

Figure 1: Overall model diagram

semantics and syntax also deviate significantly from English (Leimgruber, 2009, 2011). We categorize the challenges and formalize their interpretation using Universal Dependencies (Nivre et al., 2016), which extends to the creation of a Singlish dependency treebank with 1,200 sentences.

Based on the intricate relationship between Singlish and English, we build a Singlish parser by leveraging knowledge of English syntax as a basis. This overall approach is illustrated in Figure 1. In particular, we train a basic Singlish parser with the best off-the-shelf neural dependency parsing model using biaffine attention (Dozat and Manning, 2017), and improve it with knowledge transfer by adopting neural stacking (Chen et al., 2016; Zhang and Weiss, 2016) to integrate the English syntax. Since POS tags are important features for dependency parsing (Chen and Manning, 2014; Dyer et al., 2015), we train a POS tagger for Singlish following the same idea by integrating English POS knowledge using neural stacking.

Results show that English syntax knowledge brings 51.50% and 25.01% relative error reduction on POS tagging and dependency parsing respectively, resulting in a Singlish dependency parser with 84.47% unlabeled attachment score (UAS) and 77.76% labeled attachment score (LAS).

We make our Singlish dependency treebank, the source code for training a dependency parser and the trained model for the parser with the best performance freely available online².

²https://github.com/wanghm92/Sing_Par

2 Related Work

Neural networks have led to significant advance in the performance for dependency parsing, including transition-based parsing (Chen and Manning, 2014; Zhou et al., 2015; Weiss et al., 2015; Dyer et al., 2015; Ballesteros et al., 2015; Andor et al., 2016), and graph-based parsing (Kiperwasser and Goldberg, 2016; Dozat and Manning, 2017). In particular, the biaffine attention method of Dozat and Manning (2017) uses deep bi-directional long short-term memory (bi-LSTM) networks for high-order non-linear feature extraction, producing the highest-performing graph-based English dependency parser. We adopt this model as the basis for our Singlish parser.

Our work belongs to a line of work on transfer learning for parsing, which leverages English resources in Universal Dependencies to improve the parsing accuracies of low-resource languages (Hwa et al., 2005; Cohen and Smith, 2009; Ganchev et al., 2009). Seminal work employed statistical models. McDonald et al. (2011) investigated delexicalized transfer, where word-based features are removed from a statistical model for English, so that POS and dependency label knowledge can be utilized for training a model for low-resource language. Subsequent work considered syntactic similarities between languages for better feature transfer (Täckström et al., 2012; Naseem et al., 2012; Zhang and Barzilay, 2015).

Recently, a line of work leverages neural network models for multi-lingual parsing (Guo et al., 2015; Duong et al., 2015; Ammar et al., 2016). The basic idea is to map the word embedding spaces between different languages into the same vector space, by using sentence-aligned bilingual data. This gives consistency in tokens, POS and dependency labels thanks to the availability of Universal Dependencies (Nivre et al., 2016). Our work is similar to these methods in using a neural network model for knowledge sharing between different languages. However, ours is different in the use of a neural stacking model, which respects the distributional differences between Singlish and English words. This empirically gives higher accuracies for Singlish.

Neural stacking was previously used for cross-annotation (Chen et al., 2016) and cross-task (Zhang and Weiss, 2016) joint-modelling on monolingual treebanks. To the best of our knowledge, we are the first to employ it on cross-lingual

feature transfer from resource-rich languages to improve dependency parsing for low-resource languages. Besides these three dimensions in dealing with heterogeneous text data, another popular area of research is on the topic of domain adaptation, which is commonly associated with cross-lingual problems (Nivre et al., 2007). While this large strand of work is remotely related to ours, we do not describe them in details.

Unsupervised rule-based approaches also offer an competitive alternative for cross-lingual dependency parsing (Naseem et al., 2010; Gillenwater et al., 2010; Gelling et al., 2012; Søggaard, 2012a,b; Martínez Alonso et al., 2017), and recently been benchmarked for the Universal Dependencies formalism by exploiting the linguistic constraints in the Universal Dependencies to improve the robustness against error propagation and domain adaptation (Martínez Alonso et al., 2017). However, we choose a data-driven supervised approach given the relatively higher parsing accuracy owing to the availability of resourceful treebanks from the Universal Dependencies project.

3 Singlish Dependency Treebank

3.1 Universal Dependencies for Singlish

Since English is the major genesis of Singlish, we choose English as the source of lexical feature transfer to assist Singlish dependency parsing. Universal Dependencies provides a set of multilingual treebanks with cross-lingually consistent dependency-based lexicalist annotations, designed to aid development and evaluation for cross-lingual systems, such as multilingual parsers (Nivre et al., 2016). The current version of Universal Dependencies comprises not only major treebanks for 47 languages but also their siblings for domain-specific corpora and dialects. With the aligned initiatives for creating transfer-learning-friendly treebanks, we adopt the Universal Dependencies protocol for constructing the Singlish dependency treebank, both as a new resource for the low-resource languages and to facilitate knowledge transfer from English.

On top of the general Universal Dependencies guidelines, English-specific dependency relation definitions including additional subtypes are employed as the default standards for annotating the Singlish dependency treebank, unless augmented or redefined when necessary. The latest English

	UD English		Singlish	
	Sentences	Words	Sentences	Words
Train	12,543	204,586	900	8,221
Dev	2,002	25,148	150	1,384
Test	2,077	25,096	150	1,381

Table 1: Division of training, development, and test sets for Singlish Treebank

corpus in Universal Dependencies v1.4³ collection is constructed from the English Web Treebank (Bies et al., 2012), comprising of web media texts, which potentially smooths the knowledge transfer to our target Singlish texts in similar domains. The statistics of this dataset, from which we obtain English syntactic knowledge, is shown in Table 1 and we refer to this corpus as UD-Eng. This corpus uses 47 dependency relations and we show below how to conform to the same standard while adapting to unique Singlish grammars.

