Second language Korean Universal Dependency treebank v1.2: Focus on data augmentation and annotation scheme refinement

Hakyung Sung
Department of Linguistics
University of Oregon
hsung@uoregon.edu

Gyu-Ho Shin
Department of Linguistics
University of Illinois Chicago
ghshin@uic.edu

Abstract

We expand the second language (L2) Korean Universal Dependencies (UD) treebank with 5,454 manually annotated sentences. The annotation guidelines are also revised to better align with the UD framework. Using this enhanced treebank, we fine-tune three Korean language models—Stanza, spaCy, and Trankit—and evaluate their performance on in-domain and out-of-domain L2-Korean datasets. The results show that fine-tuning significantly improves their performance across various metrics, thus highlighting the importance of using well-tailored L2 datasets for fine-tuning first-language-based, general-purpose language models for the morphosyntactic analysis of L2 data.

1 Introduction

The Universal Dependencies (UD) framework, designed to facilitate accessible morphosyntactic annotations (de Marneffe et al., 2021), has been applied increasingly in linguistics, particularly to annotate learner corpora. This approach supports tasks such as modeling the trajectories of second language (L2) acquisition, which often require treebanks for fine-tuning language models or evaluating their performance on L2 data. Such data are typically characterized by simpler and/or nontarget-like lexico-grammatical usages compared to those produced by first-language speakers, although these characteristics vary across L2 proficiency. Previous research has increasingly adopted the UD framework to automatically handle learner corpora in various languages, including English (Berzak et al., 2016; Kyle et al., 2022; Lyashevskaya and Panteleeva, 2017; Huang et al., 2018), Chinese (Lee et al., 2017), Italian (Di Nuovo et al., 2019, 2022), Russian (Rozovskaya, 2024), and Swedish (Masciolini et al., 2024; Masciolini, 2023; Masciolini et al., 2023), demonstrating its utility in L2 studies.

Among these efforts, recent studies in Korean have developed L2-Korean UD treebanks with language-specific morphemes and dependency tags (Sung and Shin, 2023a,b, 2024). However, two research gaps remain. First, while continuing to expand the amount of data, annotation guidelines should be iteratively updated to balance cross-linguistic standardization with the preservation of language-specific features (de Marneffe et al., 2021; Manning, 2011). Second, the effectiveness of L2-Korean-optimized models should be assessed using out-of-domain data to improve their reliability in broader contexts for which they are designed (Plank, 2016; Joshi et al., 2018).

The present study addresses these gaps with three key contributions: (1) augmenting the existing L2-Korean UD treebank (v1.1, 7,530 sentences) by adding 5,454 manually annotated sentences with Korean-specific morphemes and UD annotations; (2) revising dependency annotation guidelines extensively to better align with the language-general UD framework, while implementing minor adjustments to the guidelines to better reflect the linguistic properties of Korean; and (3) fine-tuning and evaluating three Korean language models in both in-domain and out-of-domain contexts using the updated L2-Korean UD treebank (v1.2, 12,984 sentences, see Appendix for XPOS and DEPREL tag distributions).

2 Related works

A line of studies have established approaches for morpheme and dependency annotations in L2 Korean. Sung and Shin (2023b) provided preliminary guidelines for Korean morpheme annotations, addressing the need to parse morphemes taking into account the agglutinative nature of Korean morphosyntax, where a single word often combines lexical morphemes (e.g., noun, verb) and func-

Proceedings of the Third Workshop on Resources and Representations for Under-Resourced Languages and Domains (RESOURCEFUL 2025), pages 13–19
March 2, 2025

tional morphemes (e.g., postpositions, tense-aspect-modality markers). Expanding this work, Sung and Shin (2024) introduced detailed UD annotation guidelines to handle Korean-specific dependency cases such as particles and coordination.

Sung and Shin (2023a) fine-tuned morpheme parsers optimized for L2 Korean and evaluated them on in-domain and out-of-domain datasets, demonstrating the importance of high-quality input for fine-tuning L2-Korean language models. However, those studies did not include training or evaluating dependency tags. Additionally, their fine-tuning strategy was relatively simple, relying solely on one Korean pre-trained model.

