

Fine-Grained Analysis of Cross-Linguistic Syntactic Divergences

Dmitry Nikolaev^{‡*} Ofir Arviv[†] Taelin Karidi[†] Neta Kenneth[†]

Veronika Mitnik[†] Lilja Maria Saeboe[†] Omri Abend[†]

[‡]Stockholm University [†]Hebrew University of Jerusalem

dnikolaev@fastmail.com

{ofir.arviv|taelin.karidi|neta.kenneth|

veronika.mitnik|liljama.saeboe|omri.abend}@mail.huji.ac.il

Abstract

The patterns in which the syntax of different languages converges and diverges are often used to inform work on cross-lingual transfer. Nevertheless, little empirical work has been done on quantifying the prevalence of different syntactic divergences across language pairs. We propose a framework for extracting divergence patterns for any language pair from a parallel corpus, building on Universal Dependencies (UD; Nivre et al., 2016). We show that our framework provides a detailed picture of cross-language divergences, generalizes previous approaches, and lends itself to full automation. We further present a novel dataset, a manually word-aligned subset of the Parallel UD corpus in five languages, and use it to perform a detailed corpus study. We demonstrate the usefulness of the resulting analysis by showing that it can help account for performance patterns of a cross-lingual parser.

1 Introduction

The assumption that the syntactic structure of a sentence is predictably related to the syntactic structure of its translation has deep roots in NLP, notably in cross-lingual transfer methods, such as annotation projection and multi-lingual parsing (Hwa et al., 2005; McDonald et al., 2011; Kozhevnikov and Titov, 2013; Rasooli and Collins, 2017, *inter alia*), as well as in syntax-aware machine translation (MT; Birch et al., 2008; Williams et al., 2016; Bastings et al., 2017). Relatedly, typological parameters that provide information on the dimensions of similarity between grammars of different languages were found useful for a variety of NLP applications (Ponti et al., 2019). For example, neural MT in low-resource settings has been shown to benefit from bridg-

* Work mostly done while at the Hebrew University of Jerusalem.




Figure 1: An example En-Ko sentence pair exhibiting a divergence, where an En *nmod* path corresponds to a Ko *acl:recl+nsbj* path (paths are given in bold red). The En preposition *by* is not considered a content word and is not aligned with the Ko verb *gigohan*. See § 3 for further details.

ing morphosyntactic differences in parallel training data by different types of preprocessing, such as reordering (Zhou et al., 2019) and hand-coded syntactic manipulations (Ponti et al., 2018).

Nevertheless, little empirical work has been done on systematically quantifying the type and prevalence of syntactic divergences across languages. Moreover, previous work generally classified divergences into a small set of divergence classes, often based on theoretical considerations (Dorr, 1994) or on categorical (“hard”) typological features selected in an ad-hoc manner, and left basic questions, such as how often POS tags are preserved in translation and what syntactic structures are likely correspondents of different syntactic relations, largely unaddressed. See § 2.

We propose a language-neutral, fine-grained definition of *cross-linguistic morphosyntactic divergences* (CLMD) that allows for their extraction using a syntactically annotated, content-word-aligned parallel corpus. Concretely, we classify CLMD based on the edge labels on the dependency paths between corresponding pairs of content words (§ 3.2). See Figure 1 for an example.¹

¹It may appear that divergences recoverable by means of UD edge labels are purely *syntactic* and not morphosyntactic. However, this is not the case: the domain of “pure syntax” is not well defined in a non-theoretical perspective, and many phenomena we are dealing with, e.g. a switch from a direct-

We further conduct a detailed corpus study, manually aligning content words in a subset of the PUD corpus (Zeman et al., 2017) over five language pairs—English-French (En-Fr), English-Russian (En-Ru), English-Chinese (En-Zh), English-Korean (En-Ko), and English-Japanese (En-Jp) (§ 3.1)—and analyze the prevalence of divergences by types (§ 4). The resulting resource can be useful for MT research, by guiding the creation of challenge sets focusing on particular constructions known to be cross-linguistically divergent, as well as by guiding preprocessing of source-side sentences for better MT performance.²

The emerging CLMD provide information not only on the macro-structure of the grammar (e.g., whether the language is pro-dropping), but also on parameters specific to certain lexical classes (e.g., modal verbs) and probabilistic tendencies (e.g., Japanese tends to translate sequences of events, expressed in English using coordinating conjunctions, with subordinate clauses). See § 5.

Further experiments demonstrate the methodology’s applicative potential. First, we show that the proposed methodology can be straightforwardly automated by replacing manual parses and alignments with automatically induced ones (§ 7). We present a study done on a larger En-Zh corpus, which yields results similar to those obtained manually. Secondly, we show that the reported distribution over divergence types is predictive of the performance patterns of a zero-shot parser (§ 8).

2 Related Work

Comparing syntactic and semantic structures over parallel corpora is the subject of much previous work. Dorr et al. (2010) compiled a multiply-parallel corpus and annotated it with increasingly refined categories in an attempt to abstract away from syntactic detail but did not report any systematic measurement of the distribution of divergences. Šindlerová et al. (2013), Xue et al. (2014), Sulem et al. (2015), and Damonte and Cohen (2018) studied divergences over semantic graphs and argument-structure phenomena, while a related line of work examined divergences in discourse phenomena (Šoštarić et al., 2018). Other works studied the ability of a given grammar formalism to capture CLMD in a parallel corpus (e.g.,

object construction to an oblique construction, often involve morphological processes, such as adding a case ending.

²The resource can be found at <https://github.com/macleginn/exploring-clmd-divergences>

Søgaard and Wu, 2009). However, none of these works defined a general methodology for extracting and classifying CLMD.

