From Genesis to Creole Language: Transfer Learning for Singlish Universal Dependencies Parsing and POS Tagging

HONGMIN WANG, University of California Santa Barbara, USA

JIE YANG, Brigham and Women's Hospital & Harvard Medical School, USA

YUE ZHANG, School of Engineering, Westlake University, China

Singlish can be interesting to the computational linguistics community both linguistically, as a major low-resource creole based on English, and computationally, for information extraction and sentiment analysis of regional social media. In our conference paper, Wang et al. (2017), we investigated part-of-speech (POS) tagging and dependency parsing for Singlish by constructing a treebank under the Universal Dependencies scheme and successfully used neural stacking models to integrate English syntactic knowledge for boosting Singlish POS tagging and dependency parsing, achieving the state-of-the-art accuracies of 89.50% and 84.47% for Singlish POS tagging and dependency, respectively. In this work, we substantially extend Wang et al. (2017) by enlarging the Singlish treebank to more than triple the size and with much more diversity in topics, as well as further exploring neural multi-task models for integrating English syntactic knowledge. Results show that the enlarged treebank has achieved significant relative error reduction of 45.8% and 15.5% on the base model, 27% and 10% on the neural multi-task model, and 21% and 15% on the neural stacking model for POS tagging and dependency parsing, respectively. Moreover, the state-of-the-art Singlish POS tagging and dependency parsing accuracies have been improved to 91.16% and 85.57%, respectively. We make our treebanks and models available for further research.

CCS Concepts: • **Computing methodologies** → **Artificial intelligence**; **Natural language processing**; **Language resources**;

Additional Key Words and Phrases: Dependency parsing, universal dependencies, part-of-speech tagging, transfer learning, neural stacking, multi-task network, creole language, Singlish

ACM Reference format:

Hongmin Wang, Jie Yang, and Yue Zhang. 2019. From Genesis to Creole Language: Transfer Learning for Singlish Universal Dependencies Parsing and POS Tagging. *ACM Trans. Asian Low-Resour. Lang. Inf. Process.* 19, 1, Article 1 (May 2019), 29 pages.

<https://doi.org/10.1145/3321128>

1 INTRODUCTION

Languages evolve temporally and geographically, both in vocabulary as well as in syntactic structures. When major languages such as English or French are adopted in another culture as the

Authors' addresses: H. Wang, University of California Santa Barbara, 2104 Harold Frank Hall, Santa Barbara, CA, 93106-5110; email: hongmin_wang@cs.ucsb.edu; J. Yang, Brigham and Women's Hospital & Harvard Medical School, 75 Francis street, Boston, MA 02115, USA; email: jieynlp@gmail.com; Y. Zhang (corresponding author), Westlake University, Institute for Advanced Study, 310024, 18 Shilongshan Road, Hangzhou, Zhejiang, China; email: yue.zhang@wias.org.cn.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Association for Computing Machinery.

2375-4699/2019/05-ART1 \$15.00

<https://doi.org/10.1145/3321128>

primary language, they often mix with existing languages or dialects in that culture and evolve into a stable language called a creole. Examples of creoles include the French-based Haitian Creole and Colloquial Singaporean English (Singlish) (Mian-Lian and Platt 1993), an English-based creole. While the majority of the natural language processing (NLP) research attention has been focused on the major languages, little work has been done on adapting the components to creoles. One notable body of work originated from the featured translation task of the EMNLP 2011 Workshop on Statistical Machine Translation (WMT11) to translate Haitian Creole SMS messages sent during the 2010 Haitian earthquake. This work highlights the importance of NLP tools on creoles in crisis situations for emergency relief (Hewavitharana et al. 2011; Hu et al. 2011).

Singlish is one of the major languages in Singapore, with borrowed vocabulary and grammars¹ from a number of languages, including Malay, Tamil, and Chinese dialects such as Hokkien, Cantonese, and Teochew (Leimgruber 2009, 2011), and it has been increasingly used in written forms on web media. However, fluent English speakers unfamiliar with Singlish would find the creole hard to comprehend (Harada 2009). Correspondingly, fundamental English NLP components such as POS taggers and dependency parsers perform poorly on such Singlish texts as shown in Tables 2 and 4. One other example is that Seah et al. (2015) adapted the Socher et al. (2013) sentiment analysis engine to the Singlish vocabulary but failed to adapt the parser. Since dependency parsers are important for tasks such as information extraction (Miwa and Bansal 2016) and discourse parsing (Li et al. 2015), this hinders the development of such downstream applications for Singlish in written forms and thus makes it crucial to build a dependency parser that can perform well natively on Singlish.

To address this issue, we took the first attempt starting with investigating the linguistic characteristics of Singlish and specifically the causes of difficulties for understanding Singlish with English syntax. We found that, despite the obvious attribute of inheriting a large portion of basic vocabularies and grammars from English, Singlish not only imports terms from regional languages and dialects, its lexical semantics and syntax also deviate significantly from English (Leimgruber 2009, 2011). We further categorized the challenges and formalized their interpretation using Universal Dependencies (Nivre et al. 2016) and created a Singlish dependency treebank with 1,200 sentences.

Based on the intricate relationship between Singlish and English, we built a Singlish parser by leveraging knowledge of English syntax as a basis. This overall approach is illustrated in Figure 1(a). In particular, we trained a basic Singlish parser with one of the best off-the-shelf neural dependency parsing model using biaffine attention (Dozat and Manning 2017) and improved 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 also trained a POS tagger for Singlish following the same idea by integrating English POS knowledge using neural stacking.

Results by our conference paper 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).

In this article, we further extend the work by substantially enlarging the Singlish dependency treebank to more than triple the size, using data with much more diversity in terms of topics and time span from multiple local Internet forums. We show that this leads to more substantiations in Singlish syntactic constructions as illustrations by case analyzes in Section 3.4 and a summarization of all Singlish terms and their meanings in Tables B4 to B6 in Appendix B.

¹We follow Leimgruber (2011) in using “grammar” to describe “syntactic constructions,” and we do not differentiate between the two expressions in this article.

Fig. 1. Overall model diagrams: (a) Neural Stacking Model and (b) Multi-task Model.

Furthermore, we explore neural multi-task models for transferring English knowledge to Singlish dependency parsing, as shown in Figure 1(b), which is done by joint training of a Singlish and an English parser with a shared module to enable knowledge sharing. The same has been applied to POS tagging, and it has achieved significant improvement over the base models for both tasks.