3 Dataset

Building upon the previous L2-Korean UD annotation projects (Sung and Shin, 2023b, 2024), we continued annotating L2-Korean sentences using a subset of data from the same source (Park and Lee, 2016).[1] For the out-of-domain testing, we annotated additional data from the KoLLA dataset (Lee, 2022), which was designed to analyze Korean learner language with a focus on particle error annotations.[2]

Along with the annotations, we refined the annotation guidelines, implementing major revisions to better align with the language-general UD annotation scheme and minor adjustments to morpheme annotations. Together, the updated L2-Korean UD treebank (v1.2) comprises (# sents = 12,984): (1) additional data augmented and annotated using the revised scheme (# sents = 4,532); (2) revised data from the previous project (Sung and Shin, 2024), updated with the new annotation scheme (# sents = 7,530); (3) data sourced from the KoLLA dataset (Lee, 2022), annotated with the revised scheme for the out-of-domain testing (# sents = 922).

3.1 Refining annotation guidelines

Carefully curated linguistic annotations balance two key challenges: maintaining consistency and ensuring accuracy. Manning (2011) highlighted the challenges involving POS labeling, noting the inherent ambiguities and unclear boundaries between word classes, which complicate the definitive assignment of labels. Such intrinsic ambiguities can degrade the performance of taggers when training language models. Therefore, systematic checks and guideline refinements are essential for achieving optimal annotations.

For L2-Korean annotations, Sung and Shin (2024) emphasized dependency annotations grounded in language-specific justifications, building upon earlier studies of Korean dependency annotations (Lee et al., 2019; Kim et al., 2018; Seo et al., 2019). However, the previous annotation scheme did not fully conform to the language-general UD framework and exhibited notable mismatches between tags, particularly `conj`, `flat`, and `aux`. To address these issues, we revised the previous dependency annotation guidelines to better align with the language-general UD conventions, thus enhancing global applicability. Below, we outline two key areas of major changes implemented.

3.1.1 Following the left-to-right rule

The UD framework enforces a strict left-to-right rule for coordination to ensure consistency and cross-linguistic applicability in morphosyntactic annotations (Nivre et al., 2016; de Marneffe et al., 2021). This approach originates from the Stanford-typed dependencies for English (de Marneffe et al., 2006), which serve as the foundation for the universal dependency representation (McDonald et al., 2013).

Coordination Coordination (`conj`) is handled by consistently attaching the coordinating conjunction to the head of the first conjunct. The leftmost conjunct is designated as the head, with subsequent conjuncts and the coordinating conjunction depending on it.[3]

Initially, Sung and Shin (2024) assigned the head to the right-headed structure in complex clauses or noun phrase conjunctions. For instance, in complex clauses, the head was assigned to the predicate, often resulting in a right-headed structure. This approach was driven by the nature of the Korean connective marker (e.g., -고 [*-ko*]), which signifies conjunction and is logically tagged as `conj` (p. 3748). However, in line with the current UD guidelines, we revised the previous approach to strictly follow the left-to-right head structure, consistent with the UD's left-headed coordination. Now, the connective marker -고 (*-ko*) is tagged as `root`, and

[1]The source data became unavailable as of September 2024.

[2]The dataset is publicly available at: `https://cl.indiana.edu/~kolla/`

[3]This approach, while widely adopted, has raised some questions, as noted by Gerdes and Kahane (2016), where the selection of the first conjunct as the head is made without extensive justifications (p. 7).

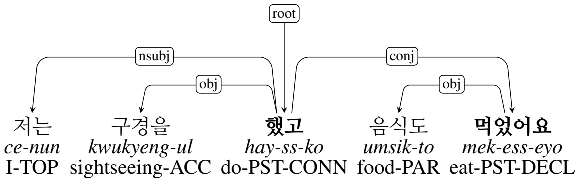

Figure 1: Coordination (Left-headed)
'I looked around and ate some foods.'

the final predicate receives the `conj` tag (Figure 1).[4]

Flat Flat (`flat`) is used when no single element in an expression can be clearly identified as the head. Similar to the case of coordination, in this structure, the leftmost element is treated as the head, with all subsequent components attached to it as equals. This applies to expressions such as "John Smith" or "San Francisco," where no one part dominates the meaning of the whole.