The only previous work we are aware of to use UD for identifying CLMD is (Wong et al., 2017), which addresses Mandarin-Cantonese divergences by comparing the marginal distribution of syntactic categories on both sides (without alignment). Relatedly, Deng and Xue (DX17; 2017) aligned phrase structure trees over an En-Zh parallel corpus. Notwithstanding the similarity in the general approach, we differ from DX17 in (i) specifically targeting content words, (ii) relying on UD, which is standardized cross-linguistically and allows to simplify the alignment process by focusing on the level of words,³ and (iii) addressing multiple language pairs. It should be noted that the classification of divergences presented in DX17 is rather coarse-grained. Of the seven classes in their study, five (Transitivity, Absence of function words, Category mismatch, Reordering, and Dropped elements) reflect local syntactic differences; one (Lexical encoding) covers many-to-one/one-to-many alignments and non-literal word translations; and the remaining residual type (Structural paraphrase) indiscriminately covers more substantial CLMD. We address this limitation and propose a methodology that automatically derives fine-grained CLMD from aligned annotated corpora and enables straightforward computation of their type statistics.

3 Fine-grained Classification of CLMD

In this section, we present a novel cross-linguistic dataset that provides a high-resolution overview of morphosyntactic differences between pairs of languages and a formal definition of morphosyntactic divergences formulated based on it.

Divergences in the syntax of sentences and their translations can stem from a number of reasons. Setting aside semantic divergences, which are differences in the content expressed by the source and the target (Carpuat et al., 2017; Vyas et al., 2018), the remaining varieties of divergences are essentially different ways to express the same content (Fisiak, 1984; Boas, 2010), which we call CLMD.

We define CLMD empirically to be recurrent divergence patterns in the syntactic structures of sentences and their translations. While content

³Their alignment process involved bottom-up and a top-down passes, sometimes yielding contradictory results.

differences may account for some of the observed syntactic divergences, by aiming for recurring patterns we expect to filter out most such cases, as they are subject to fewer grammatical constraints and should thus not yield systematic patterns of morphosyntactic divergence.

It is harder to distinguish between translation artifacts and CLMD in translated sentences that are due to the genuine differences between grammar and usage. However, translated texts are usually characterized by a higher degree of morphosyntactic transfer and rarely portray the target language as more different from the source language than it needs to be (Koppel and Ordan, 2011; Volansky et al., 2015). Therefore, we do not expect to find spurious recurrent morphosyntactic-divergence patterns introduced by the process of translation.

3.1 The Manually Aligned PUD Resource

Universal Dependencies (UD) is a framework for treebank annotation, whose objectives include satisfactory analyses of individual languages, providing a suitable basis for bringing out cross-linguistic parallelism, suitability for rapid consistent annotation and accurate automatic parsing, ease of comprehension by non-linguists, and effective support for downstream tasks. See Appendix A for a glossary of UD terms.

An important feature of the dependency analysis in UD is that content words are considered the principal components of dependency relations. Within this framework, function words are generally dependents of the content word they relate to most closely. The primacy of content words brings out cross-linguistic parallels that would be harder to detect with other annotation frameworks since function words are highly variable across languages. Importantly, dependency paths between content words do not generally contain function words. As a result, by comparing paths across languages, differences in the surface realization are often masked, and argument structure and linkage differences emphasized.

For example, a preposition accompanying a verb may be dropped in translation if the corresponding verb is transitive (cf. *went around the world* in En vs. *oboshel_{went.around} mir_{world}* in Ru). As prepositions modify the head noun in UD prepositional phrases, the dependency path between the verb and the head noun is not altered.

The Parallel Universal Dependencies (PUD) corpus consists of 1000 sentences translated into various languages by professional translators.⁴

In this paper, we study the Russian, French, Chinese, Japanese, and Korean versions of the PUD corpus, which were each aligned with the corresponding English corpus.⁵ Each parallel corpus was aligned by a human annotator, proficient in the language of the corpus and in English. The UD tokenization is adopted in all cases. Due to the difficulty in finding annotators proficient in pairs of these languages, our annotation takes English as the source language. However, it is possible to obtain an approximate alignment between any pair of these languages, pivoting through English.

Only content words are aligned, so as to sidestep the inherently ambiguous nature of aligning function words across divergent constructions. For details on the function/content distinction we apply to words, see Appendix B. We restrict the alignment to include connected components of the following types: (1) one-to-one alignments, i.e., where a single content word is aligned with another single content word; (2) many-to-one alignments, where multiple source words are aligned with a single target word; (3) one-to-many alignments, where a single source word is aligned with multiple target words. Where a source multi-word expression is translated with a target multi-word expression, we align their headwords, to indicate that their subtrees are in correspondence with one another (e.g., English *with this* and French *par conséquent*). Most of the content words in the corpora were aligned in a one-to-one alignment, which accounts for around 90% of aligned En tokens across the corpora.

3.2 Defining Divergences using UD

We present a framework for defining and investigating translation divergences across a variety of language pairs using UD. The framework operates on a sentence-aligned parallel corpus, where both sides are annotated with UD and content words in corresponding sentences are aligned.

Let $T_s = (V_s, E_s)$, $T_t = (V_t, E_t)$ be a pair

⁴Half of the sentences in the corpus are taken from news articles and the other half from Wikipedia. 750 of the sentences were originally in English, 100 in German, 50 in French, 50 in Italian, and 50 in Spanish. All sentences were translated to other languages via English.

⁵Occasionally highly divergent translations prohibited constructing an alignment. 999 sentences were aligned for En-Fr, En-Zh, and En-Jp, 995 for En-Ru, and 884 for En-Ko.

of UD trees over corresponding sentences, and let $CW_s \subset V_s$ and $CW_t \subset V_t$ be the sets of content words in T_s and T_t respectively. Let $A \subset CW_s \times CW_t$ be a token-level alignment, consisting of one-to-one, many-to-one, and one-to-many alignments. There are two ways to restrict the definition of correspondences between nodes and edges in T_s and T_t : (1) by considering only one-to-one edges or (2) by defining a one-to-one correspondence $A' \subset A$ by traversing all many-to-one alignments $C = \{(v_1, u), (v_2, u), \dots, (v_k, u)\} \subset A$, and selecting for A' only (v_i, u) , where v_i is the highest node in T_s among the nodes in C . The same is done for one-to-many alignments.⁶ The first approach is preferable for analyzing syntactic-path correspondences and was followed in this presentation. The second approach is more suitable for analysis of POS mappings, where headwords are more prominent.