3.2 Challenges and Solutions for Annotating Singlish

The deviations of Singlish from English come from both the lexical and the grammatical levels (Leimgruber, 2009, 2011), which bring challenges for analysis on Singlish using English NLP tools. The former involves imported vocabularies from the first languages of the local people and the latter can be represented by a set of relatively localized features which collectively form 5 unique grammars of Singlish according to Leimgruber (2011). We find empirically that all these deviations can be accommodated by applying the existing English dependency relation definitions while ensuring consistency with the annotations in other non-English UD treebanks, which are explained with examples as follows.

Imported vocabulary: Singlish borrows a number of words and expressions from its non-English origins (Leimgruber, 2009, 2011), such as “Kiasu”, which originates from Hokkien meaning “very anxious not to miss an opportunity”.⁴ These imported terms often constitute out-of-vocabulary (OOV) words with respect to a standard English treebank and result in difficulties for using English-trained tools on Singlish. All borrowed words are annotated based on their usages in Singlish, which mainly inherit the POS from their genesis languages. Table A4 in Appendix A

³Only guidelines for Universal Dependencies v2 but not the English corpus is available when this work is completed.

⁴Definition by the Oxford living Dictionaries for English.

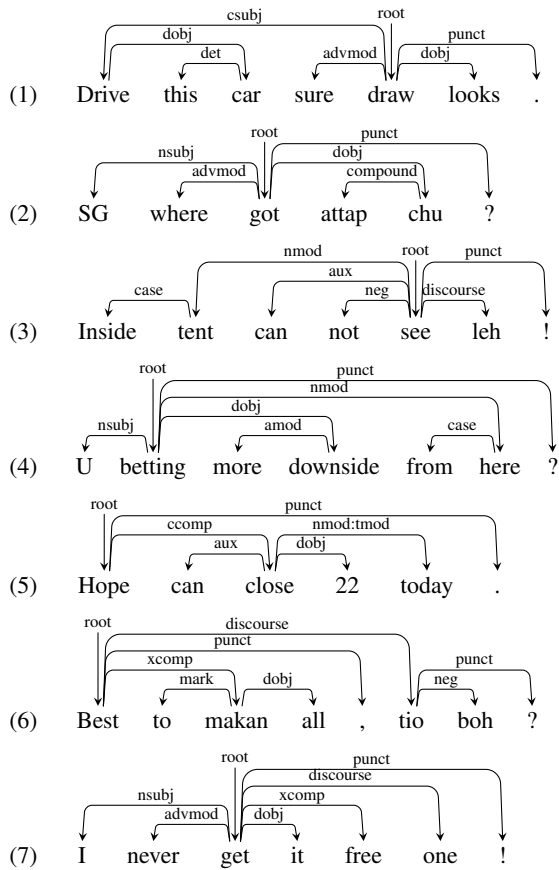


Figure 2: Unique Singlish grammars. (Arcs represent dependencies, pointing from the head to the dependent, with the dependency relation label right on top of the arc)

summarizes all borrowed terms in our treebank.

Topic-prominence: This type of sentences start with establishing its topic, which often serves as the default one that the rest of the sentence refers to, and they typically employ an object-subject-verb sentence structure (Leimgruber, 2009, 2011). In particular, three subtypes of topic-prominence are observed in the Singlish dependency treebank and their annotations are addressed as follows:

First, topics framed as clausal arguments at the beginning of the sentence are labeled as “csubj” (clausal subject), as shown by “Drive this car” of (1) in Figure 2, which is consistent with the dependency relations in its Chinese translation.

Second, noun phrases used to modify the predicate with the absence of a preposition is regarded as a “nsubj” (nominal subject). Similarly, this is a common order of words used in Chinese and one example is the “SG” of (2) in Figure 2.

Third, prepositional phrases moved in front are still treated as “nmod” (nominal modifier) of their

intended heads, following the exact definition but as a Singlish-specific form of exemplification, as shown by the “Inside tent” of (3) in Figure 2.

Although the “dislocated” (dislocated elements) relation in UD is also used for preposed elements, but it captures the ones “that do not fulfill the usual core grammatical relations of a sentence” and “not for a topic-marked noun that is also the subject of the sentence” (Nivre et al., 2016). In these three scenarios, the topic words or phrases are in relatively closer grammatical relations to the predicate, as subjects or modifiers.

Copula deletion: Imported from the corresponding Chinese sentence structure, this copula verb is often optional and even deleted in Singlish, which is one of its diagnostic characteristics (Leimgruber, 2009, 2011). In UD-Eng standards, predicative “be” is the only verb used as a copula and it often depends on its complement to avoid copular head. This is explicitly designed in UD to promote parallelism for zero-copula phenomenon in languages such as Russian, Japanese, and Arabic. The deleted copula and its “cop” (copula) arcs are simply ignored, as shown by (4) in Figure 2.

NP deletion: Noun-phrase (NP) deletion often results in null subjects or objects. It may be regarded as a branch of “Topic-prominence” but is a distinctive feature of Singlish with relatively high frequency of usage (Leimgruber, 2011). NP deletion is also common in pronoun-dropping languages such as Spanish and Italian, where the anaphora can be morphologically inferred. In one example, “Vorrei ora entrare brevemente nel merito.”⁵, from the Italian treebank in UD, “Vorrei” means “I would like to” and depends on the sentence root, “entrare”, with the “aux”(auxiliary) relation, where the subject “I” is absent but implicitly understood. Similarly, we do not recover such relations since the deleted NP imposes negligible alteration to the dependency tree, as exemplified by (5) in Figure 2.

Inversion: Inversion in Singlish involves either keeping the subject and verb in interrogative sentences in the same order as in statements, or tag questions in polar interrogatives (Leimgruber, 2011). The former also exists in non-English languages, such as Spanish and Italian, where the subject can prepose the verb in questions (La-

⁵In English: (I) would now like to enter briefly on the merit (of the discussion).

housse and Lamiroy, 2012). This simply involves a change of word orders and thus requires no special treatments. On the other hand, tag questions should be carefully analyzed in two scenarios. One type is in the form of “isn’t it?” or “haven’t you?”, which are dependents of the sentence root with the “parataxis” relation.⁶ The other type is exemplified as “right?”, and its Singlish equivalent “tio boh?” (a transliteration from Hokkien) are labeled with the “discourse” (discourse element) relation with respect to the sentence root. See example (6) in Figure 2.