In the previous L2-Korean UD annotation scheme, the core principle for assigning the head was based on the presence of particles, reflecting how they function in determining the grammatical roles of nouns in Korean—core arguments (subject, object) or non-core arguments (obliques) within a clause (Sohn, 1999). However, to conform to the UD framework's left-to-right rule, we rigorously revised all flat relations to follow this directionality. This revision affected the majority of naming conventions and combinations of names with titles in our annotated data, as described in Figure 2.

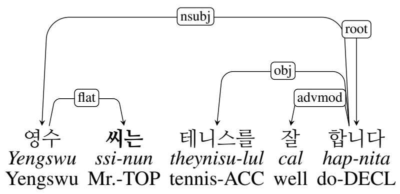

Figure 2: Flat (Left-headed)
'Youngsoo is good at tennis.'

3.1.2 Treatment of auxiliary verbs

The revised annotation scheme strictly adheres to the UD guidelines for Korean, limiting the annotation of auxiliary verbs to five specific forms.[5] These forms include (1) the affirmative copula 이- (*i*-, "to be"), which is treated as a separate auxiliary even when it functions as a suffix to a nominal predicate;[6] (2) the negative copula 않- (*anh*-, "to not be"), annotated as `AUX` in negative clauses; (3) the affirmative auxiliary 있- (*iss*-, "to be"), used as an auxiliary in affirmative clauses or to indicate progressive aspect; (4) the necessitative modal 하- (*ha*-, "must, should"), which functions as a modal auxiliary expressing necessity; and (5) the desiderative modal 싶- (*sip*-, "will, want"), which serves as a modal auxiliary expressing a desire or intention. Verbs with auxiliary-like meanings outside this set were tagged as adverbial clause modifiers (`advcl`).

3.2 Annotation process

The annotation was conducted by five native Korean speakers, each holding at least an undergraduate degree in Korean linguistics. To manage the workload and ensure comprehensive coverage, the annotators were divided into two groups, with each sentence independently annotated by a pair from one group. The annotators worked independently to minimize bias and preserve the integrity of their individual assessments, without interim adjudication meetings to resolve disagreements. When discrepancies arose between the initial pair of annotators, a third annotator, and if necessary, a fourth, were involved sequentially. Inter-annotator reliability was assessed for the initial annotation pairs (before the adjudication process) using the augmented dataset (# sents = 4,532, Table 1).

Annotation	Cohen's *Kappa*
LEMMA	0.964
XPOS	0.908
HEAD	0.892
DEPREL	0.927

Table 1: Inter-annotator reliability

4 Experiment

4.1 Model training

We evaluated four language models against L2-Korean morphosyntactic annotation tasks, drawing upon user-friendly NLP toolkits designed for multilingual applications in fundamental NLP tasks:

[4]We also revised noun phrase conjunctions, as in examples such as 사과와 바나나 (*sakwa-wa panana*, "apple and banana"), where 사과 (*sakwa*, "apple") is the head and 바나나 (*panana*, "banana") depends on it, with the coordinating conjunction -와 (-*wa*, "and") linking the two.

[5]`https://universaldependencies.org/ko/index.html`

[6]When 이- follows a noun and precedes a sentence-final functional morpheme (e.g., -다 -*ta*, as in 친구이다 *chinkwu-i-ta, "is a friend"*), we assigned it the `root` tag, simplifying the earlier practice of using a special `root:cop` tag.

(1) **Baseline:** Stanza-Korean (GSD package) (Qi et al., 2020) was used as a benchmark without fine-tuning. It aligns with both the Sejong tag set and the UD framework; (2) ***Stanza*:** We fine-tuned Stanza-Korean (GSD), which employs a biLSTM architecture (Huang et al., 2015) to model sequential dependencies. Fine-tuning allows the model to better capture localized morphosyntactic patterns in L2-Korean data by leveraging the tagging scheme and linguistic patterns encoded in the pre-existing GSD package; (3) ***spaCy*:** We fine-tuned spaCy (Honnibal et al., 2020), which uses its `tok2vec` layer to generate token-level embeddings from sub-word features. Fine-tuning in spaCy benefits from pre-trained word vectors and built-in lexical resources, making it well-suited for modeling specific lexico-grammatical nuances; (4) ***Trankit*:** We fine-tuned Trankit (Van Nguyen et al., 2021), which uses a transformer-based architecture (XLM-RoBERTa, Conneau et al., 2020) pre-trained on 100 languages. Fine-tuning a custom pipeline in *Trankit* using the *TPipeline* class enables the model to capture long-range dependencies and complex syntactic structures. All models were trained using their default hyperparameter settings to ensure a fair comparison.