We then define *Corresponding Syntactic Relations* (CSR) as a pair (R_s, R_t) such that R_s and R_t are dependency paths in T_s and T_t , and such that the origin and endpoint of R_s are in CW_s and the origin and endpoint of R_t are their aligned tokens in CW_t according to A' . If the origin or the endpoint of R_s do not have a corresponding node in T_t , R_s does not have a corresponding relation in T_t . The types of R_s and R_t are the sequence of labels on the edges of the paths, optionally along with their directionality in the tree (linear order is not taken into account). Without loss of generality, we assume that R_s begins at the leftmost word of the pair in the En sentence, and R_t by definition begins at the target word corresponding to the leftmost source word. For brevity, we only present results where directionality is not taken into account. Relations are thus written as sequences of UD edge labels separated by the '+' sign.

Token pairs that do not share a POS tag and CSR not of the same type are said to form a divergence. One-to-many and many-to-one alignments are another form of divergence.

4 Empirical Study of Divergences

We apply the proposed methodology to the aligned PUD. We compare syntactic relations, analyzing correspondences between POS tags as well as cor-

⁶Cases of non-unique highest nodes are generally rare in PUD and are thus excluded to simplify the analysis. The only frequent case is the Fr discontinuous negation marker *ne... pas*, generally corresponding to *not*.

respondences between single-edge relations in English and target-side dependency paths.

4.1 Parts of Speech

We begin by examining the mappings of the POS of corresponding tokens (see Appendix C.1 for the full percentage and count matrices). We find that En POS tags of content words are mostly stable in translations to Fr and Ru (sums of the values on the main diagonals account for 78 and 77% of the total number of word pairs respectively). Notable exceptions are the negative particle *not*, which is in a one-to-many alignment in French with *ne* and *pas*, certain types of auxiliaries analyzed as verbs in both Ru and Fr, and proper nouns, which often get mapped to Fr nouns (cf. § 9 and discussion in [Samardžić et al., 2010]).

The En-Zh matrix presents more divergences with only 65% of the alignment edges connecting tokens with the same POS. 11% of nouns were translated as verbs (the reverse mapping is found, albeit to a lesser extent, in all three corpora). Most of such cases involve names of actions and agents (*borrowing*, *ruler*, etc.). En negative particles are split between Zh adverbs, verbs, and auxiliaries; adjectives are quite often mapped to nouns, which form parts of compounds (e.g., *social media* → *shèjiāo méitǐ*, lit. 'social-interaction media'). Adpositions involving spatial relations (the only type of adpositions we consider as content words) are predominantly mapped to adverbs.

The En-Ko matrix is even more divergent: only 62% of the alignment edges connect matching POS. The most striking property of the En-Ko POS matrix is that NOUN serves as a "sink" for other POS: 27% of En adverbs, 56% of En adjectives, and 60.5% of En verbs correspond to Ko nouns. For example, En *trying (to do something)* corresponds to Ko *misu* 'attempt'. As we will show in the next section, this is due to drastic syntactic divergences in En-Ko.

The En-Jp matrix is similar: 62.4% of the edges connect matching POS. Verbs are mostly translated as verbs (58.1%), which shows more affinity between En and Jp basic clause structure. However, adjectives still mostly turn into nouns (53.7%), and adverbs are quite likely to get translated by a noun (16.4% vs. 25.8% for adverb→adverb).

Both Ko and especially Jp tend to leave En pronouns unaligned (15% and 59% respectively), up-

holding their reputations as “radically pro-drop” languages (Neeleman and Szendrői, 2007). Interestingly, Zh, another classical example of this phenomenon, loses only 9% of the pronouns. Ru, a mildly pro-drop language, loses 4% of the pronouns, while the non-pro-drop Fr loses only 2%. This demonstrates the fine granularity of distinctions an empirical approach to CLMD can yield.

4.2 Divergences in Syntactic Relations

Table 2 presents the matrices of target-side syntactic relations that correspond to single-edge source-side relations in the five parallel corpora.

Several observations can be made. First, the En-Fr and En-Ru matrices are similar and are dominated by the elements on the main diagonal (60% of the total number of edges in En-Fr and 55% in En-Ru). An exception are `compound` (which in En are mostly noun compounds), as Ru does not have a truly productive nominal compounding process and Fr compounds are annotated as other relations in UD (Kahane et al., 2017). The other three matrices are less dominated by the entries on the main diagonal (46% of the alignments in En-Zh, 32% in En-Ko, 25.8% in En-Jp) and show higher entropy in most rows, especially in `nmod`, `amod`, `obl`, and `xcomp`, `compound` again being a notable exception (entropy matrices for all relations can be found in Appendix D).

Adverbial clauses (`advcl`) have relatively low values on main diagonals and a high percentage of single edges corresponding to multi-edge paths. This reflects the wide semantic range of `advcl`: in addition to modifying the matrix predicate (*died by drowning*), they can also denote sequential and parallel events (*published a paper sparking a debate*). The latter two cases naturally give rise to `conj` and complement clauses (cf. *published a paper and sparked a debate* / *published a paper to spark a debate*), the most common other path in En-Ru and En-Fr respectively. As we show in § 5.2, there is also a converse phenomenon: sequences of events represented using coordinated clauses, `ccomp`, or `xcomp` in En are translated with `advcl` in East Asian languages.

Of particular interest are the differences between En-Ko and En-Jp confusion matrices. Japanese and Korean are largely similar from the point of view of language typology (SOV word order, topic prominence, agglutinative morphology), but there are also important differences on

the level of usage. Thus, the adjective class in Korean is less productive, and translations often resort to relative clauses for the purposes of nominal modification. Another difference is the fact that Japanese has few compounds as those are usually translated as `nmod` with a genitive particle, while Korean translates nearly all En compounds as compounds. See the discussion of further differences in the next section.

