

Compositional Evaluation on Japanese Textual Entailment and Similarity

Hitomi Yanaka¹ and Koji Mineshima²

¹The University of Tokyo, Japan ²Keio University, Japan

hyanaka@is.s.u-tokyo.ac.jp, minesima@abelard.flet.keio.ac.jp

Abstract

Natural Language Inference (NLI) and Semantic Textual Similarity (STS) are widely used benchmark tasks for compositional evaluation of pre-trained language models. Despite growing interest in linguistic universals, most NLI/STS studies have focused almost exclusively on English. In particular, there are no available multilingual NLI/STS datasets in Japanese, which is typologically different from English and can shed light on the currently controversial behavior of language models in matters such as sensitivity to word order and case particles. Against this background, we introduce JSICK, a Japanese NLI/STS dataset that was manually translated from the English dataset SICK. We also present a stress-test dataset for compositional inference, created by transforming syntactic structures of sentences in JSICK to investigate whether language models are sensitive to word order and case particles. We conduct baseline experiments on different pre-trained language models and compare the performance of multilingual models when applied to Japanese and other languages. The results of the stress-test experiments suggest that the current pre-trained language models are insensitive to word order and case marking.

1 Introduction

Natural Language Inference (NLI) (Dagan et al., 2006; Bowman et al., 2015) and Semantic Textual Similarity (STS) (Agirre et al., 2016) tasks are well positioned to serve as a basic benchmark for natural language understanding. With the recent progress of deep neural networks, including pre-trained language models such as BERT (Devlin et al., 2019), the development of benchmark datasets has centered on large crowdsourced English datasets, such as SNLI (Bowman et al., 2015) and MultiNLI (Williams et al., 2018). Since there has been an increasing need for benchmark datasets in linguistic universals (Linzen, 2020), general language understanding frameworks including NLI and STS for languages other than

English have been provided (Conneau et al., 2018; Liang et al., 2020; Le et al., 2020; Shavrina et al., 2020; Xu et al., 2020; Seelawi et al., 2021; Park et al., 2021).

Another recent line of work has investigated whether models are sensitive to shuffled word order, but the conclusions are controversial (Ravfogel et al., 2019; Sinha et al., 2021a; Sinha et al., 2021b; Gupta et al., 2021; Pham et al., 2021; White and Cotterell, 2021). One characteristic of human-like language understanding is that humans can understand sentences according to their word meanings and syntactic structures, then recognize their semantic relationships (Frege, 1963; Katz and Fodor, 1963; Montague, 1973; Janssen and Partee, 1997). Since previous work has demonstrated the usefulness of analyzing the generalization ability of models in challenging NLI in English (Naik et al., 2018; Glockner et al., 2018; McCoy et al., 2019; Rozen et al., 2019; Goodwin et al., 2020; Yanaka et al., 2021), we should continue this line of research in other languages.

Against this background, we provide a Japanese NLI/STS dataset to analyze language models in compositional inference across languages. Our motivations for focusing on Japanese are two-fold. First, Japanese is a high-resource language that has typologically different characteristics from English (Joshi et al., 2020), yet it has not been included in previous cross-lingual (Real et al., 2018; Hu et al., 2020; Ham et al., 2020; Wijnholds and Moortgat, 2021) or multilingual (Conneau et al., 2018) NLI datasets. This raises the question of whether models perform inference differently in Japanese and other languages. Second, Japanese has case markers and free word-order (Hinds, 1986; Shibatani, 1990), phenomena that pose interesting challenges for multilingual NLI. While shuffling data usually changes its meaning, the meaning of a Japanese sentence can be preserved even when the order of noun phrase (NP) arguments is swapped. By analyzing model behavior with scrambling phenomena that preserve case

relations in a sentence and particle-swapping phenomena that change case relations, we can analyze whether the model can distinguish transformations that change sentence meanings and perform compositional inferences.

This paper has three contributions. First, we provide JSICK as a compositional Japanese NLI/STS dataset by manually translating the English SICK dataset (Marelli et al., 2014). Compared with recent crowdsourced NLI datasets, SICK facilitates identification of which compositional linguistic phenomena are key to a given inference. Such a controlled structure is suited to transforming sentences for further analyses of model behavior. In addition, SICK has been translated into non-English languages (Real et al., 2018; Wijnholds and Moortgat, 2021), allowing cross-language comparisons on a sizeable parallel NLI corpus.