As a significantly extended version of our conference work, we show in this article that the current state-of-the-art Singlish dependency parsing has been improved to 85.57% UAS and 79.12% LAS and POS tagging accuracy has been improved from 89.50% to 91.16%. 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 available² to facilitate future research.

The article is organized as follows: Section 2 discusses the related works in the areas of transfer learning for cross-lingual parsing and neural stacking and multi-task models, Section 3 demonstrates the adaption of Universal Dependencies to Singlish during the dataset construction and extension processes, Section 4 describes the base, neural stacking, and multi-task models for POS tagging and experiment results using the original and extended treebanks, Section 5 describes the different models for Singlish parsing followed by discussions on various experiments in Section 6, and, finally, Section 7 concludes the article.

2 RELATED WORK

Neural networks have led to significant performance improvement for dependency parsing, including transition-based parsing (Andor et al. 2016; Ballesteros et al. 2015; Chen and Manning 2014; Dyer et al. 2015; Weiss et al. 2015; Zhou et al. 2015) and graph-based parsing (Dozat and Manning 2017; Kiperwasser and Goldberg 2016). 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 one of the highest-performing graph-based English dependency parsers. 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 (Cohen and Smith 2009; Ganchev et al. 2009; Hwa et al. 2005). Seminal work employed statistical models. McDonald et al. (2011) investigated delexicalized transfer, where word-based

²https://github.com/wanghm92/Sing_Par.

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 (Naseem et al. 2012; Täckström et al. 2012; Zhang and Barzilay 2015).

Recently, a line of work leverages neural network models for multi-lingual parsing (Ammar et al. 2016; Duong et al. 2015; Guo et al. 2015). 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 and a multi-task learning model, which respect the distributional differences between Singlish and English words. This empirically gives higher accuracies for Singlish.

Neural stacking and neural multi-task models were previously used for cross-annotation (Chen et al. 2016) and cross-task (Zhang and Weiss 2016) joint-modeling on monolingual treebanks. We were the first to employ neural stacking on cross-lingual feature transfer from resource-rich languages to improve dependency parsing for low-resource languages and we further explore neural multi-task models on this task in this work. Besides these three dimensions in dealing with heterogeneous text data, another popular area of research is on the topic of domain adaption, 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 detail.

Unsupervised rule-based approaches also offer a competitive alternative for cross-lingual dependency parsing (Gelling et al. 2012; Gillenwater et al. 2010; Martínez Alonso et al. 2017; Naseem et al. 2010; Søgaard 2012a, 2012b) and recently have 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 adaption (Martínez Alonso et al. 2017). However, we choose a data-driven supervised approach given the relatively higher parsing accuracy due 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 76 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 corpus in Universal Dependencies v2.3³ has six treebanks from different domains, and the main collection is constructed from the English Web Treebank (Bies et al. 2012), comprising web

³As of April 16, 2019, <http://universaldependencies.org/>.

Table 1. Division of Training, Development, and Test Sets for the Singlish Treebank STB-ACL

	UD English		Singlish		Extension	
	Sentences	Words	Sentences	Words	Sentences	Words
Train	12,543	204,586	900	8,221	3,050	27,368
Dev	2,002	25,148	150	1,384	150	1,384
Test	2,077	25,096	150	1,381	150	1,381

media texts such as blogs and online reviews, which potentially smooths the knowledge transfer to our target Singlish texts in similar domains. However, to be consistent with our previous work in terms of treebank extension and fair comparison for the experiments, we still adopt the annotation standards compatible with the v1.4 Universal Dependencies English treebanks⁴ and leave the version conversion to future works. 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 that collectively form five 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 means “*very anxious not to miss an opportunity*” in Hokkien,⁵ and “*makan*” of (1) in Figure 2, which means “*eat*” in Malay. 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. Tables B2 and B3 in Appendix B summarize all borrowed terms in the first version of the Singlish treebank. Tables B4 to B6 in Appendix B summarize all borrowed terms as well as their meanings looked up from the Singlish resources listed in Section 3.3 and the *Urban Dictionary*.⁶

Topic-prominence: This type of sentence starts 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 (2) and “*Earn some kopi money*” of (3) in Figure 2, which is consistent with the dependency relations in its Chinese translation.

⁴ Available at <https://lindat.mff.cuni.cz/repository/xmlui/handle/11234/1-1827>.

⁵ Definition by the *Oxford Living Dictionaries for English*.

⁶ <https://www.urbandictionary.com/>.

Fig. 2. Unique Singlish grammars. Arcs represent dependencies, pointing from the head to the dependent and the label on top. POS tags are below the words. (English translations available at Appendix C).

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 (4) 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 (5) 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 (6) in Figure 2. Copula deletion is also broadly shared in social media texts such as Tweets (Liu et al. 2018; O’Connor et al. 2018; Sanguinetti et al. 2018, 2017), which arises along with the prosperity of social media in recent two decades. The purpose of dropping components in social text media is mainly for conciseness of expression. However, the copula deletion phenomenon in Singlish mainly results from environmental influence (Leimgruber 2009), such as mixing Chinese dialects with English, which has been prevailing throughout its history.

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 a 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,*”⁷ 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 the possibly missing “*it*” between “*Hope*” and “*can*” in (7) and the “*you*” between “*Why*” and “*dun*”(don’t) in (8) 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 (Lahousse and Lamiroy 2012). This simply involves a change of word orders and thus requires no special treatments. However, 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.⁸ The other type is exemplified as “*right?*,” and its Singlish equivalent “*tio boh?*” (a transliteration from Hokkien) is labeled with the “*discourse*” (discourse element) relation with respect to the sentence root. See example (9) in Figure 2.

⁷In English: (I) would now like to enter briefly on the merit (of the discussion).

⁸In 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.

Discourse particles: Usage of clausal-final discourse particles, which originates from Hokkien and Cantonese, is one of the most typical features of Singlish (Leimgruber 2009, 2011; Lim 2007). All discourse particles that appear in our treebank are summarized in Table B1 in Appendix B 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 shown by the “*leh*” of (5) and (8) 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 (10) in Figure 2.

3.3 Data Selection and Annotation

On top of the released first version of the Singlish treebank, we extend it to more than triple the size with much more diversity in terms of topics and time span.

Data Source: Singlish is used in written form mainly in social media and local Internet forums. After comparison, we first chose the *SG Talk Forum*⁹ as the data source due to its relative abundance in Singlish contents. Specifically, they crawled¹⁰ 84,459 posts using the *Scrapy* framework¹¹ from pages dated to December 25, 2016, retaining sentences of length between 5 and 50, which total 58,310. Sentences were reversely sorted according to the log-likelihood of the sentence given by an English language model trained using the *KenLM* toolkit¹² (Heafield et al. 2013) 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*,¹³ *A Dictionary of Singlish and Singapore English*,¹⁴ and the *Singlish Vocabulary* Wikipedia page.¹⁵ 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. Finally, 1,200 sentences were chosen 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,¹⁶ and then annotated using the Dependency Viewer.¹⁷ 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. We name the dataset **STB-ACL**.