5 Qualitative Analysis of Divergences

In this section, we analyze prominent cases of divergences revealed by applying our method, attempting to demonstrate how fine-grained CLMD may be detected from the confusion matrices and shedding light on what challenges are involved in bridging these divergences (e.g., for the purposes of MT or cross-lingual transfer). Some of the divergences arise due to real differences between grammars; others are largely due to inconsistent application of the UD methodology.

5.1 Nominal and Adjectival Modifiers

When inspecting the translation correspondents of adjectives, we find that while in En-Fr and En-Ru the adjective classes are mostly overlapping, this is not the case for Zh, Ko and Jp. Instead, translation into these languages shifts probability mass from adjectives to nouns: nouns are hardly ever translated to adjectives, but adjectives are more likely to be translated to nouns than remain adjectives. This trend is related to a preference to translate adjectives into possessives (e.g., *Korean company* → Jp: *Kankoku no kaisha* lit. ‘a company of Korea’) or compounds (e.g., *European market* → Jp: *Ōshū ichiba* lit. ‘Europe market’).

5.2 Nominal Subjects

The confusion matrix shows that En `nsubj` demonstrates very different multi-edge mappings into European languages (Fr and Ru) as opposed to East Asian ones (Zh, Jp, and Ko). The “most common other path” for both Russian and French is `xcomp+nsubj`, which is easy to explain: PUD corpora of these languages “demote” fewer auxiliary predicates than English (criteria for demotion are formulated in terms of superficial syntax and differ between languages) and more often place the dependent predicates as the root. Therefore, in constructions like *he could do something* the direct edge between the subject and the verb of the

dependent clause is replaced with two edges going through the modal predicate.⁷

In Zh, Ko and Jp, however, there is another issue: sequential events described using coordinated conjuncts and `xcomp` in En are analyzed as being described with temporal or causal subordinate clauses (*Kipling met and fell in love with Florence Garrard* → Ko: *Kipeulring manna_{meet.subordinate} sarange_{in.love} ppajyeosseumyeo_{fell}*, lit. ‘having met, fell in love’). This makes the direct `nsubj` edge in En correspond to an Ko `nsubj` edge within a subordinated clause, and thus a `nsubj+advcl` path. Given that not all coordinated verbs are translated using a subordinate clause in Ko and Zh, bridging these divergences is likely to require more than a simple tree-transformation rule but possibly refinement of UD’s categories to more abstract linkage types.

5.3 Modal Auxiliaries in Korean

UD treats En modal verbs, such as *can* or *may*, as `aux`, which are dependent on the lexical verb (e.g., *could* $\xleftarrow{\text{aux}}$ *do*). Corresponding verbs in other languages are often treated as simple verbs (for example, all Ru modal verbs are simple verbs in UD).

Even more drastically, Ko routinely expresses this semantics by using an existential construction with the literal meaning ‘(for) X there was a possibility of doing Y’ (instead of *X could do Y*), which converts the En `aux` into `nsubj+acl`. In this case, a tree transformation seems to be sufficient to bridge this divergence.

5.4 `nmod`→`acl+X` in Korean

Ko also differs from other languages in the extent that it uses relative clauses for nominal modification. Table 2 shows that `nmod` has a high percentage for “other” mappings (48%). Investigation of this long tail shows that to a large extent it consists of `acl`-based constructions: `acl+advmod`, `acl+nsubj`, `acl+obj`. Added to `acl`, the cumulative share of `acl`-based constructions is on par with `compound`, the main correspondent of this relation for non-possessive `nmod` (possessive `nmod` are the only ones that map to `nmod` in Ko). This discrepancy is due to the fact that Ko nearly

⁷Cf. also 23 En `nsubj` edges mapped to Ru `nsubj+obl`. Inspection of these sentences reveals that the CLMD can be ascribed to metaphorical usage (e.g., the sense of *read* employed in *the post reads* has no direct correspondent in Ru). Some such cases can be disambiguated using existing annotation schemes.

obligatorily adds contextually-predictable predicates to oblique relations such as *actions [taken] in Crimea* or *people [being] without children*. The Korean PUD does not demote these verbs to functional-word status (such an approach is advocated for in [Gerdes and Kahane, 2016]) but turns them into clause-heading verbs, thus yielding an `acl+X` divergence.⁸

Language		ru	fr	zh	ko	jp
Thematic	Full	25	4	8	5	5
	<code>nsubj to obj/obl</code>	78	57	43	25	53
Promotional		0	0	0	0	0
Demotional		10	2	4	19	1
Structural		83	67	17	0	35
Conflational		10	5	5	6	2
Categorical	<code>nsubj+obj</code>	8	12	23	11	4
	<code>nsubj+(i) obj/obl</code>	51	34	25	11	13
#Sentences		995	999	999	884	999

Table 1: Prevalence in absolute counts of translation divergences as defined in Dorr (1994). Row headings are explained in § 6.

6 Revisiting Dorr’s Divergences

We turn to show that the types of the seminal classification of divergences from Dorr (1994) can be, with a single exception, recast in terms of UD divergences. We quote the original formulations of the divergences illustrated through English-Spanish or English-German examples.

Thematic divergence: E: *I like Mary* ⇔ S: *María me gusta a mí* ‘Mary pleases me.’ In UD, this corresponds to the situation when the original `obj` or `obl` becomes the `nsubj` and vice versa. The divergence will correspond to a CSR of type (*nsubj, obj*) or (*nsubj, obl*). A “full” thematic divergence will also involve the inverse divergence (*obj, subj*) or (*obl, subj*).

⁸The list of examples we can discuss goes on. For example, while investigating the cross-linguistic patterning of English `advcl`, we noticed that it often gets mapped to `ccomp` in French and `acl` in Russian. Both divergent annotations seem to be erroneous as the sentences they appear in are covered by the definition of `advcl` provided in the UD manual. However, the French case is interesting in that the source `advcl` in question can be characterized semantically: instead of denoting a secondary action, they reflect a sequence of events or parallel scenes (e.g., *Columbus sailed across the Atlantic... sparking a period of European exploration of the Americas*). Another problem is presented by multi-word expressions analyzed as proper nouns where all tokens have the same POS tag. The UD manual advises to retain the original parts of speech in proper nouns consisting of phrases (e.g., *Cat on a Hot Tin Roof*) but allows to treat “words that are etymologically adjectives” as PROP in names such as *the Yellow Pages*. When such names are translated, PROP get reanalyzed as ADJ, NOUN, etc., producing spurious CLMD.