Second, we create a stress-test dataset for JSICK to investigate whether language models capture word order and case particles in Japanese. We created the stress-test dataset by transforming syntactic structures of JSICK sentence pairs, where we analyze whether models consider word order and case particles when predicting entailment labels and similarity scores.

Third, for the baseline evaluation of pre-trained language models, we compare performance between different pre-trained language models on JSICK. We also compare the performance of multilingual pre-trained language models on SICK datasets of different languages, including JSICK. We also provide an in-depth analysis of sensitivity to word order and case particles based on the JSICK stress-test dataset. The analysis results suggest that both Japanese and multilingual models are surprisingly inattentive to word order and case marking. Our dataset is publicly available at <https://github.com/verypluming/JSICK>.

2 Related Work

Standard NLI benchmarks have been mainly developed for English. Recently, large crowdsourced NLI datasets derived from image captions, such as SNLI (Bowman et al., 2015), and those targeting multi-genre sentences, like MultiNLI (Williams et al., 2018), have been widely used to evaluate neural models. For linguistics-oriented datasets,

FraCaS (Cooper et al., 1994) is a manually collected NLI test set involving linguistic phenomena studied in formal semantics, and SICK (Marelli et al., 2014) is a larger and more naturalistic NLI/STS dataset made from captions focusing on compositional inference. Unlike SNLI and MultiNLI, SICK was designed by linguistic experts so as to not require dealing with aspects beyond the scope of compositional inference (e.g., world knowledge, named entities, and multiword expressions) but to cover a variety of combinations of lexical, syntactic, and semantic phenomena. The SICK dataset thus allows systematic assessments of the reasoning ability of models on compositional inference. For STS, the SemEval 2012–2017 (Agirre et al., 2016; Cer et al., 2017) competitions provided English, Arabic, and Spanish STS datasets including SICK.

With the development of multilingual pre-trained language models, general language understanding frameworks for languages other than English have been created (Liang et al., 2020; Le et al., 2020; Shavrina et al., 2020; Xu et al., 2020; Seelawi et al., 2021; Park et al., 2021), and NLI datasets have been multilingualized. Conneau et al. (2018) provided a cross-lingual NLI (XNLI) corpus by translating MultiNLI into 15 languages, including languages with few language resources such as Swahili and Urdu and languages with flexible word order such as Russian and German. Ham et al. (2020) translated MultiNLI into Korean to create KorNLI. However, Japanese is not included in these datasets. In addition, since sentences in MultiNLI are usually longer than those in SICK and contain multiword expressions beyond the scope of compositional inference, it is unrealistic to carefully transform syntactic structures of sentence pairs in XNLI to create a stress-test dataset like ours.

As examples of other non-English datasets, OCNLI (Hu et al., 2020) is a Chinese NLI dataset built from original multi-genre resources. FarsTail (Amirkhani et al., 2020) is a Persian NLI dataset containing sentences from university exams. There have also been attempts to translate the SICK dataset into Portuguese (Real et al., 2018) and Dutch (Wijnholds and Moortgat, 2021), so our Japanese SICK dataset will contribute to a multilingual SICK dataset that will allow controlled, cross-lingual analyses of the compositional abilities of language models.

Regarding Japanese NLI datasets, a Japanese SNLI dataset (Yoshikoshi et al., 2020) was constructed by using machine translation to translate the English SNLI dataset into Japanese and automatically filtering out unnatural sentences, but methods that employ machine translation are still problematic in that they can produce unnatural sentences. The Japanese Realistic Textual Entailment Corpus (Hayashibe, 2020) is a crowdsourced dataset containing Japanese hotel reviews. However, linguistic phenomena in these Japanese datasets demonstrate limited diversity because the sentences they contain are restricted to simple structures. Kawazoe et al. (2017) provided JSeM, a manually curated test set including a Japanese version of FraCaS to diagnose inference systems from a formal semantics perspective. We produced an NLI dataset by asking experts to translate the SICK dataset into Japanese, thus maintaining both sentence naturalness and compositions of linguistic phenomena.