⁹<http://sgTalk.com>.

¹⁰Using *Python* package *beatifulsoup4* available at: <https://pypi.python.org/pypi/beautifulsoup4>.

¹¹<https://scrapy.org/>.

¹²Trained using the *afp_eng* and *xin_eng* sources of English Gigaword Fifth Edition (*Gigaword*).

¹³<http://72.5.72.93/html/lexec.php>.

¹⁴<http://www.singlishdictionary.com>.

¹⁵https://en.wikipedia.org/wiki/Singlish_vocabulary.

¹⁶<http://www.nltk.org/api/nltk.tokenize.html>.

¹⁷<http://nlp.nju.edu.cn/tanggc/tools/DependencyViewer.exe>.

Extension: With the success on improving the Singlish dependency parsing accuracy using the first released version of the treebank, we take a step further to extend the treebank comparable to the size of a UD treebank for the lower source languages. Specifically, we further apply the same procedures described above using all three major and popular forums, *SG Talk Forum*, *Hardwarezone*,¹⁸ and *SgForums*,¹⁹ by crawling all textual posts on all three forums dated until June 19, 2017. This brings two major differences in the source of data for the extended dataset:

- **Size:** The data source totals 420 million words, which is more than 10 times larger.
- **Diversity:** The three forums cover 10 times more topics. In particular, *Hardwarezone* has more than 46 sub-forums cover topics including technology events, information technology diagnosis, digital entertainment lifestyle, and and so on, *SgForums* is a general open-domain forum focusing more on daily news and casual topics.

We remove about 1% contents written in non-English characters and then tokenize and split by sentences using the Stanford CoreNLP toolkit²⁰ (Manning et al. 2014). The same selection procedure has been applied to the crawled texts to select additional 2,150 sentences with 20,000 words, which are then annotated according to the same adaption of Universal Dependencies on Singlish. This extended treebank enables more insightful syntactic analysis on Singlish-specific grammars as illustrated in Section 3.4 and helps to further boost the performance for various Singlish POS tagger and dependency parser models as shown in Sections 4 and 5. We name this extended dataset **STB-EXT**.

3.4 Case Studies on Singlish-specific Syntax

With the major extension of the Singlish treebank in both size and diversity in contents, more Singlish-specific syntactic constructions are observed and substantiated by increasing frequencies. A full list of extended imported vocabularies and Singlish expressions and their meanings are tabulated in Tables B4 to B6 in Appendix B. In this sections, we specifically discuss 6 unique Singlish syntactic constructions that are more commonly observed in the extended Singlish treebank and the adaption of Universal Dependencies on them:

- **“de”** and **“le”** are the Hanyu Pinyin Romanization(Abbreviated to: Pinyin) of common interjections in Chinese,²¹ which appears often in Singlish texts due its Chinese origin. Accordingly, they have the same POS tags, “X” or “PART” depending on their specific usage, and the same dependency relations as exemplified in the UD Chinese corpus.
- **“kenna”** is a unique word in Singlish that carries multiple senses “to get” or “going to.” The former serves as an auxiliary verb and is followed by a verb. Thus, it is annotated with the “AUX” POS tag and usually depend on the verb with the “auxpass” dependency relation. The latter serves as a verb when followed by the nominal object, and thus it has the “VERB” POS tag and often is the root of the sentence or the clause.
- **“tio”** is another unique word in Singlish that also carries multiple senses summarized in three main ways. First, if it means “to get” or “accomplish,” then it is equivalent to “kenna” and is thus annotated in the same way. Second, if it means “accurately choose,” then it is treated as a verb. Last, in the fixed expression “tio boh ?,” as mentioned before, it is analogous to “isn’t it?” in English. Thus, in this case, it is treated as an interjection word used as a discourse particle at the end of a sentence.

¹⁸<http://forums.hardwarezone.com.sg/>.

¹⁹<http://sgforums.com/forums>.

²⁰<https://stanfordnlp.github.io/CoreNLP/>.

²¹“de”: 的; “le”: 了.

Fig. 3. Base POS tagger.

- “**no need**” is a commonly used phrase in Singlish, where the POS of the word “*need*” is hard to determine to be either “*VERB*” or “*NOUN*,” since “*no*” in Singlish is often interchangeable with “*not*,” especially when used at the start of a sentence or clause where no proceeding context is present to offer necessary syntactic clues. Thus, we establish a rule by differentiating its usage in two cases. First, if “*no need*” is followed by a verb, then we consider it as a case where the particle “*to*” is considered as being omitted, and the verb is the head of an adjective clause modifying “*need*” as a “*NOUN*.” In this case, “*no*” has the POS tag of “*DET*.” However, if “*no need*” is followed by a noun, then “*need*” is naturally considered as a verb, and “*no*” has a POS tag of “*PART*,” equivalent to “*not*,” which is used more often in standard English.
- “**not**” as a common English word is extensively used in Singlish not accompanied with any auxiliary or copular verb. Some examples are “*Not ...*,” “*If not, ...*,” “*If not ... , ...*,” and the most typical “*Confirm not ...*”²² expression. According to the usages in English, we have summarized the rule of thumb as: If “*not*” is used with an auxiliary verb, then it has a POS of “*PART*.” Otherwise, when used independently, it has a POS of “*ADV*,” with exceptions when used in “*If not, ...*” and “*or not*.” In both cases, it modifies its head with a “*advcl*” dependency relation. The corresponding word in Singlish is “**bo**” or “**boh**” and they are annotated in the same way according to this rule.
- Adjectives and adverbs are often used with a **duplication** as a tone intensifier in Singlish. For example, in the sentence “*Aiya, we are comparing what are cheap cheap mah*,” “*cheap*” is duplicated to by the speaker to emphasize on such expression. This is unique in Singlish, and the closest dependency relation can be applied is “*mwe*” for multi-word expression. However, we propose that it can be a unique language-specific construction in Singlish and possibly named as “*mwe: dup*.”