Promotional divergence: E: *John usually goes home* ⇔ S: *Juan suele ir a casa* ‘John tends to go home.’ This corresponds to the situation where the original root predicate becomes an `xcomp`, and the original `advmod` takes its place as the root. Corresponding CSR type: (*advmod, xcomp*).

Demotional divergence: E: *I like eating* ⇔ G: *Ich esse gern* ‘I eat likingly.’ The original `xcomp` becomes the root predicate, and the original root predicate is demoted to the position of an `advmod`. The relevant CSR type is (*xcomp, advmod*).

Structural divergence: E: *John entered the house* ⇔ S: *Juan entró en la casa* ‘John entered in the house.’ The original `obj` becomes an `obl`. CSR type: (*obj, obl*).

Conflational divergence: E: *I stabbed John* ⇔ S: *Yo le di puñaladas a Juan* ‘I gave knife-wounds to John.’ The original root predicate is in a one-to-many alignment with a combination of a root predicate and its `obj`.

Categorial divergence: E: *I am hungry* ⇔ G: *Ich habe Hunger* ‘I have hunger.’ The original root predicate becomes an `obj`. CSR type: (*nsubj, nsubj + obj*).

Lexical divergence: E: *John broke into the room* ⇔ S: *Juan forzó la entrada al cuarto* ‘John forced (the) entry to the room.’ Divergences of this type arise whenever aligned words have at best partially overlapping semantic content and never appear on their own but always with other divergences. The information necessary to ascertain the degree of word-meaning overlap is not embedded into UD or any other cross-lingual annotation scheme; therefore we were unable to provide a formal interpretation of this type of divergence.

Frequencies of Dorr’s Divergences in PUD are presented in Table 1 (except for Lexical divergences, which are hard to formalize). It is evident that these types only account for a small portion of the encountered divergences, the point already made for En-Zh in DX17. It seems then that “hand-crafted” translation divergences, however insightful they may be, receive attention disproportionate to their empirical frequency.

7 Perspectives for Automation

One of the strengths of our approach is that it only relies on UD parses and alignments, for which automatic tools exist in many languages. To demonstrate the feasibility of an automated protocol, we conducted an analysis of the WMT14 En-Zh News Commentary corpus.⁹ We used TsinghuaAligner (Liu and Sun, 2015) and pretrained English and Chinese UD parsers from the StanfordNLP toolkit (Qi et al., 2018). To verify that the aligner we are using is adequate for the task, we aligned the En-Zh PUD corpus pair and checked the resulting precision and recall of the edges corresponding to content words.¹⁰ The results (P=0.86, R=0.32) indicate that the automated approach is able to recover around a third of the information obtained through manual alignment with reasonable precision. Importantly, we find recall to be nearly uniform for all source edge types, which suggests that the low recall can be mitigated by using a larger corpus without biasing the results.

The POS and edge-type confusion matrices built from this experiment are very similar to the ones reported in this paper (save for `compound`, which is not produced by the Stanford Zh parser), and are not reproduced here (they can be found in the Supplementary Materials).

8 Applicability for Zero-Shot Parsing

We come to demonstrate the applicability of our method for analyzing the performance of a downstream cross-lingual transfer task. We consider zero-shot cross-lingual parsing (Ammar et al., 2016; Schuster et al., 2019) as a test case and investigate to what extent the performance of a zero-shot parser on a given dependency label can be predicted from its stability in translation. As test sets we use the test sets of GSD UD corpora for the five languages (Ru, Fr, Zh, Ko, and Jp), as well as the corresponding PUD corpora. We train a parser following the setup of Mulcaire et al. (2019) and use a pretrained multilingual BERT (Devlin et al., 2019), feeding its output embeddings into a biaffine-attention neural UD parser (Dozat and Manning, 2017) trained on the English EWT corpus. We evaluate the parser’s ability to predict relation types by computing F-scores for each de-

⁹<http://www.statmt.org/wmt14/>

¹⁰These were defined here as edges with the following labels: `root, nsubj, amod, nmod, advmod, nummod, acl, advcl, xcomp, compound, flat, obj, obl`.

pendency label (save for labels corresponding to function words that were generally not aligned). Appendix E gives full implementation details.

We start by computing Spearman correlations between F-scores and the PRESERVATION indices, defined as the proportion of identity mappings in the confusion matrices for each corpus (e.g., PRESERVATION for `acl` in Ru is 0.48, while in Jp it is 0.37). The correlations are very strong for some languages and noticeable for others ($\rho = 0.62, 0.75, 0.31, 0.42, 0.77$ for Ru, Fr, Zh, Ko, and Jp respectively on GSD test sets, and $\rho = 0.7, 0.82, 0.72, 0.84, 0.68$ on PUD).

We hypothesize that the preservation of a relation in translation is related to the ability of a zero-shot parser to predict it. In order to control for obvious covariates, we introduce two control variables: (1) SOURCE-SIDE HARDNESS (test-set F-scores attained by the parser on English dependency relations) and (2) TARGET-SIDE HARDNESS (F-scores attained by a parser trained on the target-language UD GSD corpus on the target-language test set). We use a mixed-effects model with PRESERVATION, SOURCE-SIDE HARDNESS, and TARGET-SIDE HARDNESS as fixed effects, random intercepts for language, and F-scores for dependency relations as the dependent variable. We then used likelihood-ratio test to compute p -values for the difference in predictive power between the model without PRESERVATION and one with it. The p -value (using Holm correction) is highly significant (< 0.001) for the PUD corpora, and for GSD it is significant with $p = 0.02$.

These results suggest that morphosyntactic differences between languages, as uncovered by our method, play a role in the transferability of parsers across languages. This also underscores the potential utility of bridging CLMD for improving syntactic transfer across languages.