While recent works (Sinha et al., 2021a,b; Gupta et al., 2021; Pham et al., 2021; Hessel and Schofield, 2021) have shown that pre-trained language models are insensitive to word order on permuted English datasets in the standard natural language understanding benchmark GLUE (Wang et al., 2019) including NLI, other works have analyzed sentence perplexity with varying word orders and shown controversial results regarding inductive biases for word order in different languages (Ravfogel et al., 2019; White and Cotterell, 2021). For six languages, including Japanese, Yang et al. (2019) evaluated whether multilingual BERT captures word order in the translated PAWS dataset (Zhang et al., 2019), involving adversarial paraphrase identification pairs whose sentences share words but differ in word order. Experiments showed that BERT performance for Japanese is consistently worse than that for Indo-European languages. Our study deepens insight into the causes of performance differences by stress-test evaluation of NLI and STS tasks through careful manipulation of case markers, which are more challenging tasks than are two-class paraphrase identification tasks. Kuribayashi et al. (2021) reexamined the general hypothesis that language models with lower perplexity are more human-like in Japanese than in English, and the results have shown the necessity of evaluating models across languages.

Analyzing the model behavior, rather than perplexity, in transformed Japanese inference should provide new insights into the model’s sensitivity to word order.

3 JSICK Dataset Creation

3.1 Translation

The original SICK dataset uses 6,077 sentences to provide 9,927 sentence pairs (A, B). To create the JSICK dataset, we first asked an expert translator to translate the 6,077 English sentences in SICK into Japanese. The translator did not see entailment labels, instead just translating a list of English sentences sorted alphabetically. The translations were independently validated by English–Japanese bilinguals, and no examples were discarded. To avoid changing sentence meaning during translation, we prepared translation guidelines and asked the translator to translate English sentences into natural Japanese while maintaining diversity in lexical, syntactic, and semantic phenomena such as hypernym–hyponym relations, active–passive alternations, and quantification in the original English sentences. Note that sentences in the JSICK dataset contain some translations that are unnatural due to cultural factors, but reflecting culture in translation is beyond the scope of analyzing the compositional inference ability of models.

The guidelines explain how to translate linguistic phenomena, including indefinite and definite articles, singular and plural nouns, passive verbs, negation, and quantification. We also asked the translator to try to keep word orders as consistent as possible with the original sentences. Some instructions from our guidelines are given in detail below.

Indefinite/Definite Articles The following describes our instructions regarding the distinction between indefinite and definite articles. The distinction between indefinite and definite articles is an important phenomenon that affects interpretations of quantification (Hawkins, 1978; Heim, 1982). However, since Japanese does not have articles (Hinds, 1986; Shibatani, 1990), it is not obvious how to translate indefinite articles as in (1) and definite articles as in (2) into Japanese. We therefore translated subject NPs as bare noun phrases, using the particle $が$ (ga) when translating the nominative case involving an indefinite article, and using the particle (topic marker) $は$ (wa)

when translating the nominative case involving a definite article. Since the majority of sentences in the SICK dataset are episodic, we can correctly translate English sentences into Japanese by the above rule.

- (1) 男性が ギターを 弾いていない
man Nom guitar Acc playing is not
'A man is not playing a guitar'
- (2) 男性は ギターを 弾いていない
man Topic guitar Acc playing is not
'The man is not playing a guitar'

Singular and Plural Nouns In examples like (3), we can translate plural nouns by adding a plurality suffix such as *たち* (*-tachi*). However, Japanese does not have a general way to form plural words like the *-s* suffix in English. Thus, as in (4), we prioritized sentence naturalness by not adding the plural suffix *たち* to the accusative case of words like *エビ* (*shrimps*).

- (3) 男性たちが 木 を 切っている
men Nom wood Acc cutting are
'Men are cutting wood'
- (4) 女性が エビ を ゆでている
woman nom Shrimps Acc boiling is
'A woman is boiling shrimps'

3.2 Validation

There are issues with the gold labels in the original SICK dataset (Bowman et al., 2015; Kalouli et al., 2017). In addition, translation from English to Japanese can change the appropriateness of the entailment label used in English (see Section 3.4). Thus, instead of using the original gold labels, we used the crowdsourcing platform Lancers¹ to re-annotate entailment labels and similarity scores for JSICK. Definitions of entailment labels (*entailment*, *contradiction*, and *neutral*) and similarity scores in a range of 1 (completely unrelated) to 5 (very related) for a pair (*A*, *B*) of sentences are the same as those for the original SICK. In our instructions, we noted that sentences *A* and *B* describe the same situation or event to avoid any indeterminacy of event and entity coreference that might cause inconsistencies in *contradiction* labels (Bowman et al., 2015).