Discourse particles: Usage of clausal-final discourse particles, which originates from Hokkien and Cantonese, is one of the most typical feature of Singlish (Leimgruber, 2009, 2011; Lim, 2007). All discourse particles that appear in our treebank are summarized in Table A3 in Appendix A with the imported vocabulary. These words express the tone of the sentence and thus have the “INTJ” (interjection) POS tag and depend on the root of the sentence or clause labeled with “discourse”, as is shown by the “leh” of (3) in Figure 2. The word “one” is a special instance of this type with the sole purpose being a tone marker in Singlish but not English, as shown by (7) in Figure 2.

3.3 Data Selection and Annotation

Data Source: Singlish is used in written form mainly in social media and local Internet forums. After comparison, we chose the *SG Talk Forum*⁷ as our data source due to its relative abundance in Singlish contents. We crawled 84,459 posts using the Scrapy framework⁸ from pages dated up to 25th December 2016, retaining sentences of length between 5 and 50, which total 58,310. Sentences are reversely sorted according to the log likelihood of the sentence given by an English language model trained using the KenLM toolkit (Heafield et al., 2013)⁹ normalized by the sentence length, so that those most different from standard English can be chosen. Among the top 10,000 sentences, 1,977 sentences contain unique Singlish vocabularies defined by *The*

*Coxford Singlish Dictionary*¹⁰, *A Dictionary of Singlish and Singapore English*¹¹, and the *Singlish Vocabulary* Wikipedia page¹². The average normalized log likelihood of these 10,000 sentences is -5.81, and the same measure for all sentences in UD-Eng is -4.81. This means these sentences with Singlish contents are 10 times less probable expressed as standard English than the UD-Eng contents in the web domain. This contrast indicates the degree of lexical deviation of Singlish from English. We chose 1,200 sentences from the first 10,000. More than 70% of the selected sentences are observed to consist of the Singlish grammars and imported vocabularies described in section 3.2. Thus the evaluations on this treebank can reflect the performance of various POS taggers and parsers on Singlish in general.

Annotation: The chosen texts are divided by random selection into training, development, and testing sets according to the proportion of sentences in the training, development, and test division for UD-Eng, as summarized in Table 1. The sentences are tokenized using the NLTK Tokenizer,¹³ and then annotated using the Dependency Viewer.¹⁴ In total, all 17 UD-Eng POS tags and 41 out of the 47 UD-Eng dependency labels are present in the Singlish dependency treebank. Besides, 100 sentences are randomly selected and double annotated by one of the coauthors, and the inter-annotator agreement has a 97.76% accuracy on POS tagging and a 93.44% UAS and a 89.63% LAS for dependency parsing. A full summary of the numbers of occurrences of each POS tag and dependency label are included in Appendix A.

4 Part-of-Speech Tagging

In order to obtain automatically predicted POS tags as features for a base English dependency parser, we train a POS tagger for UD-Eng using the baseline model of Chen et al. (2016), depicted in Figure 3. The bi-LSTM networks with a CRF layer (bi-LSTM-CRF) have shown state-of-the-art performance by globally optimizing the tag sequence (Huang et al., 2015; Chen et al., 2016).

⁶In UD: Relation between the main verb of a clause and other sentential elements, such as sentential parenthetical clause, or adjacent sentences without any explicit coordination or subordination.

⁷<http://sgTalk.com>

⁸<https://scrapy.org/>

⁹Trained using the *afp_eng* and *xin_eng* sources of English Gigaword Fifth Edition (Gigaword).

¹⁰<http://72.5.72.93/html/lexec.php>

¹¹<http://www.singlishdictionary.com>

¹²https://en.wikipedia.org/wiki/Singlish_vocabulary

¹³<http://www.nltk.org/api/nltk.tokenize.html>

¹⁴<http://nlp.nju.edu.cn/tanggc/tools/DependencyViewer.exe>




Figure 3: Base POS tagger

Based on this English POS tagging model, we train a POS tagger for Singlish using the feature-level neural stacking model of Chen et al. (2016). Both the English and Singlish models consist of an input layer, a feature layer, and an output layer.

4.1 Base Bi-LSTM-CRF POS Tagger

Input Layer: Each token is represented as a vector by concatenating a word embedding from a lookup table with a weighted average of its character embeddings given by the attention model of Bahdanau et al. (2014). Following Chen et al. (2016), the input layer produces a dense representation for the current input token by concatenating its word vector and the ones for its surrounding context tokens in a window of finite size.

Feature Layer: This layer employs a bi-LSTM network to encode the input into a sequence of hidden vectors that embody global contextual information. Following Chen et al. (2016), we adopt bi-LSTM with peephole connections (Graves and Schmidhuber, 2005).

Output layer: This is a CRF layer to predict the POS tags for the input words by maximizing the conditional probability of the sequence of tags given input sentence.

4.2 POS Tagger with Neural Stacking

We adopt the deep integration neural stacking structure presented in Chen et al. (2016). As shown in Figure 4, the distributed vector representation for the target word at the input layer of the Singlish Tagger is augmented by concatenating the emission vector produced by the English Tagger with the original word and character-based embeddings, before applying the concatenation within a context window in section 4.1. During training, loss is back-propagated to all trainable parameters

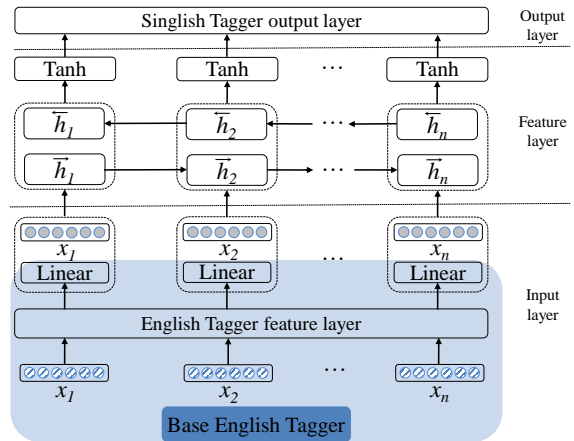


Figure 4: POS tagger with neural stacking

System	Accuracy
ENG-on-SIN	81.39%
Base-ICE-SIN	78.35%
Stack-ICE-SIN	89.50%

Table 2: POS tagging accuracies

in both the Singlish Tagger and the pre-trained feature layer of the base English Tagger. At test time, the input sentence is fed to the integrated tagger model as a whole for inference.

