-trained models on 90% of the Wikipedia corpus. Like humans, neither model had any issues with SS and PP sentences, which do not have an attractor. The results for SP and PS sentences are shown in Figure 4. The comparison between prepositional and relative modifiers shows that the single-task model was much more likely to make errors when the attractor was in a relative clause (whereas humans are not sensitive to this distinction). This asymmetry was substantially mitigated, though not completely eliminated, by CCG pre-training.

Our second set of sentences was based on the experimental materials of Wagers et al. (2009). We adapted them by deleting the relativizer and creating two preambles from each sentence in the original experiment:

- (7) EMBEDDED VERB: The player(s) the coach(es)...
- (8) MAIN CLAUSE VERB: The player(s) the coach(es) like the best...

In the first preamble, the verb is expected to agree with the embedded clause subject (*the coach(es)*), whereas in the second one it is expected to agree with the main clause subject (*the player(s)*).

Figure 5 shows that both models made very few errors predicting the embedded clause verb, and more errors predicting the main clause verb. The relative improvement of the pre-trained model compared to the single-task one is more modest in these sentences, possibly because the single-task model does better to begin with on these sentences than on the Bock and Cutting (1992) ones. This in turn may be because the attractor immediately precedes the verb in Bock and Cutting (1992) but not in Wagers et al. (2009), and an immediately adjacent noun may be a stronger attractor. The Appendix contains additional figures tracking the predictions of the network as it processes a sample of sentences with relative clauses; it also illustrates the activation of particular units over the course of such a sentence.

5 Agreement and Language Modeling

We now turn our attention to the language modeling task. The previous experiments confirmed that



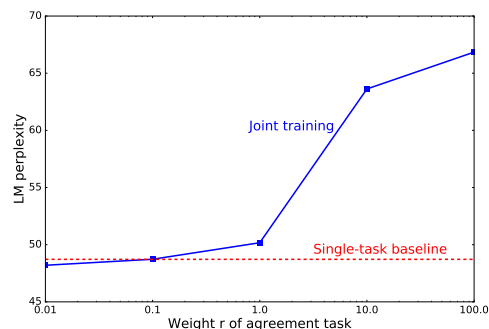


Figure 5: Accuracy on sentences based on Wagers et al. (2009). Error bars indicate standard deviation across runs.



(a) Agreement



(b) Language modeling

Figure 6: Overall results of language modeling + agreement multi-task training (trained only on sentences with an intervening noun).

agreement in sentences without attractors is easy to predict. We therefore limited ourselves in the language modeling experiments to sentences with potential attractors. Concretely, within the subset of 30% of the Wikipedia corpus, we trained our language model only on sentences with at least one noun (of any number) between the subject and the verb. There were 60680 sentences in the training set. We averaged our results over three runs. Training was stopped after 10 epochs, and the number of hidden units was set to $D = 50$.

5.1 Overall Results

The overall results are shown in Figure 6. Joint training with the LM task improves the performance of the agreement task to a significant extent, bringing accuracy up from 90.2% to 92.6% (a relative reduction of 25% in error rate). This may be due to the higher quality of the word representations that can be learned from the language modeling signal, which in turn help the model make more accurate syntactic predictions.

In the other direction, we do not obtain clear improvements in perplexity from jointly training the LM with agreement. Surprisingly, visual inspection of Figure 6b suggests that the jointly trained LM may achieve somewhat better performance than the single-task baseline for *small* values of r (that is, when the agreement task has a small effect on the overall training loss). To assess the statistical significance of this difference, we repeated the experiment with $r = 0.01$ with 20 random initializations. The standard deviation in LM loss was about 0.018, yielding a standard deviation of 0.011 for three-run averages under Gaussian assumptions. Since the difference of 0.015 between the mean LM losses of the single-task and joint training setups is of comparable magnitude, we conclude that there is no clear evidence that joint training reduces perplexity.