¹<https://www.lancers.jp/>.

The annotators were six native Japanese speakers, randomly selected from the crowdsourcing platform. The authors annotated the gold labels with ten examples in the JSICK trial set (500 examples) to provide ten test questions. We asked the annotators to fully understand the guidelines to the point where they could assign the same labels as gold labels for all ten test questions. We adopted annotations that were agreed upon by a majority vote as gold entailment labels and adopted the average of the annotation results as gold similarity scores. For entailment labels, the authors also manually checked whether the majority judgement vote was semantically valid for each example. Since recent work (Pavlick and Kwiatkowski, 2019; Gantt et al., 2020) has demonstrated the importance of information for modeling disagreements in NLI datasets, we publicly release the raw annotations with the JSICK dataset.

The average annotation time was 1 min per pair, and Krippendorff's alpha for the entailment labels was 0.65. There were 6,957 cases (70.1%) in which three annotators assigned the same entailment labels, 2,922 cases (29.5%) in which two annotators assigned the same entailment labels, and 48 cases (0.4%) in which the labels of all three annotators assigned different labels. For cases where the labels of all three annotators assigned different labels, the labels were determined by the consensus of the authors.

3.3 Linguistic Tags

To analyze the ability of models to capture various linguistic phenomena, we annotated the JSICK dataset with linguistic phenomenon tags. We provided a set of nine linguistic tags for linguistic phenomena: numerals, negation, quantification, passive voices, anaphora, conjunction, disjunction, modal, and additive particle. We automatically annotated multiple tags with premise-hypothesis pairs, using Janome² to process each premise and hypothesis sentence for morphological analysis and part-of-speech tagging. If results of either of the morphological analyses included phrase patterns related to a linguistic tag, the premise-hypothesis pair is annotated with that tag.

²<https://github.com/mocobeta/janome>.

ID	premise-hypothesis pair (A, B)	Linguistic phenomena	Entailment	Similarity
5862	A: [二人 ^{NUM}]の女性が群衆の前でダンスをし[ながら ^{CONJ}]歌っている [Two ^{NUM}] women are dancing [and ^{CONJ}] singing in front of a crowd B: [二人 ^{NUM}]の女性が[多く ^{QUANT}]の人の前でダンスをし[ながら ^{CONJ}]歌っている [Two ^{NUM}] women are dancing [and ^{CONJ}] singing in front of [many ^{QUANT}] people	QUANT/NUM/CONJ	entailment	4.7
352	A: [別 ^{NUM}]の犬を追いかけけている[か ^{DISJ}]、または口に棒きれをくわえている 犬は[一 ^{NUM}]匹もい[ない ^{NEG}] There is [no ^{NEG}] dog chasing another [or ^{DISJ}] holding a stick in its mouth B: 犬が[別 ^{NUM}]の犬を追いかけけていて、口に棒きれをくわえている A dog is chasing [another ^{NUM}] and is holding a stick in its mouth	NUM/DISJ/NEG	contradiction	3.9
765	A: その子供は幸せ[そう ^{MODAL}]に雪の中で滑っている The kid is [happily ^{MODAL}] sliding in the snow B: 雪で覆わ[れ ^{PAS}]た丘の上にいる男の子が赤いジャケット[と ^{CONJ}]黒い帽子を身に着け、ひざまづいて滑っている A boy on a hill [covered ^{PAS}] in snow is wearing a red jacket [and ^{CONJ}] a black hat and is sliding on his knees	MODAL/CONJ/PAS	neutral	2.5

Table 1: Examples from the JSICK dataset. Each ID corresponds to the ID in the original SICK dataset.

Phenomenon	Train	Dev	Test	Total
Numeral (NUM)	1,374	151	1,513	3,038
Negation (NEG)	1,096	107	1,140	2,343
Quantification (QUANT)	698	81	744	1,523
Passive (PAS)	649	83	695	1,427
Anaphora (ANA)	612	74	700	1,386
Conjunction (CONJ)	558	69	640	1,267
Disjunction (DISJ)	364	53	428	845
Modal (MODAL)	62	4	69	135
Additive particle (ADDP)	7	1	13	21

Table 2: Distribution of linguistic tags in JSICK.