4.3 Results

We use the publicly available source code¹⁵ by Chen et al. (2016) to train a 1-layer bi-LSTM-CRF based POS tagger on UD-Eng, using 50-dimension pre-trained SENNA word embeddings (Collobert et al., 2011). We set the hidden layer size to 300, the initial learning rate for Adagrad (Duchi et al., 2011) to 0.01, the regularization parameter λ to 10^{-6} , and the dropout rate to 15%. The tagger gives 94.84% accuracy on the UD-Eng test set after 24 epochs, chosen according to development tests, which is comparable to the state-of-the-art accuracy of 95.17% reported by Plank et al. (2016). We use these settings to perform 10-fold jackknifing of POS tagging on the UD-Eng training set, with an average accuracy of 95.60%.

Similarly, we trained a POS tagger using the Singlish dependency treebank alone with pre-trained word embeddings on The Singapore Component of the International Corpus of English (ICE-SIN) (Nihilani, 1992; Ooi, 1997), which consists of both spoken and written texts. However, due to limited amount of training data, the

¹⁵<https://github.com/chenhongshen/NNHetSeq>

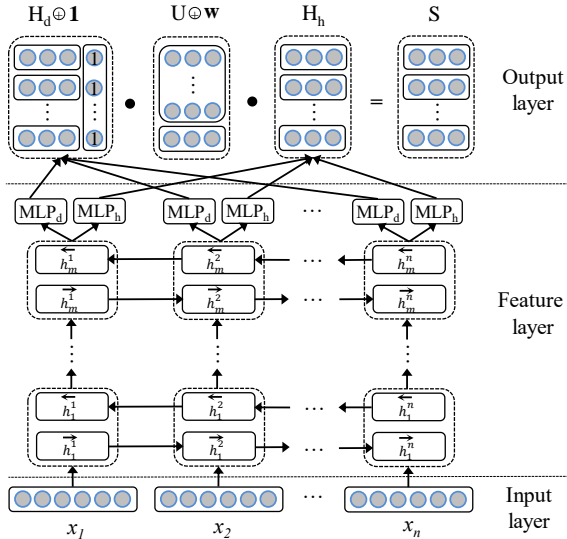


Figure 5: Base parser

tagging accuracy is not satisfactory even with a larger dropout rate to avoid over-fitting. In contrast, the neural stacking structure on top of the English base model trained on UD-Eng achieves a POS tagging accuracy of 89.50%¹⁶, which corresponds to a 51.50% relative error reduction over the baseline Singlish model, as shown in Table 2. We use this for 10-fold jackknifing on Singlish parsing training data, and tagging the Singlish development and test data.

5 Dependency Parsing

We adopt the Dozat and Manning (2017) parser¹⁷ as our base model, as displayed in Figure 5, and apply neural stacking to achieve improvements over the baseline parser. Both the base and neural stacking models consist of an input layer, a feature layer, and an output layer.

5.1 Base Parser with Bi-affine Attentions

Input Layer: This layer encodes the current input word by concatenating a pre-trained word embedding with a trainable word embedding and POS tag embedding from the respective lookup tables.

Feature Layer: The two recurrent vectors produced by the multi-layer bi-LSTM network from each input vector are concatenated and mapped to multiple feature vectors in lower-dimension space by a set of parallel multilayer perceptron (MLP)

¹⁶We empirically find that using ICE-SIN embeddings in neural stacking model performs better than using English SENNA embeddings. Similar findings are found for the parser, of which more details are given in section 6.

¹⁷<https://github.com/tdozat/Parser>

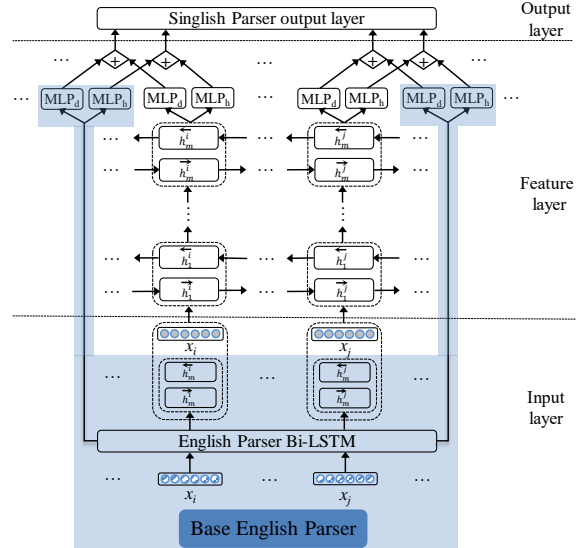


Figure 6: Parser with neural stacking

layers. Following Dozat and Manning (2017), we adopt Cif-LSTM cells (Greff et al., 2016).

Output Layer: This layer applies biaffine transformation on the feature vectors to calculate the score of the directed arcs between every pair of words. The inferred trees for input sentence are formed by choosing the head with the highest score for each word and a cross-entropy loss is calculated to update the model parameters.

5.2 Parser with Neural Stacking

Inspired by the idea of feature-level neural stacking (Chen et al., 2016; Zhang and Weiss, 2016), we concatenate the pre-trained word embedding, trainable word and tag embeddings, with the two recurrent state vectors at the last bi-LSTM layer of the English Tagger as the input vector for each target word. In order to further preserve syntactic knowledge retained by the English Tagger, the feature vectors from its MLP layer is added to the ones produced by the Singlish Parser, as illustrated in Figure 6, and the scoring tensor of the Singlish Parser is initialized with the one from the trained English Tagger. Loss is back-propagated by reversely traversing all forward paths to all trainable parameter for training and the whole model is used collectively for inference.

6 Experiments

6.1 Experimental Settings

We train an English parser on UD-Eng with the default model settings in Dozat and Manning (2017).