5.2 Grammaticality of LM Predictions

To evaluate the syntactic abilities of an RNN trained as a language model, Linzen et al. (2016) proposed to perform the agreement task by comparing the probability under the learned LM of the correct and incorrect verb forms, under the assumption that all other things being equal a grammatical sequence should have a higher probability than an ungrammatical one (Lau et al., 2016; Le Godais et al., 2017). For instance, if the sentence starts with *the dogs*, we compute:

$$\hat{p}_{\text{correct}} = \frac{\hat{p}(w_2 = \text{are} | w_{0:1} = \text{the dogs})}{\hat{p}(w_2 = \text{are} | \dots) + \hat{p}(w_2 = \text{is} | \dots)} \quad (3)$$

The prediction for the agreement task is derived by thresholding \hat{p}_{correct} at 0.5.

Is the LM learned in the joint training setup with high r more aware of subject-verb agreement than a single-task LM? Note that this is not a circular question: we are not asking whether the explicit agreement prediction output layer can perform the agreement task — that would be unsurprising — but whether joint training with this task rearranges the probability distributions that the LM defines over the entire vocabulary in a way that is more consistent with English grammar.

As the method outlined in Equation 3 may be sensitive to the idiosyncrasies of the particular verb being predicted, we also explored an unlexicalized way of performing the task. Recall that since we replace uncommon words by their POS




Figure 7: Language model agreement evaluation. Red bars indicate the results obtained on the single-task LM model, blue bars those obtained in a joint training setup with $r = 100$.

tags, POS tags are part of our lexicon. We can use this fact to compare the LM probabilities of the POS tags for the correct and incorrect verb forms: in the example of the preamble *the dogs*, the correct POS would be VBP and the incorrect one VBZ.

The results can be seen in Figure 7. The accuracy of the LM predictions from the jointly trained models is almost as high as that obtained through the agreement model itself. Conversely, the single-task model trained only on language modeling performed only slightly better than chance, and worse than our last noun baseline (recall that the dataset only included sentences with an intervening noun between the subject and the verb, though possibly of the same number as the subject). Predictions based on POS tags are somewhat worse than predictions based on the specific verb. In summary, while joint training with the explicit agreement task does not noticeably reduce language model perplexity, it does help the LM capture syntactic dependencies: the ranking of upcoming words is more consistent with the constraints of English syntax.

6 Conclusions

Previous work has shown that the syntactic representations developed by RNNs that are trained on the agreement prediction task are sufficient for the majority of sentences, but break down in more complex sentences (Linzen et al., 2016, 2017). These deficiencies could be due to fundamental limitations of the architecture, which can only be addressed by switching to more expressive archi-

tures (Socher, 2014; Grefenstette et al., 2015; Dyer et al., 2016). Alternatively, they could be due to insufficient supervision signal in the agreement prediction task, for example because relative clauses with agreement attractors are infrequent in a natural corpus.

We showed that additional supervision from pre-training on syntactic tagging tasks such as CCG supertagging can help the RNN develop more effective syntactic representations which substantially improve its performance on complex sentences, supporting the second hypothesis.

The syntactic representations developed by the RNNs were still not perfect even in the multi-task setting, suggesting that stronger inductive biases expressed as richer representational assumptions may lead to further improvements in syntactic performance. The weaker performance on complex sentences in the single-task setting indicates that the inductive bias inherent in RNNs is insufficient for learning adequate syntactic representations from unannotated strings; improvements due to a stronger inductive bias are therefore likely to be particularly pronounced in languages for which parsed corpora are small or unavailable. Finally, the strong syntactic supervision required to promote sophisticated syntactic representations in RNNs may limit their viability as models of language acquisition in children (though children may have sources of supervision that were not available to our models).

We also explored whether multi-task training with the agreement task can improve performance on more standard NLP tasks. We found that it can indeed lead to improved supertagging accuracy when there is a limited amount of training data for that task; this form of weak syntactic supervision can be used to improve parsers for low-resource languages for which only small treebanks are available.