Gold label	Train	Dev	Test	Total
JSICK-NLI				
Entailment	969 (21.5)	122 (24.4)	1,088 (22.1)	2,179 (22.0)
Contradiction	743 (16.5)	80 (16.0)	797 (16.2)	1,620 (16.3)
Neutral	2,788 (62.0)	298 (59.6)	3,042 (61.7)	6,128 (61.7)
SICK-NLI				
Entailment	1,299 (28.9)	144 (28.8)	1,414 (28.7)	2,857 (28.8)
Contradiction	665 (14.8)	74 (14.8)	720 (14.6)	1,459 (14.7)
Neutral	2,536 (56.4)	282 (56.4)	2,793 (56.7)	5,611 (56.5)
JSICK-STs				
1-2	614 (13.6)	71 (14.2)	651 (13.2)	1,336 (13.4)
2-3	1,164 (25.9)	111 (22.2)	1,248 (25.3)	2,523 (25.4)
3-4	1,373 (30.5)	155 (31.0)	1,587 (32.2)	3,115 (31.4)
4-5	1,349 (30.0)	163 (32.6)	1,441 (29.2)	2,955 (29.7)
SICK-STs				
1-2	436 (9.7)	37 (7.4)	451 (9.2)	924 (9.3)
2-3	635 (14.1)	69 (13.8)	674 (13.7)	1,378 (13.9)
3-4	1,742 (38.7)	192 (38.4)	1,965 (39.9)	3,899 (39.3)
4-5	1,687 (37.5)	202 (40.4)	1,837 (37.3)	3,726 (37.5)
Total	4,500	500	4,927	9,927

Table 3: Distribution of JSICK and SICK sentence pairs for each gold entailment label and similarity score. Numbers in parentheses are percentages of the entire dataset.

Table 1 shows examples of linguistic tagging in the JSICK dataset, and Table 2 shows the distribution of linguistic tags. Table 5 shows the results of comparing the percentage of linguistic tags in the JSICK test data and the two existing

Similarity	Entailment	Contradiction	Neutral	Total
1-2	0 (0.0)	5 (0.3)	1,331 (21.7)	1,336
2-3	6 (0.3)	160 (9.9)	2,357 (38.4)	2,523
3-4	293 (13.4)	633 (39.1)	2,189 (35.7)	3,115
4-5	1,880 (86.3)	822 (50.7)	253 (4.1)	2,955
Total	2,179	1,620	6,128	

Table 4: Distribution of JSICK sentence pairs across NLI and STS tasks.

Phenomenon	JSICK	JSNLI	JRTEC
Numeral	1,513 (30.7)	1,030 (26.3)	47 (1.2)
Negation	1,140 (23.1)	66 (1.7)	291 (7.5)
Quantification	744 (15.1)	298 (7.6)	185 (4.8)
Passive	695 (14.1)	226 (5.8)	89 (2.3)
Anaphora	700 (14.2)	487 (12.4)	72 (1.9)
Conjunction	640 (13.0)	922 (23.5)	136 (3.5)
Disjunction	428 (8.7)	168 (4.3)	65 (1.7)
Modal	69 (1.4)	103 (2.6)	11 (0.3)
Additive particle	13 (0.3)	6 (0.2)	39 (1.0)
Test total	4,927	3,916	3,885

Table 5: Comparison of linguistic tags between JSICK and previous Japanese NLI datasets. Numbers in parentheses are percentages of the entire test set.

large Japanese NLI datasets mentioned in Section 2, Japanese SNLI (JSNLI) and the Japanese Realistic Textual Entailment Corpus (JRTEC). Compared with previous datasets, JSICK contains more linguistic phenomena, including numerals, negation, quantification, passive voice, anaphora, and disjunction. This indicates that the distribution of linguistic phenomena in the JSICK dataset is well balanced.