4.2 Dataset split

The updated L2-Korean UD treebank (v1.2) was divided into subsets for training and evaluation. The training set contained 9,649 sentences, while the development set, comprising 1,208 sentences, was used for fine-tuning and model optimization. The test set, which included 1,205 sentences, was used to evaluate in-domain performance. Additionally, an out-of-domain test set comprising 922 sentences was designated to assess the models' robustness and generalizability to data beyond the training space.

4.3 Evaluation Metrics

To evaluate these models, we measured F1 scores across the following metrics: XPOS, LEMMA, UAS (Unlabeled Attachment Score), and LAS (Labeled Attachment Score).

4.4 Results

The fine-tuned models effectively improved their performance across various metrics for both in-domain and out-of-domain datasets. For the in-domain L2K-UD-test set, Trankit outperformed other models in XPOS, UAS, and LAS, while Stanza achieved the best LEMMA score despite trailing overall. In the out-of-domain KoLLA treebank, Trankit again excelled in XPOS, UAS, and LAS, demonstrating its generalizability beyond the traning space. Stanza consistently performed best in the LEMMA metric, indicating its strong lexical capabilities even with domain shifts.

Dataset	Metric	Baseline	Stanza	spaCy	Trankit
L2K-UD-test (in-domain)	XPOS	82.44	89.72	83.15	**91.81**
	LEMMA	89.61	**95.64**	87.97	88.84
	UAS	76.72	85.53	82.21	**92.28**
	LAS	60.69	80.36	75.21	**89.13**
KoLLA (out-of-domain)	XPOS	77.79	81.87	71.21	**84.51**
	LEMMA	88.03	**91.01**	79.64	86.90
	UAS	72.30	81.17	74.48	**88.93**
	LAS	58.53	75.14	63.56	**85.45**

Table 2: Evaluation metrics

5 Discussion and future directions

We expanded the L2-Korean UD treebank with refined annotation schemes to improve model performance after fine-tuning. Using this treebank, we fine-tuned three models—Stanza, spaCy, and Trankit—and evaluated their performance in both in-domain and out-of-domain contexts. The evaluation results showed significant performance improvements across various metrics, underscoring the value of using an L2 dataset for fine-tuning. Among the models, Trankit's transformer-based architecture outperformed the others in XPOS, UAS, and LAS across both test datasets, demonstrating its effectiveness of capturing morphosyntactic features in L2-Korean data. The fine-tuned models and relevant documentations are available at `https://github.com/NLPxL2Korean/UD-KSL`. The treebank will be updated at `https://github.com/UniversalDependencies/UD_Korean-KSL`.

Although both Trankit and Stanza employ a character-based seq2seq model (Van Nguyen et al., 2021), Stanza's superior lemmatization performance compared to Trankit can be attributed to two primary factors. First, Stanza includes a dictionary-based lemmatizer (Qi et al., 2020), which may have strengthened its ability to handle a wide variety of morphological patterns. Second, as noted earlier, Stanza uniquely leverages a model that was pre-trained on L1 data (UD-Korean GSD) before being fine-tuned on the current L2 data, which appears to enable it to capitalize on prior lemmatization knowledge for more accurate predictions.

To fully harness the potential of transformer-based architectures in fine-tuning L2-Korean models, future L2-Korean UD treebanks could adopt two complementary strategies. One approach involves combining L2-Korean data drawn from various genres or diverse learner backgrounds. The other centers on refining the match between universal UPOS tags and language-specific XPOS tags through expert revisions to enhance UPOS to boost their effectiveness for lemmatization within the seq2seq framework.

Acknowledgments

The authors gratefully acknowledge Hee-June Koh, Chanyoung Lee, Youkyung Sung for their contributions to manual annotations and discussions for the enhancement of annotation guidelines. This research was supported by the 2024 Korean Studies Grant Program of the Academy of Korean Studies (AKS-2024-R-012).

References

Yevgeni Berzak, Jessica Kenney, Carolyn Spadine, Jing Xian Wang, Lucia Lam, Keiko Sophie Mori, Sebastian Garza, and Boris Katz. 2016. https://aclanthology.org/P16-1070.pdf Universal dependencies for learner english. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 737–746.