	Sentences	Words	Vocabulary
GloVe6B	N.A.	6000m	400,000
Giga100M	57,000	1.26m	54,554
ICE-SIN	87,084	1.26m	40,532

Table 3: Comparison of the scale of sources for training word embeddings

Trained on	System	UAS	LAS
English	ENG-on-SIN	75.89	65.62
Singlish	Baseline	75.98	66.55
	Base-Giga100M	77.67	67.23
	Base-GloVe6B	78.18	68.51
	Base-ICE-SIN	79.29	69.27
Both	ENG-plus-SIN	82.43	75.64
	Stack-ICE-SIN	84.47	77.76

Table 4: Dependency parser performances

It achieves an UAS of 88.83% and a LAS of 85.20%, which are close to the state-of-the-art 85.90% LAS on UD-Eng reported by Ammar et al. (2016), and the main difference is caused by us not using fine-grained POS tags. We apply the same settings for a baseline Singlish parser. We attempt to choose a better configuration of the number of bi-LSTM layers and the hidden dimension based on the development set performance, but the default settings turn out to perform the best. Thus we stick to all default hyper-parameters in Dozat and Manning (2017) for training the Singlish parsers.

We experimented with different word embeddings, as with the raw text sources summarized in Table 3 and further described in section 6.2. When using the neural stacking model, we fix the model configuration for the base English parser model and choose the size of the hidden vector and the number of bi-LSTM layers stacked on top based on the performance on the development set. It turns out that a 1-layer bi-LSTM with 900 hidden dimension performs the best, where the bigger hidden layer accommodates the elongated input vector to the stacked bi-LSTM and the fewer number of recurrent layers avoids over-fitting on the small Singlish dependency treebank, given the deep bi-LSTM English parser network at the bottom. The evaluation of the neural stacking model is further described in section 6.3.

System	UAS	LAS
Base-ICE-SIN	77.00	66.69
Stack-ICE-SIN	82.43	73.96

Table 5: Dependency parser performances by the 5-cross-fold validation

6.2 Investigating Distributed Lexical Characteristics

In order to learn characteristics of distributed lexical semantics for Singlish, we compare performances of the Singlish dependency parser using several sets of pre-trained word embeddings: GloVe6B, large-scale English word embeddings¹⁸; ICE-SIN, Singlish word embeddings trained using GloVe (Pennington et al., 2014) on the ICE-SIN (Nihilani, 1992; Ooi, 1997) corpus; Giga100M, a small-scale English word embeddings trained using GloVe (Pennington et al., 2014) with the same settings on a comparable size of English data randomly selected from the English Gigaword Fifth Edition for a fair comparison with ICE-SIN embeddings.

First, the English Giga100M embeddings marginally improve the Singlish parser from the baseline without pre-trained embeddings and also using the UD-Eng parser directly on Singlish, represented as “ENG-on-SIN” in Table 4. With much more English lexical semantics being fed to the Singlish parser using the English GloVe6B embeddings, further enhancement is achieved. Nevertheless, the Singlish ICE-SIN embeddings lead to even more improvement, with 13.78% relative error reduction, compared with 7.04% using the English Giga100M embeddings and 9.16% using the English GloVe6B embeddings, despite the huge difference in sizes in the latter case. This demonstrates the distributional differences between Singlish and English tokens, even though they share a large vocabulary. More detailed comparison is described in section 6.4.

6.3 Knowledge Transfer Using Neural Stacking

We train a parser with neural stacking and Singlish ICE-SIN embeddings, which achieves the best performance among all the models, with a UAS of 84.47%, represented as “Stack-ICE-SIN” in Table 4, which corresponds to 25.01% relative error reduction compared to the baseline. This demonstrates that knowledge from English can be successfully incorporated to boost the Singlish parser. To further evaluate the effectiveness of the neural stacking model, we also trained a base model with the combination of UD-Eng and the Singlish tree-

¹⁸Trained with Wikipedia 2014 the Gigaword. Downloadable from <http://nlp.stanford.edu/data/glove.6B.zip>

	Topic Prominence		Copula Deletion		NP Deletion		Discourse Particles		Others	
Sentences	15		19		21		51		67	
	UAS	LAS	UAS	LAS	UAS	LAS	UAS	LAS	UAS	LAS
ENG-on-SIN	78.15	62.96	66.91	56.83	72.57	64.00	70.00	59.00	78.92	68.47
Base-Giga100M	77.78	68.52	71.94	61.15	76.57	69.14	85.25	77.25	73.13	60.63
Base-ICE	81.48	72.22	74.82	63.31	80.00	73.71	85.25	77.75	75.56	64.37
Stack-ICE	87.04	76.85	77.70	71.22	80.00	75.43	88.50	83.75	84.14	76.49

Table 6: Error analysis with respect to grammar types

bank, represented as “ENG-plus-SIN” in Table 4, which is still outperformed by the neural stacking model. Besides, we performed a 5-cross-fold validation for the base parser with Singlish ICE-SIN embeddings and the parser using neural stacking, where half of the held-out fold is used as the development set. The average UAS and LAS across the 5 folds shown in Table 5 and the relative error reduction on average 23.61% suggest that the overall improvement from knowledge transfer using neural stacking remains consistent. This significant improvement is further explained in section 6.4.

6.4 Improvements over Grammar Types

To analyze the sources of improvements for Singlish parsing using different model configurations, we conduct error analysis over 5 syntactic categories¹⁹, including 4 types of grammars mentioned in section 3.2²⁰, and 1 for all other cases, including sentences containing imported vocabularies but expressed in basic English syntax. The number of sentences and the results in each group of the test set are shown in Table 6.

The neural stacking model leads to the biggest improvement over all categories except for a tie UAS performance on “NP Deletion” cases, which explains the significant overall improvement.

Comparing the base model with ICE-SIN embeddings with the base parser trained on UD-Eng, which contain syntactic and semantic knowledge in Singlish and English, respectively, the former outperforms the latter on all 4 types of Singlish grammars but not for the remaining samples. This suggests that the base English parser mainly contributes to analyzing basic English syntax, while the base Singlish parser models unique Singlish grammars better.