4 PART-OF-SPEECH TAGGING

As previously mentioned, since POS tags are important features for dependency parsing, we first obtain the automatically predicted POS tags. Specifically, we investigate two approaches (i.e., neural stacking and neural multi-task learning), which leverage both the English and Singlish POS tagging datasets. As shown in Figure 3, the bi-LSTM network with a CRF layer (bi-LSTM-CRF) is chosen as the base model, it has shown state-of-the-art performance by globally optimizing the tag sequence (Chen et al. 2016; Huang et al. 2015). For the neural stacking approach, we train a POS tagger for UD-Eng using the base model. 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

²²Equivalent to: *I’m sure that ... is/do not*.

Fig. 4. Neural stacking POS tagger.

the English and Singlish models consist of an input layer, a feature layer, and an output layer. For the neural multi-task structure shown in Figure 5, the input layer and feature layer are shared by both the English and Singlish models, and each model has distinct output layers. Same with the multi-task model of Chen et al. (2016), model parameters are trained by both tasks.

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: As shown in Figure 3, this is a linear layer over the feature layer, followed by 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. The hidden state vectors from the linear layer are called the emission vectors.

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 vectors produced by the English Tagger with the original word and character-based embeddings, before applying the concatenation within a context window described in Section 4.1. During training, the loss is back-propagated to all trainable parameters in both the Singlish Tagger and the pre-trained layers of the base English Tagger. At test time, the input sentence is fed to the whole integrated tagger model for inference.

4.3 POS Tagger with Neural Multi-task Learning

Following Chen et al. (2016), we investigate the multi-task learning for Singlish POS tagging. Figure 5 shows the multi-task structure, the English tagger and Singlish tagger share the same

Fig. 5. Multi-task POS tagger.

input layer and feature layer but use different output layers. In the training phase, sentences from both datasets are both fed into the same input layer and feature layer but flow through the respective output layer of each task to compute the score of label sequences. The loss of each POS tagger is back-propagated to the output layer of the corresponding task and the shared feature layer and the input layer. The testing stage is similar to the base tagger but only the tagger output layer corresponding to the language of the input sentence is used.

4.4 Results

Base English POS tagger: we followed Chen et al. (2016) and used the publicly available source code²³ 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). The hidden layer size was set 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). These settings were used to perform 10-fold jackknifing of POS tagging on the UD-Eng training set, with an average accuracy of 95.60%. This English POS tagger is used as the base English tagger in the neural stacking model described in Section 4.2.

Base Singlish POS tagger: Similarly, we trained a base Singlish POS tagger using the Singlish 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 Singlish texts. Due to the limited amount of training data, the POS tagging accuracy on **STB-ACL** is significantly lower compared to all other models even with a larger dropout rate to avoid over-fitting during experiments, as shown by *Base-ICE-SIN* in Table 2.²⁴ However, the POS tagging accuracy is significantly improved from 78.35% to 88.27% using **STB-EXT** with 45.82% relative error reduction, which shows the effectiveness of the increased scale of training data.

Neural stacking: The neural stacking model has achieved the best accuracy of 89.50% using the **STB-ACL**,²⁵ as is shown by *Stack-ICE-SIN* in Table 2, which corresponds to a 51.50%

²³<https://github.com/chenhongshen/NNHetSeq>.

²⁴In order for fair comparison with the results trained on **STB-ACL** and also to compare the relative amounts of improvement across different models, we restrict all our experiments to the same dev and test sets despite their relative small size. We leave for future explorations to perform significance test and experiments with the different dataset splits on **STB-EXT**.

²⁵We 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.

Table 2. POS Tagging Accuracies (%) for Models Trained on **STB-ACL** and **STB-EXT**

System	STB-ACL	STB-EXT
ENG-on-SIN	81.39	81.39
Base-ICE-SIN	78.35	88.27
Base-ICE-ENG+SIN	87.76	89.72
Stack-ICE-SIN	89.50	90.73
Multitask-ICE-SIN	87.83	91.16
Integrated-ICE-SIN	87.91	91.09

relative error reduction over its baseline Singlish model. Similarly, when training the neural stacking model on the **STB-EXT**, the POS tagging accuracy is further enhanced by 1.23% absolute accuracy improvement and 11.71% relative error reduction over its baseline Singlish POS tagger.

Multi-task model: The multi-task models can also improve the POS tagger on both **STB-ACL** and **STB-EXT** significantly, as shown by *Multitask-ICE-SIN* in Table 2. Comparing with the neural stacking model, the multi-task model has lower accuracy when trained on **STB-ACL** but outperforms the neural stacking model on **STB-EXT**.

This shows that the multi-task model benefits more from the extension of training data even with only one shared input layer compared to separate input layers in the neural stacking model. The reason may be that the **STB-EXT** dataset can provide sufficient samples for the multi-task network structure to extract more transferable features between English and Singlish, which leads to a better performance.²⁶

We also experimented with the integrated model by Chen et al. (2016) that utilizes both neural stacking and multi-task structures, as shown by *Integrated-ICE-SIN* in Table 2. However, it does not achieve further improvements, which is different from the observation by Chen et al. (2016). Moreover, simply combining the UD-Eng and **STB-ACL** or **STB-EXT** datasets and training base POS tagger models yields improvements over the two baseline models train on a single language but still under-performs the stacking and multi-task models on both STBs. This again shows the availability of transferable knowledge from English to Singlish and also the better effectiveness of the transfer learning models.

To produce the automatically predicted POS tags for all train, development, and test datasets, we use the best POS tagger, *Stack-ICE-SIN*, on **STB-ACL** and *Multitask-ICE-SIN* on **STB-EXT** to perform fivefold jackknifing on the training sets and label the development and test dataset, respectively.

²⁶To further prove this consistency of improvement brought by dataset extension, we pre-trained another set of word embeddings using all the 420-million-word raw texts crawled from the Internet forums, called *Forum-SIN*. We used it with the the multi-task model and it leads to further improvements of Singlish POS tagging accuracy to 88.83% using the **STB-ACL** dataset and 91.45% using the **STB-EXT** dataset, respectively. The *Forum-SIN* embeddings to a certain extent compensates the domain difference between the *ICE-SIN* embeddings and the Singlish dependency treebanks, which is more significant for the smaller **STB-ACL** dataset, as shown by the difference scale of improvement over *Multitask-ICE-SIN*: 1% for **STB-ACL** but 0.29% for **STB-EXT**. However, the *Forum-SIN* embeddings trained on forum texts are less concentrated on Singlish contents compared to *ICE-SIN*, which was manually constructed with strict content control. However, the *Forum-SIN* is advantageous over the *ICE-SIN* embeddings by not only bridging the domain differences but also incorporating the context information of all texts annotated in the two treebanks during pre-training.