Alexis Conneau, Alexei Baevski, Ronan Collobert, Abdelrahman Mohamed, and Michael Auli. 2020. Unsupervised cross-lingual representation learning for speech recognition. *arXiv preprint arXiv:2006.13979*.

Elisa Di Nuovo, Cristina Bosco, Alessandro Mazzei, Manuela Sanguinetti, et al. 2019. Towards an italian learner treebank in universal dependencies. In *CEUR workshop proceedings*, volume 2481, pages 1–6. CEUR-WS.

Elisa Di Nuovo, Manuela Sanguinetti, Alessandro Mazzei, Elisa Corino, and Cristina Bosco. 2022. Valico-ud: Treebanking an italian learner corpus in universal dependencies. *IJCoL. Italian Journal of Computational Linguistics*, 8(8-1).

Kim Gerdes and Sylvain Kahane. 2016. Dependency annotation choices: Assessing theoretical and practical issues of universal dependencies. In *LAW X (2016) The 10th Linguistic Annotation Workshop: 131*.

Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. spacy: Industrial-strength natural language processing in python.

Yan Huang, Akira Murakami, Theodora Alexopoulou, and Anna Korhonen. 2018. https://www.jbe-platform.com/content/journals/10.1075/ijcl.16080.hua Dependency parsing of learner english. *International Journal of Corpus Linguistics*, 23(1):28–54.

Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional lstm-crf models for sequence tagging. *arXiv preprint arXiv:1508.01991*.

Vidur Joshi, Matthew Peters, and Mark Hopkins. 2018. https://doi.org/10.18653/v1/P18-1110 Extending a parser to distant domains using a few dozen partially annotated examples. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1190–1199, Melbourne, Australia. Association for Computational Linguistics.

Hansaem Kim et al. 2018. *2018년 국어 말뭉치 연구 및 구축 [2018 Korean language corpus research and construction]*. National Institute of Korean Language, Republic of Korea.

Kristopher Kyle, Masaki Eguchi, Aaron Miller, and Theodore Sither. 2022. https://aclanthology.org/2022.bea-1.7/ A dependency treebank of spoken second language english. In *Proceedings of the 17th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2022)*, pages 39–45.

Chanyoung Lee, Tae hwan Oh, and Hansam Kim. 2019. 한국어 보편 의존 구문 분석 (universal dependencies) 방법론 연구 [a study on universal dependency annotation for korean]. *언어사실과 관점 [Language Facts and Perspectives]*, 47:141–175.

John Lee, Herman Leung, and Keying Li. 2017. https://aclanthology.org/W17-0408/ Towards universal dependencies for learner chinese. In *Proceedings of the NoDaLiDa 2017 Workshop on Universal Dependencies (UDW 2017)*, pages 67–71.

Sun-Hee Lee. 2022. Corpus-based research and ksl. In *The Routledge Handbook of Korean as a Second Language*, pages 257–276. Routledge.

Olga Lyashevskaya and Irina Panteleeva. 2017. Realec learner treebank: annotation principles and evaluation of automatic parsing. In *Proceedings of the 16th International Workshop on Treebanks and Linguistic Theories*, pages 80–87.

Christopher D Manning. 2011. Part-of-speech tagging from 97% to 100%: is it time for some linguistics? In *International conference on intelligent text processing and computational linguistics*, pages 171–189. Springer.

Marie-Catherine de Marneffe, Bill MacCartney, Christopher D Manning, et al. 2006. Generating typed dependency parses from phrase structure parses. In *Lrec*, volume 6, pages 449–454.

Marie-Catherine de Marneffe, Christopher D Manning, Joakim Nivre, and Daniel Zeman. 2021. Universal dependencies. *Computational linguistics*, 47(2):255–308.

Arianna Masciolini. 2023. A query engine for l1-l2 parallel dependency treebanks. In *Proceedings of the 24th Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 574–587.

Arianna Masciolini, Emilie Francis, and Maria Irena Szawerna. 2024. Synthetic-error augmented parsing of swedish as a second language: Experiments with word order. In *Proceedings of the Joint Workshop on Multiword Expressions and Universal Dependencies (MWE-UD)@ LREC-COLING 2024*, pages 43–49.