Similar trends are also observed on the base model using the English Giga100M embeddings, but the overall performances are not as good as

using ICE-SIN embeddings, especially over basic English syntax where it undermines the performance to a greater extent. This suggests that only limited English distributed lexical semantic information can be integrated to help modelling Singlish syntactic knowledge due to the differences in distributed lexical semantics.

7 Conclusion

We have investigated dependency parsing for Singlish, an important English-based creole language, through annotations of a Singlish dependency treebank with 10,986 words and building an enhanced parser by leveraging on knowledge transferred from a 20-times-bigger English treebank of Universal Dependencies. We demonstrate the effectiveness of using neural stacking for feature transfer by boosting the Singlish dependency parsing performance to from UAS 79.29% to UAS 84.47%, with a 25.01% relative error reduction over the parser with all available Singlish resources. We release the annotated Singlish dependency treebank, the trained model and the source code for the parser with free public access. Possible future work include expanding the investigation to other regional languages such as Malay and Indonesian.

Acknowledgments

Yue Zhang is the corresponding author. This research is supported by IGDSS1603031 from Temasek Laboratories@SUTD. We appreciate anonymous reviewers for their insightful comments, which helped to improve the paper, and Zhiyang Teng, Jiangming Liu, Yupeng Liu, and Enrico Santus for their constructive discussions.

¹⁹Multiple labels are allowed for one sentence.

²⁰The “Inversion” type of grammar is not analyzed since there is only 1 such sentence in the test set.

References

- Waleed Ammar, George Mulcaire, Miguel Ballesteros, Chris Dyer, and Noah Smith. 2016. [Many languages, one parser](#). *Transactions of the Association of Computational Linguistics* 4:431–444. <http://aclweb.org/anthology/Q16-1031>.
- Daniel Andor, Chris Alberti, David Weiss, Aliaksei Severyn, Alessandro Presta, Kuzman Ganchev, Slav Petrov, and Michael Collins. 2016. [Globally normalized transition-based neural networks](#). In *Proceedings of the ACL 2016*. Association for Computational Linguistics, pages 2442–2452. <https://doi.org/10.18653/v1/P16-1231>.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. [Neural machine translation by jointly learning to align and translate](#). *arXiv preprint abs/1409.0473*. <http://arxiv.org/abs/1409.0473>.
- Miguel Ballesteros, Chris Dyer, and A. Noah Smith. 2015. [Improved transition-based parsing by modeling characters instead of words with lstms](#). In *Proceedings of the EMNLP 2015*. Association for Computational Linguistics, pages 349–359. <https://doi.org/10.18653/v1/D15-1041>.
- Ann Bies, Justin Mott, Colin Warner, and Seth Kulick. 2012. English web treebank ldc2012t13 .
- Danqi Chen and Christopher Manning. 2014. [A fast and accurate dependency parser using neural networks](#). In *Proceedings of the EMNLP 2014*. Association for Computational Linguistics, pages 740–750. <https://doi.org/10.3115/v1/D14-1082>.
- Hongshen Chen, Yue Zhang, and Qun Liu. 2016. [Neural network for heterogeneous annotations](#). In *Proceedings of the EMNLP 2016*. Association for Computational Linguistics, pages 731–741. <http://aclweb.org/anthology/D16-1070>.
- Shay Cohen and A. Noah Smith. 2009. [Shared logistic normal distributions for soft parameter tying in unsupervised grammar induction](#). In *Proceedings of the NAACL-HLT 2009*. Association for Computational Linguistics, pages 74–82. <http://aclweb.org/anthology/N09-1009>.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. [Natural language processing \(almost\) from scratch](#). *Journal of Machine Learning Research* 12:2493–2537. <http://dl.acm.org/citation.cfm?id=2078186>.
- Timothy Dozat and Christopher D. Manning. 2017. [Deep biaffine attention for neural dependency parsing](#). In *International Conference on Learning Representations 2017*. volume abs/1611.01734. <http://arxiv.org/abs/1611.01734>.
- John C. Duchi, Elad Hazan, and Yoram Singer. 2011. [Adaptive subgradient methods for online learning and stochastic optimization](#). *Journal of Machine Learning Research* 12:2121–2159. <http://dl.acm.org/citation.cfm?id=2021068>.
- Long Duong, Trevor Cohn, Steven Bird, and Paul Cook. 2015. [A neural network model for low-resource universal dependency parsing](#). In *Proceedings of the EMNLP 2015*. Association for Computational Linguistics, pages 339–348. <https://doi.org/10.18653/v1/D15-1040>.
- Chris Dyer, Miguel Ballesteros, Wang Ling, Austin Matthews, and A. Noah Smith. 2015. [Transition-based dependency parsing with stack long short-term memory](#). In *Proceedings of the ACL-IJCNLP 2015*. Association for Computational Linguistics, pages 334–343. <https://doi.org/10.3115/v1/P15-1033>.
- Kuzman Ganchev, Jennifer Gillenwater, and Ben Taskar. 2009. [Dependency grammar induction via bitext projection constraints](#). In *Proceedings of the ACL-IJCNLP 2009*. Association for Computational Linguistics, pages 369–377. <http://aclweb.org/anthology/P09-1042>.
- Douwe Gelling, Trevor Cohn, Phil Blunsom, and Joao Graca. 2012. [The pascal challenge on grammar induction](#). In *Proceedings of the NAACL-HLT Workshop on the Induction of Linguistic Structure*. Association for Computational Linguistics, pages 64–80. <http://www.aclweb.org/anthology/W12-1909>.
- Jennifer Gillenwater, Kuzman Ganchev, João Graça, Fernando Pereira, and Ben Taskar. 2010. [Sparsity in dependency grammar induction](#). In *Proceedings of the ACL 2010 (Short Papers)*. Association for Computational Linguistics, pages 194–199. <http://www.aclweb.org/anthology/P10-2036>.
- Alex Graves and Jürgen Schmidhuber. 2005. [Frame-wise phoneme classification with bidirectional lstm and other neural network architectures](#). *Neural Networks* 18(5):602–610.
- K. Greff, R. K. Srivastava, J. Koutnk, B. R. Steunebrink, and J. Schmidhuber. 2016. [Lstm: A search space odyssey](#). *IEEE Transactions on Neural Networks and Learning Systems* PP(99):1–11. <https://doi.org/10.1109/TNNLS.2016.2582924>.
- Jiang Guo, Wanxiang Che, David Yarowsky, Haifeng Wang, and Ting Liu. 2015. [Cross-lingual dependency parsing based on distributed representations](#). In *Proceedings of the ACL-IJCNLP 2015*. Association for Computational Linguistics, pages 1234–1244. <https://doi.org/10.3115/v1/P15-1119>.
- Shinichi Harada. 2009. [The roles of singapore standard english and singlish](#). *Information Research* 40:70–82.