3.4 Dataset

Table 3 shows that the distribution of JSICK dataset gold labels is almost the same as that for the English SICK dataset. The distribution of JSICK sentence pairs across NLI and STS

Type	Example (ID)
Singular-Plural	A: ある人がキノコをナイフで切っている (A person is cutting a mushroom with a knife) a person Nom mushroom Acc knife with cutting is
	B: ある人がいくつかのキノコを切っている (A person is cutting some mushrooms) a person Nom some mushroom Acc cutting is ID: JSICK-2590
Lexical Gap	A: ある人がピンク色のロープで岩を登っている (A person is climbing a rock with a rope, which is pink) a person Nom pink is rope with rock Acc climbing is
	B: 一人の男性がロープでその崖を登っている (One man is climbing the cliff with a rope) one man Nom rope with the cliff Acc climbing is ID: JSICK-648

Table 6: Examples of linguistic factors that cause differences in entailment labels between English and Japanese.

tasks (Table 4) also follows the same trend as in SICK; similarity scores for the *entailment* and *contradiction* cases tend to be in the range of 3 to 5, while *neutral* similarity scores are distributed.

The most common cases where translation changed the entailment label from that in the original SICK dataset were those where the labels are changed to *neutral*. There were 242 such examples, due to grammatical differences between English and Japanese. Table 6 shows some typical examples. One major grammatical difference that can change entailment labels is the distinction between singular and plural NPs. In English, the plural form *mushrooms* explicitly indicates that there is more than one mushroom. By contrast, there is no grammatical singular-plural marking in Japanese (Nakanishi and Tomioka, 2004), so the bare noun キノコ (‘mushroom’) can be interpreted as either singular or plural. This caused split entailment judgments among the annotators. Other types of discrepancy are due to various lexical gaps. For instance, in Lexical Gap Example B in Table 6, the English word *man* can be applied to both men and women, while its natural counterpart (男性) in Japanese does not have such a generic meaning. As a result, the entailment label for this example is *neutral*, rather than *entailment*.

4 Baseline Experiments

4.1 Experimental Setup

In this study, we experimented with two Japanese pre-trained language models: Japanese BERT (jaBERT; Devlin et al., 2019) pre-trained on Japanese Wikipedia, and Japanese RoBERTa (jaRoBERTa; Liu et al., 2019) pre-trained on Japanese Wikipedia and the Japanese portion of CC-100: Monolingual Datasets from Web Crawl Data (Conneau et al., 2020).

For jaBERT, we investigated performance differences between the BERT-base model³ pre-trained with 17 million sentences from Wikipedia articles and the BERT-large model⁴ pre-trained with 30 million sentences. The configuration was the same as that for the original BERT model. To check whether methods for tokenization and masked language modeling (MLM) affect model performance, we compared three settings for the BERT-base model. In the (SUBWORD) setting, the model processes input texts with word-level tokenization by the MeCab morphological parser with a standard Japanese dictionary IPAdic (Asahara and Matsumoto, 2003), followed by WordPiece subword tokenization (Schuster and Nakajima, 2012). The vocabulary size was 32,000. In the (WHOLE) setting, the subword model was trained with whole-word masking enabled for the MLM objective. In the (CHAR) setting, the model processed texts with word-level tokenization based on the IPAdic, followed by character-level tokenization.

For jaRoBERTa, we compared the performance of the base model⁵ and the large model.⁶ The input text was segmented into words by the Japanese morphological analyzer Juman++ (Morita et al., 2015; Tolmachev et al., 2018), and each word was tokenized using SentencePiece.⁷

We also analyzed differences in behaviors of the Japanese and multilingual pre-trained language models. As multilingual models, we used the multilingual BERT model (mBERT) trained

³<https://huggingface.co/cl-tohoku/bert-base-japanese>.

⁴<https://huggingface.co/cl-tohoku/bert-large-japanese>.

⁵<https://huggingface.co/nlp-waseda/roberta-base-japanese>.

⁶<https://huggingface.co/nlp-waseda/roberta-large-japanese>.

⁷<https://github.com/google/sentencepiece>.