9 Discussion

The presented methodology gives easy access to different levels of analysis. On one hand, by focusing on content words, the approach abstracts away from much local-syntactic detail (such as re-ordering or adding/removing function words). At the same time, the methodology and datasets provide means to investigate essentially any kind of well-defined CLMD. Indeed, since function words in UD tend to be dependents of content words, we may analyze the former by considering the distri-

bution of function word types that each type of content word has. Moreover, sub-typing dependency paths based on their linear direction can allow investigating word-order differences.¹¹

Other than informing the development of cross-lingual transfer learning, our analysis directly supports the validation of UD annotation. For example, we reveal inconsistencies in the treatment of multi-word expressions across languages. Thus, the translation of many NPs with adjectival modifiers, such as *Celtic sea* or *episcopal church*, are analyzed as compounds. Languages such as Ru, lacking a truly productive nominal-compound relation, carve this class up based mostly on the POS of the dependent element (e.g., *episcopal* corresponds to a Ru `amod`), its semantic class (e.g., compounds with cardinal directions are Ru `amods`), and whether the dependent element itself has dependents (these mostly correspond to Ru `nmods`). Our method can be used to detect and bridge such inconsistencies.

In conclusion we note that our analysis suggests that considerable entropy in the mapping between the syntactic relations of the source and target sides can be reduced by removing inconsistencies in the application of UD, and perhaps more importantly by refining UD with semantic distinctions that will normalize corresponding constructions across languages to have a similar annotation. This will simultaneously advance UD’s stated goal of “bringing out cross-linguistic parallelism across languages” and, as our results on zero-shot parsing suggest, make it more useful for cross-linguistic transfer.

Acknowledgments

We thank Nathan Schneider for helpful comments and anonymous reviewers for useful feedback. This work was supported by the Israel Science Foundation (grant no. 929/17).

¹¹Surface-Syntactic Universal Dependencies may be a better fit for such type of analysis, see (Gerdes et al., 2018).

	acl	advcl	advmod	amod	appos	ccomp	compound	conj	fixed	flat	nmod	nsubj	nummod	obj	obl	parataxis	xcomp	Collapsed	Other	MCOP (%)	MCOP		
English-Russian	acl	48	1	2		1					9	2		2	1	1	4	1	29	5	nmod+acl		
	advcl	5	32	1		4		5			1									37	2	advcl+xcomp	
	advmod			63	2						1							5	24	2	2	advmod+nummod	
	amod	3	1	77							5		1						5	7	2	det	
	appos					36			1	13	18	1	1						1	29	6	nmod+flat	
	ccomp	1	2	1		49				1				1			6	6		33	11	ccomp+xcomp	
	compound				33	1		1		6	23		1		1				20	13	3	nmod+nmod	
	conj								76										1	21	2	conj+conj	
	fixed									15	8								69	8	8	nmod+flat	
	flat				3	1		1		75	6								4	9	3	flat+flat	
	nmod			1	4	1				1	59		1	1	3					29	6	det	
	nsubj											1	71		2	2				22	4	nsubj+xcomp	
	nummod			1	5					1	2			82	1				7	2	1	compound+nummod	
	obj			1	1					1	2		8	4	57	9			4	15	4	iobj	
	obl		1	5						4	2		3	50					1	34	7	iobj	
	parataxis		1				1		7								59			31	3	acl+obl	
	xcomp		3	5			4		1		1	2		2	8			51	2	22	8	iobj	
English-French	acl	34	1	4	1	17		1			4			1	1		8	1	26	3	acl+xcomp		
	advcl	1	44	1		9		2							7	2	2			32	4	xcomp+advcl	
	advmod			63	1						2			1	3		1	6	21	4	4	advmod+xcomp	
	amod			1	75	1					9							5	7	1	1	det	
	appos			2	69	1		1			6									21	3	appos+nmod	
	ccomp	1				52								2		5	11			29	10	ccomp+xcomp	
	compound			20	4		7	1		3	42		2	1	1				14	6	1	nmod+nmod	
	conj							77			1								1	18	2	nmod+conj	
	fixed			17	17					17									17	33	17	appos+nummod	
	flat			2	15		2			65	2		4						4	7	3	appos+flat	
	nmod		1	2							66		1	1	3				1	24	3	nmod+nmod	
	nsubj											1	79		1	1				18	5	nsubj+xcomp	
	nummod			1	1						17			71					3	5	3	det	
	obj										5	1			72	8			1	3	9	2	obj+nmod
	obl		2							2	1		6	66					1	1	20	3	obl+nmod
	parataxis	1	1			5		12							6	66		49		31	3	obl+parataxis	
	xcomp		3	3		1	2				1	1		3	4			71	4	9	2	obj+amod	
English-Chinese	acl	32	1	2		2					1	8		2	1		1	1	48	3	obj+advcl		
	advcl		21			2									1			17		57	7	dep	
	advmod	1	1	38	1		1					1		3	4		2	9	38	3	3	advmod+obj	
	amod	4		2	17		31				5	3	3						14	18	3	compound+compound	
	appos	1				50	8	2		2	5	2						1	3	27	2	compound+compound	
	ccomp		1	2			37											5		55	4	ccomp+advcl	
	compound				2		42	1		2	3	1							33	14	8	compound+compound	
	conj		10				1	45											2	40	5	dep	
	fixed					2				2	2								94	2	2	appos+flat	
	flat					9	9	1		47									9	25	18	appos+flat	
	nmod	2		1	1		18				24	3	1	4	3			3	39	3	3	compound+compound	
	nsubj											60		3	1					35	8	nsubj+advcl	
	nummod			1						1			43						23	33	24	nummod+clf	
	obj	1	1			1	1					1		59	4			1	4	27	7	ccomp+nsubj	
	obl		1	2		2	1				1	4		16	19				2	2	50	6	obj+advcl
	parataxis		2	1		7												1		88	22	dep	
	xcomp	1	3	7		1	15					1		6	1			26	3	38	13	aux	
English-Korean	acl	37	2	2		2					1	1		2					54	4	acl+nsubj+acl		
	advcl	2	18	7							1		2						71	6	6	advmod+acl	
	advmod	1	2	43	1		3				1		1					4	43	8		aux	
	amod	8		2	12		38				1	8	2	2					10	19	5	det	
	appos	27	1		8		17			3	2	2								40	4	acl+compound	
	ccomp		30	3		2				2										64	8	advcl+advcl	
	compound			1			57	1		3	2	1		1					23	11	2	compound+compound	
	conj	1	12	1			1	53		1	1								1	31	2	conj+compound	
	fixed						29												43	29	14	compound+compound	
	flat	4					17			60	1	1							5	13	3	compound+flat	
	nmod	1		3		1	17				25	3	1	1					1	45	6	acl+advmod	
	nsubj	1		2			1					47		3						45	7	nsubj+advcl	
	nummod			1			2						68						1	27	10	nummod+nmod	
	obj	1	1	6			3					4		52					2	30	6	aux+obj	
	obl	1	1	32			2					4		5					1	54	4	advmod+compound	
	parataxis		2	21																77	12	advcl+advcl	
	xcomp	2	10	14		7	7					1		6					2	50	9	aux	
English-Japanese	acl	33	4			1					7	1		1				4	45	5	nmod+acl		
	advcl		15										4						15	56	7	obj+acl	
	advmod	1	2	15	1		4				1	2	2	1	2				13	46	9	aux	
	amod	9			8		22				16	1	2	1					27	13	2	compound+compound	
	appos	3	2			9	18				21								2	34	11	nmod+nmod	
	ccomp	7	7				13													60	7	aux	
	compound			1			31				15								39	10	3	compound+compound	
	conj	1	5		1						31								2	48	13	nmod+nmod	
	fixed						2				2									90	4	2	case
	flat						14				41								18	20	7	nmod+compound	
	nmod						8				40			1	1				7	35	7	nmod+nmod	
	nsubj						1				2	43		1	2				5	39	5	nsubj+advcl	
	nummod	1					4						52		1				8	21	13	nummod+compound	
	obj		1				4				6	8		44	1				10	21	5	iobj	
	obl			2			2				4	1		3	18				3	53	14	iobj	
	parataxis		6				6													75	13	advcl+obl	
	xcomp	7	11	4			7						2		4	2			4	43	15		