Finally, for language modeling, multi-task training with the agreement task did not reduce perplexity, but did improve the grammaticality of the predictions of the language model (as measured by the relative ranking of grammatical and ungrammatical verb forms); such a language model that favors grammatical sentences may produce more natural-sounding text.

Acknowledgments

We thank Emmanuel Dupoux for discussion. This research was supported by the European Research Council (grant ERC-2011-AdG 295810 BOOTPHON), the Agence Nationale pour la Recherche (grants ANR-10-IDEX-0001-02 PSL and ANR-10-LABX-0087 IEC) and the Israeli Science Foundation (grant number 1555/15).


References

- Bart Bakker and Tom Heskes. 2003. Task clustering and gating for Bayesian multitask learning. *Journal of Machine Learning Research* 4:83–99.
- Srinivas Bangalore and Aravind K. Joshi. 1999. Supertagging: An approach to almost parsing. *Computational Linguistics* 25(2):237–265.
- Joachim Bingel and Anders Søgaard. 2017. Identifying beneficial task relations for multi-task learning in deep neural networks. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*. Association for Computational Linguistics, Valencia, Spain, pages 164–169.
- Kathryn Bock and J. Cooper Cutting. 1992. Regulating mental energy: Performance units in language production. *Journal of Memory and Language* 31(1):99–127.
- Kathryn Bock and Erica L. Middleton. 2011. Reaching agreement. *Natural Language & Linguistic Theory* 29(4):1033–1069.
- Kathryn Bock and Carol A. Miller. 1991. Broken agreement. *Cognitive Psychology* 23(1):45–93.
- Rich Caruana. 1998. Multitask learning. In Sebastian Thrun and Lorien Pratt, editors, *Learning to learn*, Kluwer Academic Publishers, Boston, pages 95–133.
- François Chollet. 2015. Keras. <https://github.com/fchollet/keras>.
- Ronan Collobert and Jason Weston. 2008. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proceedings of the 25th International Conference on Machine Learning*. New York, NY, USA, pages 160–167.
- Chris Dyer, Adhiguna Kuncoro, Miguel Ballesteros, and A. Noah Smith. 2016. Recurrent neural network grammars. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, pages 199–209.