Therefore, to avoid the extra variety and preserve the fair comparison between only different models and datasets regarding knowledge transfer between language syntactic constructions, we make this set of embeddings available and leave further investigation on domain adaption to future work.

Fig. 6. Base parser model.

5 DEPENDENCY PARSING

We adopt the Dozat and Manning (2017) parser²⁷ as our base model, as displayed in Figure 6, and apply neural stacking and neural multi-task learning 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. However, the multi-task model consist of shared input and feature layers but has distinct output layers for the English and Singlish parsers, respectively.

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) layers. Following Dozat and Manning (2017), we adopt the Cif-LSTM cells by Greff et al. (2015).

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. 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 7, and the scoring tensor of the Singlish Parser is initialized with the one from

²⁷<https://github.com/tdozat/Parser>.

Fig. 7. Parser with neural stacking.

Fig. 8. Parser with multitask structure.

the trained English Tagger. The 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.

5.3 Parser with Neural Multi-task Learning

Similarly to the multi-task POS tagging structure, we also investigate the parser with neural multi-task joint training with English corpus. Figure 8 shows the structure of neural multi-task parser. However, since the multi-layer bi-LSTM network serves as the general feature extractor that captures the contextual information from for each input word by combining the hidden state vectors from both directions through concatenation, the task-specific layers start from the set of MLP

Table 3. Comparison of the Scale of Sources for Training Word Embeddings

	Sentences	Words	Vocabulary
GloVe6B	N.A.	6,000m	400,000
Giga100M	57,000	1.26m	54,554
ICE-SIN	87,084	1.26m	40,532

layers that maps the feature vectors in different lower-dimension spaces for different tasks (in this case English and Singlish parsers, respectively). This is then followed by distinct output layers for constructing different parse trees for English and Singlish, respectively. This network structure is different from the neural stacking model described in Section 5.2 in two ways. The first is that the input layer and the feature layer is shared in the multi-task model instead of separate for the neural stacking model, which aims to capture the common syntactic features shared by both English and Singlish. Second, the model is jointly trained using both the English and Singlish treebanks instead of pre-training on the English base parser model and loaded to the base components in the neural stacking model. This aims to provide regularization to the input and feature layers that encode the transferable syntactic feature between English and Singlish from training on the English treebank while the Singlish-specific output layers are trained using the Singlish treebank.

6 EXPERIMENTS

6.1 Experimental Settings

We train a base English parser on UD-Eng with the default model settings in Dozat and Manning (2017). It achieves a 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 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 one-layer bi-LSTM with 900 hidden dimension performs the best on the **STB-ACL** dataset and two-layer bi-LSTM with 600 hidden dimension performs the best on the **STB-EXT** dataset. The bigger hidden layers accommodate the elongated input vector to the stacked bi-LSTM and the fewer number of recurrent layers helps to avoid over-fitting on the relatively small Singlish dependency treebanks compared to UD-Eng given the deep bi-LSTM English parser network at the bottom. The evaluation of the neural stacking and multi-task models is further described in below.

6.2 Investigating Distributed Lexical Characteristics

To learn characteristics of distributed lexical semantics for Singlish, we compare performances of the base Singlish dependency parser trained on **STB-ACL** using several sets of pre-trained word embeddings with the raw text sources summarized in Table 3: *GloVe6B*, large-scale English word embeddings,²⁸ and *ICE-SIN*, Singlish word embeddings trained using GloVe (Pennington et al.

²⁸Trained with Wikipedia 2014 the Gigaword. Downloadable from <http://nlp.stanford.edu/data/glove.6B.zip>.

Table 4. Dependency Parser Performances Measured by UAS and LAS for Models Trained on English Only, Singlish Only, and Both English and Singlish

Trained on English				
System	UAS		LAS	
ENG-on-SIN	75.89		65.62	

Trained on Singlish				
System	STB-ACL		STB-EXT	
	UAS	LAS	UAS	LAS
Baseline	75.98	66.55	79.71	71.99
Base-Giga100M	77.67	67.23	—	—
Base-GloVe6B	78.18	68.51	83.02	75.21
Base-ICE-SIN	79.29	69.27	82.51	75.30

Trained on Singlish and English				
System	STB-ACL		STB-EXT	
	UAS	LAS	UAS	LAS
ENG-plus-SIN	82.43	75.64	83.28	75.81
Stack-ICE-SIN	84.47	77.76	85.57	79.12
Multitask-ICE-SIN	79.12	70.46	82.60	76.06

Two Singlish training sets are used: **STB-ACL** and **STB-EXT**.

2014) on the ICE-SIN (Nihilani 1992; Ooi 1997) corpus. These two sets of embeddings capture the distributional semantics from English and Singlish, respectively. However, due to the significant difference in the size of the corpus used for training the embeddings, in order for a fair comparison, we have trained another set of embeddings-*Giga100M*, a small-scale English word embeddings trained using GloVe (Pennington et al. 2014) with the same settings as the **ICE-SIN** embeddings on a comparable size of English data randomly selected from the English Gigaword Fifth Edition.

As shown in Table 4, 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.

Similar improvements has been observed for **STB-EXT**, as shown in Table 4, and both *ICE-SIN* and *GloVe6B* embeddings have lead to significant improvement over the baseline model. Furthermore, the Singlish *ICE-SIN* embeddings still yield comparable performance to the English *GloVe6B* embeddings with slightly lower UAS and slightly higher LAS on **STB-EXT**. This indicates that the distributional semantics from both the English and Singlish are beneficial for Singlish parsing.

The *Multi-task-SIN* yields lower performance on **STB-ACL** but comparable results with *ENG-plus-SIN* as shown in Table 4. With the additional output layer and corresponding parameters, the

Table 5. Dependency Parser
Performances by the
Five-cross-fold Validation

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

Table 6. Error Analysis with Respect to Grammar Types

	Topic Prominence		Copula Deletion		NP Deletion		Discourse Particles		Others	
Percentage	8.65%		10.98%		12.14%		29.48%		38.73%	
	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

multi-task model confronts more limitations using the smaller **STB-ACL** but on the other hand is better at leveraging more training data to improve the performance.²⁹

6.3 Knowledge Transfer Using Neural Stacking

We train a parser with the neural stacking model and Singlish *ICE-SIN* embeddings on the **STB-ACL** dataset, 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 comparable baseline, *Base-ICE-SIN*. 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 treebank, represented as *ENG-plus-SIN* in Table 4, which is still outperformed by the neural stacking model.

Besides, we performed a five-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 five 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.