References

- Waleed Ammar, George Mulcaire, Miguel Ballesteros, Chris Dyer, and Noah A. Smith. 2016. Many languages, one parser. *Transactions of the Association for Computational Linguistics*, 4:431–444.
- Joost Bastings, Ivan Titov, Wilker Aziz, Diego Marcheggiani, and Khalil Simaan. 2017. Graph convolutional encoders for syntax-aware neural machine translation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1957–1967, Copenhagen, Denmark. Association for Computational Linguistics.
- Alexandra Birch, Miles Osborne, and Philipp Koehn. 2008. Predicting success in machine translation. In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 745–754.
- Hans Boas, editor. 2010. *Contrastive studies in construction grammar*. John Benjamins.
- Marine Carpuat, Yogarshi Vyas, and Xing Niu. 2017. Detecting cross-lingual semantic divergence for neural machine translation. In *Proceedings of the First Workshop on Neural Machine Translation*, pages 69–79.
- Marco Damonte and Shay B. Cohen. 2018. Cross-lingual abstract meaning representation parsing. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1146–1155. Association for Computational Linguistics.
- Dun Deng and Nianwen Xue. 2017. Translation divergences in chinese–english machine translation: An empirical investigation. *Computational Linguistics*, 43(3):521–565.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Bonnie J. Dorr. 1994. Machine translation divergences: a formal description and proposed solution. *Computational linguistics*, 20(4):597–635.
- Bonnie J Dorr, Rebecca J Passonneau, David Farwell, Rebecca Green, Nizar Habash, Stephen Helmreich, Eduard Hovy, Lori Levin, Keith J Miller, Teruko Mitamura, et al. 2010. Interlingual annotation of parallel text corpora: a new framework for annotation and evaluation. *Natural Language Engineering*, 16(3):197–243.
- Timothy Dozat and Christopher D. Manning. 2017. Deep biaffine attention for neural dependency parsing. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Jacek Fisiak, editor. 1984. *Contrastive linguistics: Prospects and problems*. Mouton.
- Kim Gerdes, Bruno Guillaume, Sylvain Kahane, and Guy Perrier. 2018. SUD or Surface-Syntactic Universal Dependencies: An annotation scheme near-isomorphic to UD. In *Universal Dependencies Workshop 2018*, Brussels, Belgium.
- Kim Gerdes and Sylvain Kahane. 2016. Dependency annotation choices: Assessing theoretical and practical issues of Universal Dependencies. In *Proceedings of the 10th Linguistic Annotation Workshop*.
- Rebecca Hwa, Philip Resnik, Amy Weinberg, Clara Cabezas, and Okan Kolak. 2005. Bootstrapping parsers via syntactic projection across parallel text. *Natural Language Engineering*, 11:311–325.
- Sylvain Kahane, Marine Courtin, and Kim Gerdes. 2017. Multi-word annotation in syntactic treebanks-propositions for universal dependencies. In *Proceedings of the 16th International Workshop on Treebanks and Linguistic Theories*, pages 181–189.
- Moshe Koppel and Noam Ordan. 2011. Translationese and its dialects. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 1318–1326, Portland, Oregon, USA. Association for Computational Linguistics.
- Mikhail Kozhevnikov and Ivan Titov. 2013. Cross-lingual transfer of semantic role labeling models. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1190–1200, Sofia, Bulgaria. Association for Computational Linguistics.
- Yang Liu and Maosong Sun. 2015. Contrastive unsupervised word alignment with non-local features. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI’15, page 2295–2301. AAAI Press.
- Ryan McDonald, Slav Petrov, and Keith Hall. 2011. Multi-source transfer of delexicalized dependency parsers. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 62–72, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- Phoebe Mulcaire, Jungo Kasai, and Noah A. Smith. 2019. Low-resource parsing with crosslingual contextualized representations. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 304–315, Hong Kong, China. Association for Computational Linguistics.