- Jeffrey L. Elman. 1991. Distributed representations, simple recurrent networks, and grammatical structure. *Machine Learning* 7(2-3):195–225.
- Julie Franck, Gabriella Vigliocco, and Janet Nicol. 2002. Subject-verb agreement errors in French and English: The role of syntactic hierarchy. *Language and Cognitive Processes* 17(4):371–404.
- Edward Grefenstette, Karl Moritz Hermann, Mustafa Suleyman, and Phil Blunsom. 2015. Learning to transduce with unbounded memory. In *Advances in Neural Information Processing Systems* 28. pages 1828–1836.
- Kazuma Hashimoto, Caiming Xiong, Yoshimasa Tsuruoka, and Richard Socher. 2016. A joint many-task model: Growing a neural network for multiple NLP tasks. In *NIPS 2016 Continual Learning and Deep Networks Workshop*.
- Julia Hockenmaier and Mark Steedman. 2007. CCG-bank: A corpus of CCG derivations and dependency structures extracted from the Penn Treebank. *Computational Linguistics* 33(3):355–396.
- Rafal Jozefowicz, Oriol Vinyals, Mike Schuster, Noam Shazeer, and Yonghui Wu. 2016. Exploring the limits of language modeling. *arXiv preprint arXiv:1602.02410*.
- Jey Han Lau, Alexander Clark, and Shalom Lappin. 2016. Grammaticality, acceptability, and probability: A probabilistic view of linguistic knowledge. *Cognitive Science*.
- Gaël Le Godais, Tal Linzen, and Emmanuel Dupoux. 2017. Comparing character-level neural language models using a lexical decision task. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*. Association for Computational Linguistics, Valencia, Spain, pages 125–130.
- Mike Lewis, Kenton Lee, and Luke Zettlemoyer. 2016. LSTM CCG parsing. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. pages 221–231.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. 2016. Assessing the ability of LSTMs to learn syntax-sensitive dependencies. *Transactions of the Association for Computational Linguistics* 4:521–535.
- Tal Linzen, Yoav Goldberg, and Emmanuel Dupoux. 2017. Agreement attraction errors in neural networks. In *Proceedings of the CUNY Conference on Human Sentence Processing*.
- Mitchell P. Marcus, Mary Ann Marcinkiewicz, and Beatrice Santorini. 1993. Building a large annotated corpus of English: The Penn Treebank. *Computational Linguistics* 19(2):313–330.
- Héctor Martínez Alonso and Barbara Plank. 2017. When is multitask learning effective? Semantic sequence prediction under varying data conditions. In *Proceedings of the Conference of the European Chapter of the Association for Computational Linguistics*.
- Tomas Mikolov, Martin Karafiát, Lukas Burget, Jan Cernocký, and Sanjeev Khudanpur. 2010. Recurrent neural network based language model. In *Proceedings of Interspeech*.
- Richard Socher. 2014. *Recursive Deep Learning for Natural Language Processing and Computer Vision*. Ph.D. thesis, Stanford University.
- Anders Søgaard and Yoav Goldberg. 2016. Deep multi-task learning with low level tasks supervised at lower layers. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Association for Computational Linguistics, Berlin, Germany, pages 231–235.
- Adrian Staub. 2009. On the interpretation of the number attraction effect: Response time evidence. *Journal of Memory and Language* 60(2):308–327.
- Mark Steedman. 2000. *The syntactic process*. MIT Press.
- Martin Sundermeyer, Ralf Schlüter, and Hermann Ney. 2012. LSTM neural networks for language modeling. In *Proceedings of the 13th Annual Conference of the International Speech Communication Association (INTERSPEECH)*. pages 194–197.
- Theano Development Team. 2016. [Theano: A Python framework for fast computation of mathematical expressions](http://arxiv.org/abs/1605.02688). *arXiv e-prints* abs/1605.02688. <http://arxiv.org/abs/1605.02688>.
- Ashish Vaswani, Yonatan Bisk, Kenji Sagae, and Ryan Musa. 2016. Supertagging with LSTMs. In *Proceedings of NAACL-HLT*. pages 232–237.
- Matthew W. Wagers, Ellen F. Lau, and Colin Phillips. 2009. Agreement attraction in comprehension: Representations and processes. *Journal of Memory and Language* 61(2):206–237.
- Wenduan Xu, Michael Auli, and Stephen Clark. 2015. CCG supertagging with a recurrent neural network. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*. Association for Computational Linguistics, Beijing, China, pages 250–255.

A Appendix


This appendix presents figures based on sentences with relative clause (see Section 4.4). Figure 8 tracks the word-by-word predictions that the single-task model and the pre-trained model make for three sample sentences; the grammatical ground truth is indicated with a dotted black line. Overall, the pre-trained model is closer to the ground truth than the single-task model, even in cases where both models ultimately make the correct prediction (Figure 8b). Figures 8a and 8c show cases in which an attractor in an embedded clause misleads the single-task but not the pre-trained one. Finally, Figure 9 shows a sample of four units that appear to track interpretable aspects of the sentence.



(a) Bock and Cutting (1992): PS




(b) Wagers et al. (2009): PS




(c) Wagers et al. (2009): SP


Figure 8: Probability of a plural prediction after each word in the sentence for three sample sentences. The black dotted line indicates the grammatical ground truth.




(a) Unit 30: approximately tracks the number of the currently relevant subject



(b) Unit 50: only active within noun phrases



(c) Unit 73: represents of the number of the main clause subject



(d) Unit 86: approximately tracks the number of the currently relevant subject

Figure 9: Activations of a sample of interpretable units throughout an example sentence from Wagers et al. (2009), for all four number configurations.