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 categories³⁰ using **STB-ACL**, including four types of grammars mentioned in Section 3.2,³¹ and one 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.

²⁹We tried with different corpus weighting ratio between UD-Eng and empirically found 1:1 works the best.

³⁰The percentages add up to more than 100%, since multiple labels are allowed for one sentence.

³¹The “Inversion” type of grammar is not analyzed, since there is only one such sentence in the test set.

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 four 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 to model Singlish syntactic knowledge due to the differences in distributed lexical semantics.

6.5 Knowledge Transfer Using Neural Multi-task Learning

As is shown by *Multitask-ICE-SIN* in Table 4, *Multitask-ICE-SIN* has significantly lower accuracies compared with *ENG-plus-SIN*, since **STB-ACL** is not large enough compared to UD_Eng for the parser to extract sufficient transferable syntactic knowledge between Singlish and English by means of multi-task learning. This is substantially mitigated by the enlarged size of the Singlish training data when using **STB-EXT**, which leads to comparable accuracies with *ENG-plus-SIN*. This demonstrates the effectiveness of the extended dataset and the ability of the multi-task network structure on better knowledge transfer compared with the simple data mixing approach on a larger training dataset.

However, the Singlish parser with the multi-task model trained on both the two datasets has lower accuracies than the parser with the neural stacking structure (*Stack-ICE-SIN*). The main reason is that the separate input layer in the neural stacking model is better at capturing language-specific syntactic information compared to the shared input layer jointly trained by two languages. This defect can be partially compensated by the enlargement of the training dataset but the performance is still not satisfactory. Another hypothesis is that the domain difference between the Singlish *ICE-SIN* embeddings and the Singlish dependency treebanks is magnified when using a shared input layer.

6.6 Significance of Dataset Extension

Figure 9 has illustrated that all parsers trained on the **STB-EXT** dataset yield better results,³² which has proven the quality of the **STB-EXT** dataset. Besides, their consistency in the ranking of the parsing accuracies further substantiated our findings in our conference article. Furthermore, we find that the improvement from training data extension is more significant when applying the multi-task model (LAS from 70.46% to 76.06% with 18.96% relative error reduction) than the neural stacking parser (LAS from 77.76% to 79.12% with 6.12% relative error reduction). This is consistent with our previous observation that the multi-task model benefits more from the extension of training data compared to the neural stacking model in POS tagging task.

By extending the size of the training data using **STB-EXT**, the neural stacking parser is also improved significantly with 17.50% relative error reduction compared to *Base-ICE-SIN* and with 13.70% relative error reduction compared to *ENG-plus-SIN*, as is shown by *Stack-ICE-SIN* in Table 4. This is consistent with the POS tagging accuracy improvement described in Section 4 and has led to the new SOTA on Singlish parsing accuracies with UAS 85.57 and LAS 79.12.

³²Performance measured by LAS has the same trend as UAS and is excluded for simplicity.

Fig. 9. Comparison on dependency parser performances between STB-ACL and STB-EXT.

7 CONCLUSION

We have investigated dependency parsing for Singlish, an important English-based creole language, through annotations of a Singlish dependency treebank with 30,986 words in total and building an enhanced parser by leveraging on knowledge transferred from a 7-times-bigger English treebank of Universal Dependencies. We demonstrate the effectiveness of our dataset extension by improvements in experiment results using all base, neural stacking and multi-task models. Besides, we specifically explored multi-task models that boost the Singlish POS tagging accuracy and dependency parsing performance by joint learning transferable features between English and Singlish. We release the extended Singlish dependency treebank, **STB-EXT**, the trained model and the source code for the parser with free public access. Possible future work includes expanding the investigation to other regional languages such as Malay and Indonesian, and designing of better feature selection mechanisms to assist transfer learning between English and its low-resource creole languages.

APPENDICES

A STATISTICS OF SINGLISH DEPENDENCY TREEBANK V1

Table A1. Statistics I: Number of Occurrence of All POS Tags in STB-ACL

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 A2. Statistics II: Number of Occurrence of All
Dependency Labels in STB-ACL

Dependency labels			
acl	37	dobj	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:npm	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

B IMPORTED VOCABULARIES OF THE SINGLISH DEPENDENCY TREEBANK

Table B1. List of Discourse Particles

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 B2. List of Imported Vocabularies (A-M)

A-B		
act blur	ah beng	ah ne
angpow	arrowed	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
chio	chiak liao bee / jiao liao bee	
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
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
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	

Table B3. List of Imported
Vocabularies (N-Z)

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 B4. List of Singlish Terms and Their Meanings (A-D)

Term	Meaning
ah beng(=ah lian)	bully
ah kong	grandpa
ah tiong	mainland Chinese immigrants
alamak	oh no!
atas	upstairs/high class
bak kua	dried meat
bak	meat
balek	return
balek kampung	go home
balik kampong	go away
bochap(=boh chup)	nonchalant/relaxed/ignore/things that one couldn't care less all the time
boleh	can/possible
bo pian	no choice
bo hua	not worth the trouble
buat bodoh	play dumb
buay(=bo)	no
buay song	not satisfied
chai(=chye)	vegetable
char bor	female/women
cham	pitifully disastrous
cheena	a person heavily influenced by Chinese culture
chenghu	government
ching chiong	Asian(mocking)
chim(=cheem)	profound/complicated
chiu	you
chio	pretty/handsome
chope	reserve a seat/book
chow	reprimand
chua pai	big shot
cow peh cow bu(=CPCB/KPKB)	cry father, cry mother
dao	arrogant and haughty
dulan	dislike/hate

Table B5. List of Singlish Terms and Their Meanings (G-R)

Term	Meaning
garang	fierce/powerful
gian	eager
gian png	to be eager for rice, desperately craving for something very common
gong	stunned/daze
gong kia	stupid fool
hamkanchan	to curse one's family
hao lian	arrogant
hiao(=hiau)	sexy
hiu(=hew)	interest
hiong	fierce/strong
ho say	alright
hokay	okay
ian	hot guy
jiak	eat
jin	very
kay	very
kam siah	thank
kia	children/person/afraid of
kia si lang	very frightening
kong si mi	what are you saying
kong	speak
kns	just like shit
kua	dried food
lagi	additional/more
lang	person
lao kwee	embarrassment
liek(=liak) boh kiu(=kiew)	do not understand
lim peh	your father
lobang	hole/opening/contacts/tip/connections
longkong	drain
mah si(=poon)	also/even
mai	don't
mai tu liao	do not waste time
malu	humiliating
masuk	enter
mee	noodle
mensus	boss/king
moi	me
nabeh	shit/fuck/damn
nia	only
not bat an eyelid	without blinking
otang	owe someone money
pang sai	defecate
pang seh	be stood up with appointment
pai	bad
paiseh	a sense of embarrassment in the most discreet way
pang/pun	put
pok	attempting a go
pok kai	go die in the street
rast	bad