- Ad Neeleman and Kriszta Szendrői. 2007. Radical pro drop and the morphology of pronouns. *Linguistic Inquiry*, 38(4):671–714.
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajič, 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 Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 1659–1666, Portorož, Slovenia. European Language Resources Association (ELRA).
- Edoardo Maria Ponti, Helen O’Horan, Yevgeni Berzak, Ivan Vulić, Roi Reichart, Thierry Poibeau, Ekaterina Shutova, and Anna Korhonen. 2019. Modeling language variation and universals: A survey on typological linguistics for natural language processing. *Computational Linguistics*, 45(3):559–601.
- Edoardo Maria Ponti, Roi Reichart, Anna Korhonen, and Ivan Vulić. 2018. Isomorphic transfer of syntactic structures in cross-lingual nlp. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1531–1542.
- Peng Qi, Timothy Dozat, Yuhao Zhang, and Christopher D. Manning. 2018. Universal dependency parsing from scratch. In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 160–170, Brussels, Belgium. Association for Computational Linguistics.
- Mohammad Sadegh Rasooli and Michael Collins. 2017. Cross-lingual syntactic transfer with limited resources. *Transactions of the Association for Computational Linguistics*, 5:279–293.
- Tanja Samardžić, Lonke Van Der Plas, Goljihan Kashaeva, and Paola Merlo. 2010. The scope and the sources of variation in verbal predicates in English and French. In *Proceedings of the Ninth International Workshop on Treebanks and Linguistic Theories*, pages 199–211, Tartu. Northern European Association for Language Technology (NEALT).
- Tal Schuster, Ori Ram, Regina Barzilay, and Amir Globerson. 2019. Cross-lingual alignment of contextual word embeddings, with applications to zero-shot dependency parsing. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1599–1613, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jana Šindlerová, Zdeňka Urešová, and Eva Fučíková. 2013. Verb valency and argument noncorrespondence in a bilingual treebank. In *Proceedings of the Seventh International Conference Slovko*, pages 100–108.
- Anders Søgaard and Dekai Wu. 2009. Empirical lower bounds on translation unit error rate for the full class of inversion transduction grammars. In *Proceedings of the 11th International Conference on Parsing Technologies (IWPT'09)*, pages 33–36, Paris, France. Association for Computational Linguistics.
- Margita Šoštarić, Christian Hardmeier, and Sara Stymne. 2018. Discourse-related language contrasts in English-Croatian human and machine translation. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 36–48, Brussels, Belgium. Association for Computational Linguistics.
- Elior Sulem, Omri Abend, and Ari Rappoport. 2015. Conceptual annotations preserve structure across translations: A French-English case study. In *Proceedings of the 1st Workshop on Semantics-Driven Statistical Machine Translation (S2MT 2015)*, pages 11–22, Beijing, China. Association for Computational Linguistics.
- Vered Volansky, Noam Ordan, and Shuly Wintner. 2015. On the features of translationese. *Digital Scholarship in the Humanities*, 30(1):98–118.
- Yogarshi Vyas, Xing Niu, and Marine Carpuat. 2018. Identifying semantic divergences in parallel text without annotations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1503–1515, New Orleans, Louisiana. Association for Computational Linguistics.
- Philip Williams, Rico Sennrich, Matt Post, and Philipp Koehn. 2016. Syntax-based statistical machine translation. *Synthesis Lectures on Human Language Technologies*, 9(4):1–208.
- Tak-sum Wong, Kim Gerdes, Herman Leung, and John Lee. 2017. Quantitative comparative syntax on the Cantonese-Mandarin parallel dependency treebank. In *Proceedings of the Fourth International Conference on Dependency Linguistics (Depling 2017)*, pages 266–275, Pisa, Italy. Linköping University Electronic Press.
- Nianwen Xue, Ondřej Bojar, Jan Hajič, Martha Palmer, Zdeňka Urešová, and Xiuhong Zhang. 2014. Not an interlingua, but close: Comparison of English AMRs to Chinese and Czech. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 1765–1772, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Daniel Zeman, Martin Popel, Milan Straka, Jan Hajič, Joakim Nivre, Filip Ginter, Juhani Luotolahti, Sampo Pyysalo, Slav Petrov, Martin Potthast, Francis Tyers, Elena Badmaeva, Memduh Gokirmak, Anna Nedoluzhko, Silvie Cinkova, Jan Hajic jr., Jaroslava Hlavacova, Václava Kettnerová, Zdenka

Uresova, Jenna Kanerva, Stina Ojala, Anna Mäsilä, Christopher D. Manning, Sebastian Schuster, Siva Reddy, Dima Taji, Nizar Habash, Herman Leung, Marie-Catherine de Marneffe, Manuela Sanguinetti, Maria Simi, Hiroshi Kanayama, Valeria de Paiva, Kira Droganova, Héctor Martínez Alonso, Çağrı Çöltekin, Umut Sulubacak, Hans Uszkoreit, Vivien Macketanz, Aljoscha Burchardt, Kim Harris, Katrin Marheinecke, Georg Rehm, Tolga Kayadelen, Mohammed Attia, Ali Elkahky, Zhuoran Yu, Emily Pitler, Saran Lertpradit, Michael Mandl, Jesse Kirchner, Hector Fernandez Alcalde, Jana Strnadová, Esha Banerjee, Ruli Manurung, Antonio Stella, Atsuko Shimada, Sookyoung Kwak, Gustavo Mendonca, Tatiana Lando, Rattima Nitisaroj, and Josie Li. 2017. Conll 2017 shared task: Multilingual parsing from raw text to universal dependencies. In *Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 1–19, Vancouver, Canada. Association for Computational Linguistics.

Chunting Zhou, Xuezhe Ma, Junjie Hu, and Graham Neubig. 2019. Handling syntactic divergence in low-resource machine translation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1388–1394, Hong Kong, China. Association for Computational Linguistics.