- Kenneth Heafield, Ivan Pouzyrevsky, H. Jonathan Clark, and Philipp Koehn. 2013. Scalable modified kneser-ney language model estimation. In *Proceedings of the ACL 2013 (Short Papers)*. Association for Computational Linguistics, pages 690–696. <http://aclweb.org/anthology/P13-2121>.
- Sanjika Hewavitharana, Nguyen Bach, Qin Gao, Vamshi Ambati, and Stephan Vogel. 2011. *Proceedings of the Sixth Workshop on Statistical Machine Translation*, Association for Computational Linguistics, chapter CMU Haitian Creole-English Translation System for WMT 2011, pages 386–392. <http://aclweb.org/anthology/W11-2146>.
- Chang Hu, Philip Resnik, Yakov Kronrod, Vladimir Eidelman, Olivia Buzek, and B. Benjamin Bederson. 2011. *Proceedings of the Sixth Workshop on Statistical Machine Translation*, Association for Computational Linguistics, chapter The Value of Monolingual Crowdsourcing in a Real-World Translation Scenario: Simulation using Haitian Creole Emergency SMS Messages, pages 399–404. <http://aclweb.org/anthology/W11-2148>.
- Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional LSTM-CRF models for sequence tagging. *arXiv preprint abs/1508.01991*. <http://arxiv.org/abs/1508.01991>.
- Rebecca Hwa, Philip Resnik, Amy Weinberg, Clara Cabezas, and Okan Kolak. 2005. Bootstrapping parsers via syntactic projection across parallel texts. *Natural Language Engineering* 11(3):311–325. <https://doi.org/10.1017/S1351324905003840>.
- Eliyahu Kiperwasser and Yoav Goldberg. 2016. Simple and accurate dependency parsing using bidirectional lstm feature representations. *Transactions of the Association of Computational Linguistics* 4:313–327. <http://aclweb.org/anthology/Q16-1023>.
- Karen Lahousse and Béatrice Lamiroy. 2012. Word order in french, spanish and italian: A grammaticalization account. *Folia Linguistica* 46(2):387–415.
- Jakob R. E. Leimgruber. 2009. *Modelling variation in Singapore English*. Ph.D. thesis, Oxford University.
- Jakob R. E. Leimgruber. 2011. Singapore english. *Language and Linguistics Compass* 5(1):47–62. <https://doi.org/10.1111/j.1749-818X.2010.00262.x>.
- Jiwei Li, Thang Luong, Dan Jurafsky, and Eduard Hovy. 2015. When are tree structures necessary for deep learning of representations? In *Proceedings of the EMNLP 2015*. Association for Computational Linguistics, pages 2304–2314. <https://doi.org/10.18653/v1/D15-1278>.
- Lisa Lim. 2007. Mergers and acquisitions: on the ages and origins of singapore english particles. *World Englishes* 26(4):446–473.
- Héctor Martínez Alonso, Željko Agić, Barbara Plank, and Anders Søgaard. 2017. Parsing universal dependencies without training. In *Proceedings of the EACL 2017*. Association for Computational Linguistics, pages 230–240. <http://www.aclweb.org/anthology/E17-1022>.
- Ryan McDonald, Slav Petrov, and Keith Hall. 2011. Multi-source transfer of delexicalized dependency parsers. In *Proceedings of the EMNLP 2011*. Association for Computational Linguistics, pages 62–72. <http://aclweb.org/anthology/D11-1006>.
- Ho Mian-Lian and John T. Platt. 1993. *Dynamics of a contact continuum: Singaporean English*. Oxford University Press, USA.
- Makoto Miwa and Mohit Bansal. 2016. End-to-end relation extraction using lstms on sequences and tree structures. In *Proceedings of the ACL 2016*. Association for Computational Linguistics, pages 1105–1116. <https://doi.org/10.18653/v1/P16-1105>.
- Tahira Naseem, Regina Barzilay, and Amir Globerson. 2012. Selective sharing for multilingual dependency parsing. In *Proceedings of the ACL 2012*. Association for Computational Linguistics, pages 629–637. <http://aclweb.org/anthology/P12-1066>.
- Tahira Naseem, Harr Chen, Regina Barzilay, and Mark Johnson. 2010. Using universal linguistic knowledge to guide grammar induction. In *Proceedings of the EMNLP 2010*. Association for Computational Linguistics, Cambridge, MA, pages 1234–1244. <http://www.aclweb.org/anthology/D10-1120>.
- Paroo Nihilani. 1992. The international computerized corpus of english. *Words in a cultural context. Singapore: UniPress* pages 84–88.
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajic, Christopher D. Manning, Ryan McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, Reut Tsarfaty, and Daniel Zeman. 2016. Universal dependencies v1: A multilingual treebank collection. In *Proceedings of the LREC 2016*. European Language Resources Association.
- Joakim Nivre, Johan Hall, Sandra Kübler, Ryan McDonald, Jens Nilsson, Sebastian Riedel, and Deniz Yuret. 2007. The CoNLL 2007 shared task on dependency parsing. In *Proceedings of the CoNLL Shared Task Session of EMNLP-CoNLL 2007*. Association for Computational Linguistics, pages 915–932. <http://www.aclweb.org/anthology/D/D07/D07-1096>.
- Vincent B Y Ooi. 1997. *Analysing the Singapore ICE corpus for lexicographic evidence*. ENGLISH LANGUAGE & LITERATURE. <http://scholarbank.nus.edu.sg/handle/10635/133118>.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the EMNLP 2014*. Association for Computational Linguistics, pages 1532–1543. <https://doi.org/10.3115/v1/D14-1162>.