Table B6. List of Singlish Terms and Their Meanings (S–Z)

Term	Meaning
sabo	sabotage
sala(=salah)	wrong
see heng suay	see how it goes
skali(=skarly)	once
si	die
si beh(=sibei/sibeh)	very
sia	ready
siam	avoid/get lost/get out of the way of
sian	boring
siao(=seow)	mad/crazy
siao kia	crazy fool
sia suay(=siasway)	to embarrass
simi	what
sin-seh	Chinese medicine clinic
song	satisfied
suan	ridicule/insult
suay	extremely unlucky
suave	charming
swee(=sui)	beautiful
tahan	endure
taiji	matters/issues
tia(=tia gong)	listen/hearsay
tidak	not
tio	correct/got it
toh	collapse/fail
tock(=tok) kong	superb
up the lorry	bad outcome
wayang	performance on the screen
yao siew/siu	short life

C EXPLANATION IN ENGLISH FOR EXAMPLES IN FIGURE 3.4

- (1) If you are staying at CityView, you must go and eat
 - CityView: a condominium in Singapore
 - makan: eat, have meal
 - This sentence is to be understood in the context, where the posts around it were talking about some kind of food nearby CityView and this post is suggesting to try it out if you are living nearby
- (2) It'll definitely be eyecatching if you are driving this car
- (3) It is already good enough to earn some pocket money
 - kopi: coffee. coffee money can be regarded as pocket money for daily use
- (4) Are they any thatched houses in Singapore?
 - SG: short for Singapore
 - attap chu: attap means thatch³³, chu means house

³³Thatching is the craft of building a roof with dry vegetation.

- (5) You cannot see if you are inside the tent!
 - leh: a discourse particle
- (6) You are betting more downside from here?
 - U: short for you
- (7) Hope it can close at 22 today
- (8) Why don't you try storage?
 - dun: don't
- (9) It's best to eat it all, isn't it?
 - tio boh: isn't it right?
- (10) I never get it for free !
 - one: discourse particle in this case

ACKNOWLEDGMENTS

We are very grateful for the detailed and constructive comments from all three reviewers, which are helpful to make this article better.

REFERENCES

- Waleed Ammar, George Mulcaire, Miguel Ballesteros, Chris Dyer, and Noah Smith. 2016. Many languages, one parser. *Trans. Assoc. Comput. Ling.* 4 (2016), 431–444.
- 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 Annual Meeting of the Association for Computational Linguistics (ACL'16)*. Association for Computational Linguistics, 2442–2452. DOI : <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* (2014). <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 Conference on Empirical Methods in Natural Language Processing (EMNLP'15)*. Association for Computational Linguistics, 349–359. DOI : <https://doi.org/10.18653/v1/D15-1041>
- Ann Bies, Justin Mott, Colin Warner, and Seth Kulick. 2012. English Web Treebank LDC2012T13. (2012).
- Danqi Chen and Christopher Manning. 2014. A fast and accurate dependency parser using neural networks. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP'14)*. Association for Computational Linguistics, 740–750. DOI : <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 Conference on Empirical Methods in Natural Language Processing (EMNLP'16)*. Association for Computational Linguistics, 731–741.
- Shay Cohen and A. Noah Smith. 2009. Shared logistic normal distributions for soft parameter tying in unsupervised grammar induction. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT'09)*. Association for Computational Linguistics, 74–82.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *J. Mach. Learn. Res.* 12 (2011), 2493–2537.
- Timothy Dozat and Christopher D. Manning. 2017. Deep biaffine attention for neural dependency parsing. In *International Conference on Learning Representations 2017*. *arXiv preprint 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. *J. Mach. Learn. Res.* 12 (2011), 2121–2159.
- Long Duong, Trevor Cohn, Steven Bird, and Paul Cook. 2015. A neural network model for low-resource universal dependency parsing. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP'15)*. Association for Computational Linguistics, 339–348. DOI : <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 Annual Meeting of the Association for Computational Linguistics and the International Joint Conference on Natural Language Processing (ACL-IJCNLP'15)*. Association for Computational Linguistics, 334–343. DOI : <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 Annual Meeting of the Association for Computational Linguistics and the International Joint Conference on Natural Language Processing (ACL-IJCNLP'09)*. Association for Computational Linguistics, 369–377.
- Douwe Gelling, Trevor Cohn, Phil Blunsom, and Joao Graca. 2012. The PASCAL challenge on grammar induction. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies Workshop on the Induction of Linguistic Structure (NAACL-HLT'12)*. Association for Computational Linguistics, 64–80.
- Jennifer Gillenwater, Kuzman Ganchev, João Graça, Fernando Pereira, and Ben Taskar. 2010. Sparsity in dependency grammar induction. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL'10) (Short Papers)*. Association for Computational Linguistics, 194–199. <http://www.aclweb.org/anthology/P10-2036>.
- Alex Graves and Jürgen Schmidhuber. 2005. Framewise phoneme classification with bidirectional LSTM and other neural network architectures. *Neur. Netw.* 18, 5 (2005), 602–610.
- Klaus Greff, Rupesh Kumar Srivastava, Jan Koutník, Bas R. Steunebrink, and Jürgen Schmidhuber. 2015. LSTM: A search space odyssey. *CoRR* abs/1503.04069 (2015). <http://arxiv.org/abs/1503.04069>.
- Jiang Guo, Wanxiang Che, David Yarowsky, Haifeng Wang, and Ting Liu. 2015. Cross-lingual dependency parsing based on distributed representations. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics and the International Joint Conference on Natural Language Processing (ACL-IJCNLP'15)*. Association for Computational Linguistics, 1234–1244. DOI : <https://doi.org/10.3115/v1/P15-1119>
- Shinichi Harada. 2009. The roles of singapore standard english and singlish. *Inf. Res.* 40 (2009), 70–82.
- Kenneth Heafield, Ivan Pouzyrevsky, H. Jonathan Clark, and Philipp Koehn. 2013. Scalable modified Kneser-Ney language model estimation. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL'13)*. Association for Computational Linguistics, 690–696. <http://aclweb.org/anthology/P13-2121>.
- Sanjika Hewavitharana, Nguyen Bach, Qin Gao, Vamshi Ambati, and Stephan Vogel. 2011. CMU Haitian Creole-English translation system for WMT. In *Proceedings of the 6th Workshop on Statistical Machine Translation*. Association for Computational Linguistics, 386–392.
- Chang Hu, Philip Resnik, Yakov Kronrod, Vladimir Eidelman, Olivia Buzek, and B. Benjamin Bederson. 2011. The value of monolingual crowdsourcing in a real-world translation scenario: Simulation using Haitian Creole emergency SMS messages. In *Proceedings of the 6th Workshop on Statistical Machine Translation*. Association for Computational Linguistics, 399–404.
- Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional LSTM-CRF models for sequence tagging. *arXiv preprint* abs/1508.01991 (2015). <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. *Nat. Lang. Eng.* 11, 3 (Sep. 2005), 311–325. DOI : <https://doi.org/10.1017/S1351324905003840>
- Eliyahu Kiperwasser and Yoav Goldberg. 2016. Simple and accurate dependency parsing using bidirectional LSTM feature representations. *Trans. Assoc. Comput. Ling.* 4 (2016), 313–327.
- Karen Lahousse and Béatrice Lamiroy. 2012. Word order in French, Spanish and Italian: A grammaticalization account. *Folia Ling.* 46, 2 (2012), 387–415.
- Jakob R. E. Leimgruber. 2009. *Modelling Variation in Singapore English*. Ph.D. Dissertation. Oxford University.
- Jakob R. E. Leimgruber. 2011. Singapore English. *Lang. Ling. Compass* 5, 1 (2011), 47–62. DOI : <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 Conference on Empirical Methods in Natural Language Processing (EMNLP'15)*. Association for Computational Linguistics, 2304–2314. DOI : <https://doi.org/10.18653/v1/D15-1278>
- Lisa Lim. 2007. Mergers and acquisitions: On the ages and origins of Singapore English particles. *World Engl.* 26, 4 (2007), 446–473.
- Yijia Liu, Yi Zhu, Wanxiang Che, Bing Qin, Nathan Schneider, and Noah A. Smith. 2018. Parsing tweets into universal dependencies. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT'18)*, 965–975. <https://aclanthology.info/papers/N18-1088/n18-1088>
- Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014. The Stanford CoreNLP natural language processing toolkit. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL'14)*. 55–60.
- Héctor Martínez Alonso, Željko Agić, Barbara Plank, and Anders Søgaard. 2017. Parsing universal dependencies without training. In *Proceedings of the European Chapter of the Association for Computational Linguistics (EACL'17)*. Association for Computational Linguistics, 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 Conference on Empirical Methods in Natural Language Processing (EMNLP'11)*. Association for Computational Linguistics, 62–72.