Arianna Masciolini, Elena Volodina, and Dana Dannlls. 2023. Towards automatically extracting morphosyntactical error patterns from l1-l2 parallel dependency treebanks. In *Proceedings of the 18th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2023)*, pages 585–597.

Ryan McDonald, Joakim Nivre, Yvonne Quirmbach-Brundage, Yoav Goldberg, Dipanjan Das, Kuzman Ganchev, Keith Hall, Slav Petrov, Hao Zhang, Oscar Täckström, et al. 2013. https://aclanthology.org/P13-2017/ Universal dependency annotation for multilingual parsing. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 92–97.

Joakim Nivre, Marie-Catherine De Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajic, Christopher D Manning, Ryan McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, et al. 2016. Universal dependencies v1: A multilingual treebank collection. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 1659–1666.

Jungyeul Park and Jung Hee Lee. 2016. A korean learner corpus and its features. *En-e-hak [Linguistics]*, (75):69–85.

Barbara Plank. 2016. What to do about non-standard (or non-canonical) language in nlp. *arXiv preprint arXiv:1608.07836*.

Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D Manning. 2020. https://aclanthology.org/2020.acl-demos.14/ Stanza: A python natural language processing toolkit for many human languages. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 101–108.

Alla Rozovskaya. 2024. Universal dependencies for learner russian. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 17112–17119.

Saetbyol Seo et al. 2019. 한국어 보편 의존 관계 분석에 관한 제언 [a proposal on universal dependencies (v.2) annotation for korean]. *언어와 정보 [Language and Information]*, 23(1):91–122.

Ho-Min Sohn. 1999. *The Korean language*. New York, NY: Cambridge University Cambridge University Press.

Hakyung Sung and Gyu-Ho Shin. 2023a. Diversifying language models for lesser-studied languages and language-usage contexts: A case of second language korean. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 11461–11473.

Hakyung Sung and Gyu-Ho Shin. 2023b. Towards l2-friendly pipelines for learner corpora: A case of written production by l2-korean learners. In *Proceedings of the 18th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2023)*, pages 72–82.

Hakyung Sung and Gyu-Ho Shin. 2024. Constructing a dependency treebank for second language learners of korean. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 3747–3758.

Minh Van Nguyen, Viet Dac Lai, Amir Pouran Ben Veyseh, and Thien Huu Nguyen. 2021. Trankit: A light-weight transformer-based toolkit for multilingual natural language processing. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 80–90.

Appendix

XPOS	v1.1	v1.2	DEPREL	v1.1	v1.2
NNG	25338	40001	nsubj	8767	13781
VV	10219	16714	punct	8287	14066
EC	8600	13282	obl	7332	12034
EF	7541	12994	root	6866	12989
SF	7525	12948	obj	5572	9203
ETM	6694	9831	advmod	4995	7829
JKB	6366	10450	advcl	4703	8425
JX	5406	8656	acl	4501	6400
NNB	4748	7454	nmod	2059	3882
JKO	4735	7717	aux	1963	2312
MAG	4312	6774	conj	1860	2782
JKS	4136	6668	amod	1413	2176
VA	3380	5905	cc	1306	2154
XSV	3278	4761	nmod:poss	1299	1877
VX	3237	4555	det	933	1373
EP	2850	5215	case	894	1477
NNP	2847	4810	flat	854	1172
NP	2145	3548	ccomp	642	897
VCP	2083	3098	dislocated	576	1035
MM	1672	2689	mark	509	838
XSN	1467	2179	list	303	444
JKG	1329	1921	goeswith	203	235
NF	1312	2208	nummod	179	342
XSA	1199	1815	appos	128	95
MAJ	1160	1921	compound	52	112
SN	1017	1475	vocative	46	49
ETN	830	1213	parataxis	37	39
NA	818	1215	csubj	22	22
JC	685	1269	discourse	6	6
SP	607	864	fixed	6	24
XR	424	684	dep	3	5
SS	266	378			
NV	262	516			
VCN	174	251			
XPN	167	208			
NR	157	228			
SL	133	268			
JKC	122	177			
JKQ	58	86			

Table 3: Comparison of XPOS and DEPREL tag distributions in L2-Korean UD v.1.1 and v.1.2