Barbara Plank, Anders Søgaard, and Yoav Goldberg. 2016. [Multilingual part-of-speech tagging with bidirectional long short-term memory models and auxiliary loss](#). In *Proceedings of the ACL 2016 (Short Papers)*. Association for Computational Linguistics, pages 412–418. <https://doi.org/10.18653/v1/P16-2067>.

Chun-Wei Seah, Hai Leong Chieu, Kian Ming Adam Chai, Loo-Nin Teow, and Lee Wei Yeong. 2015. [Troll detection by domain-adapting sentiment analysis](#). In *18th International Conference on Information Fusion (Fusion) 2015*. IEEE, pages 792–799.

Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, D. Christopher Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the EMNLP 2013*. Association for Computational Linguistics, pages 1631–1642. <http://aclweb.org/anthology/D13-1170>.

Anders Søgaard. 2012a. [Two baselines for unsupervised dependency parsing](#). In *Proceedings of the NAACL-HLT Workshop on the Induction of Linguistic Structure*. Association for Computational Linguistics, pages 81–83. <http://www.aclweb.org/anthology/W12-1910>.

Anders Søgaard. 2012b. [Unsupervised dependency parsing without training](#). *Natural Language Engineering* 18(2):187203. <https://doi.org/10.1017/S1351324912000022>.

Oscar Täckström, Ryan McDonald, and Jakob Uszkoreit. 2012. [Cross-lingual word clusters for direct transfer of linguistic structure](#). In *Proceedings of the NAACL-HLT 2012*. Association for Computational Linguistics, pages 477–487. <http://aclweb.org/anthology/N12-1052>.

David Weiss, Chris Alberti, Michael Collins, and Slav Petrov. 2015. [Structured training for neural network transition-based parsing](#). In *Proceedings of the ACL-IJCNLP 2015*. Association for Computational Linguistics, pages 323–333. <https://doi.org/10.3115/v1/P15-1032>.

Yuan Zhang and Regina Barzilay. 2015. [Hierarchical low-rank tensors for multilingual transfer parsing](#). In *Proceedings of the EMNLP 2015*. Association for Computational Linguistics, pages 1857–1867. <https://doi.org/10.18653/v1/D15-1213>.

Yuan Zhang and David Weiss. 2016. [Stack-propagation: Improved representation learning for syntax](#). In *Proceedings of the 54th ACL*. Association for Computational Linguistics, pages 1557–1566. <https://doi.org/10.18653/v1/P16-1147>.

Hao Zhou, Yue Zhang, Shujian Huang, and Jiajun Chen. 2015. [A neural probabilistic structured-prediction model for transition-based dependency parsing](#). In *Proceedings of the ACL-IJCNLP 2015*. Association for Computational Linguistics, pages 1213–1222. <https://doi.org/10.3115/v1/P15-1117>.

A Statistics of Singlish Dependency Treebank

POS Tags					
ADJ	782	INTJ	556	PUNCT	1604
ADP	490	NOUN	1779	SCONJ	126
ADV	941	NUM	153	SYM	11
AUX	429	PART	355	VERB	1704
CONJ	167	PRON	682	X	10
DET	387	PROPN	810		

Table A1: Statistics of POS tags

Dependency labels			
acl	37	dojb	612
acl:relcl	29	expl	10
advcl	194	iobj	15
advmod	859	list	10
appos	18	mwe	105
amod	423	name	117
aux	377	neg	261
auxpass	47	nmod	398
case	463	nmod:npmod	26
cc	167	nmod:poss	153
ccomp	138	nmod:tmod	81
compound	420	nsubj	1005
compound:prt	30	nsubjpass	34
conj	238	nummod	94
cop	152	mark	275
csubj	30	parataxis	241
det	304	punct	1607
det:predet	7	remnant	17
discourse	552	vocative	41
dislocated	2	xcomp	190

Table A2: Statistics of dependency labels

ah	aiyah	ba
hah / har / huh	hiak hiak hiak	hor
huat	la / lah	lau
leh	loh / lor	ma / mah
wahlow / wah lau	wa / wah	ya ya
walaneh / wah lan eh		

Table A3: List of discourse particles

A-B		
act blur	ah beng	ah ne
angpow	arowed	ang ku kueh
angmoh/ang moh	ahpek / ah peks	atas
boh/bo	boho jiak	boh pian
buay lin chu	buen kuey	
C		
chai tow kway	chao ah beng	chap chye png
char kway teow	chee cheong fun / che cheong fen	
cheesepie	cheong / chiong	chiam / cham
chiak liao bee / jiao liao bee		chio
ching chong	chio bu / chiobu	chui
chop chop	chow-angmoh	chwee kueh
D-F		
dey	diam diam	die kock standing
die pain pain	dun	eat grass
flip prata	fried beehoon	
G		
gahmen / garment	gam	geylang
gone case	gong kia	goreng pisang
gui		
H-J		
hai si lang	heng	hiong
hoot	Hosay / ho say	how lian
jepun kia / jepun kias		
jialat / jia lak / jia lat		
K		
ka	kaki kong kaki song	
kancheong	kateks	kautim
kay kiang	kayu	kee chia
kee siao	kelong	kena / kana
kiam	kiasu	ki seow
kkj	kong si mi	kopi
kopi lui	kopi-o	kosong
koyok	ku ku bird	
L		
lagi	lai liao	laksa
lao jio kong	lao sai	lau chwee nua
liao / ler	like dat / like that	
lim peh	lobang	
M		
mahjong kaki	makan	masak masak
mati	mee	mee pok
mee rebus	mee siam	mee sua
mei mei		
N-S		
nasi lemak	pang sai	piak
sabo	sai	same same
sia	sianz / sian	sia suay
sibeh	siew dai	siew siew dai
simi taisee	soon kuey	sotong
suay / suey	swee	
T		
tahan	tak pakai	te te kee
tong	tua	tikopeh
tio	tio pian/dio pian	
talk cock / talk cock	sing song	
U-Z		
umm zai	up lorry / up one's lorry	
xiao	zhun / buay zhun	

Table A4: List of imported vocabularies