- Ho Mian-Lian and John T. Platt. 1993. *Dynamics of a Contact Continuum: Singaporean English*. Oxford University Press, New York.
- Makoto Miwa and Mohit Bansal. 2016. End-to-end relation extraction using LSTMs on sequences and tree structures. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL'16)*. Association for Computational Linguistics, 1105–1116. DOI : <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 Annual Meeting of the Association for Computational Linguistics (ACL'12)*. Association for Computational Linguistics, 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 Conference on Empirical Methods in Natural Language Processing (EMNLP'10)*. Association for Computational Linguistics, Cambridge, MA, 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 (1992), 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 International Conference on Language Resources and Evaluation (LREC'16)*. European Language Resources Association, 23–28.
- 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 Conference on Empirical Methods in Natural Language Processing (EMNLP-CoNLL'07)*. Association for Computational Linguistics, 915–932.
- Brendan T. O'Connor, Su Lin Blodgett, and Johnny Wei. 2018. Twitter universal dependency parsing for African-American and mainstream American english. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL'18)*. 1415–1425.
- Vincent B. Y. Ooi. 1997. *Analysing the Singapore ICE Corpus for Lexicographic Evidence*.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP'14)*. Association for Computational Linguistics, 1532–1543. DOI : <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 Annual Meeting of the Association for Computational Linguistics (ACL'16)*. Association for Computational Linguistics, 412–418. DOI : <https://doi.org/10.18653/v1/P16-2067>
- Manuela Sanguinetti, Cristina Bosco, Alberto Lavelli, Alessandro Mazzei, Oronzo Antonelli, and Fabio Tamburini. 2018. PoSTWITA-UD: An Italian Twitter treebank in universal dependencies. In *Proceedings of the 11th International Conference on Language Resources and Evaluation (LREC'18)*.
- Manuela Sanguinetti, Cristina Bosco, Alessandro Mazzei, Alberto Lavelli, and Fabio Tamburini. 2017. Annotating Italian social media texts in universal dependencies. In *Proceedings of the 4th International Conference on Dependency Linguistics*. 229–239. <https://aclanthology.info/papers/W17-6526/w17-6526>
- 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 *Proceedings of the 18th International Conference on Information Fusion (Fusion'15)*. IEEE, 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 Conference on Empirical Methods in Natural Language Processing (EMNLP'13)*. Association for Computational Linguistics, 1631–1642.
- Anders Søgaard. 2012a. Two baselines for unsupervised dependency parsing. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies Workshop on the Induction of Linguistic Structure (NAACL-HLT'12)*. Association for Computational Linguistics, 81–83.
- Anders Søgaard. 2012b. Unsupervised dependency parsing without training. *Nat. Lang. Eng.* 18, 2 (2012), 187–203. DOI : <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 Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT 2012)*. Association for Computational Linguistics, 477–487.
- Hongmin Wang, Yue Zhang, GuangYong Leonard Chan, Jie Yang, and Hai Leong Chieu. 2017. Universal dependencies parsing for colloquial Singaporean english. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, 1732–1744. <http://aclweb.org/anthology/P17-1159>.
- David Weiss, Chris Alberti, Michael Collins, and Slav Petrov. 2015. Structured training for neural network transition-based parsing. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics and the International*

- Joint Conference on Natural Language Processing (ACL-IJCNLP'15)*. Association for Computational Linguistics, 323–333. DOI : <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 Conference on Empirical Methods in Natural Language Processing (EMNLP'15)*. Association for Computational Linguistics, 1857–1867. DOI : <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 Annual Meeting of the Association for Computational Linguistics (ACL'16)*. Association for Computational Linguistics, 1557–1566. DOI : <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 Annual Meeting of the Association for Computational Linguistics and the International Joint Conference on Natural Language Processing (ACL-IJCNLP'15)*. Association for Computational Linguistics, 1213–1222. DOI : <https://doi.org/10.3115/v1/P15-1117>

Received July 2018; revised December 2018; accepted March 2019