Model	Train (setting)	Test	Prec	Rec	macro-F1	Acc
Japanese models						
jaRoBERTa-large	JSICK	JSICK	90.1	87.3	88.6	90.3±0.04
	JSICK (hypo-only)	JSICK	21.0	33.3	25.7	62.9±0.09
	JSNLI	JSICK	61.8	69.3	54.9	53.8±0.09
		JSNLI	94.2	94.4	94.2	94.3±0.06
	JSNLI+JSICK	JSICK	86.7	88.2	87.4	89.0±0.09
		JSNLI	93.6	93.7	93.6	93.7±0.09
jaRoBERTa-base	JSICK	JSICK	84.9	87.8	86.2	87.9±0.07
	JSICK (hypo-only)	JSICK	21.0	33.3	25.7	62.9±0.08
	JSNLI	JSICK	60.2	67.4	53.5	52.6±0.07
		JSNLI	93.0	93.0	92.9	93.0±0.06
	JSNLI+JSICK	JSICK	83.0	90.1	85.5	87.9±0.05
		JSNLI	92.4	92.5	92.4	92.5±0.03
jaBERT-large	JSICK	JSICK	86.5	84.8	85.6	87.9±0.03
	JSICK (hypo-only)	JSICK	20.6	33.3	25.5	61.8±0.09
	JSNLI	JSICK	57.8	64.9	52.5	52.2±0.05
		JSNLI	92.6	92.7	92.6	92.7±0.04
	JSNLI+JSICK	JSICK	88.1	88.7	88.4	90.0±0.03
		JSNLI	93.4	93.5	93.4	93.5±0.02
jaBERT-base (WHOLE)	JSICK	JSICK	78.6	81.8	79.9	82.4±0.05
	JSICK (hypo-only)	JSICK	48.8	42.8	44.0	58.9±0.07
	JSNLI	JSICK	57.5	63.7	52.5	52.4±0.05
		JSNLI	94.1	94.2	94.1	94.2±0.03
	JSNLI+JSICK	JSICK	84.4	87.0	85.6	87.5±0.03
		JSNLI	94.3	94.3	94.3	94.3±0.03
jaBERT-base (CHAR)	JSICK	JSICK	76.2	80.6	78.1	80.7±0.05
	JSICK (hypo-only)	JSICK	47.8	46.6	47.0	54.9±0.08
	JSNLI	JSICK	55.5	60.2	47.9	47.9±0.03
		JSNLI	90.7	90.8	90.7	90.8±0.03
	JSNLI+JSICK	JSICK	83.8	85.3	84.4	86.3±0.04
		JSNLI	90.8	90.8	90.8	90.8±0.03
jaBERT-base (SUBWORD)	JSICK	JSICK	76.9	80.9	78.5	80.8±0.06
	JSICK (hypo-only)	JSICK	49.2	38.7	37.8	60.8±0.07
	JSNLI	JSICK	57.3	63.1	48.8	48.3±0.04
		JSNLI	91.5	91.5	91.3	91.5±0.04
	JSNLI+JSICK	JSICK	84.9	82.8	83.7	86.1±0.03
		JSNLI	91.5	91.5	91.4	91.5±0.03
Multilingual models						
XLM-RoBERTa-large	JSICK	JSICK	88.2	86.5	87.2	89.1±0.10
	JSICK (hypo-only)	JSICK	52.0	51.2	50.2	56.1±0.09
	JSNLI	JSICK	61.2	68.4	54.9	53.9±0.09
		JSNLI	94.5	94.6	94.5	94.6±0.04
	JSNLI+JSICK	JSICK	89.2	89.4	89.3	90.8±0.07
		JSNLI	94.0	94.1	94.0	94.1±0.05
XLM-RoBERTa-base	JSICK	JSICK	79.3	68.1	70.2	78.5±0.08
	JSICK (hypo-only)	JSICK	40.3	43.7	45.5	56.8±0.06
	JSNLI	JSICK	56.6	63.2	51.5	51.0±0.09
		JSNLI	92.1	92.2	92.1	92.1±0.05
	JSNLI+JSICK	JSICK	85.9	86.4	86.0	88.1±0.07
		JSNLI	92.0	92.1	92.0	92.1±0.04
mBERT	JSICK	JSICK	88.2	86.4	87.3	89.2±0.08
	JSICK (hypo-only)	JSICK	44.7	36.2	32.9	58.8±0.09
	JSNLI	JSICK	58.2	65.2	52.6	51.9±0.05
		JSNLI	91.8	92.0	91.9	92.0±0.04
	JSNLI+JSICK	JSICK	87.8	87.2	87.5	89.3±0.03
		JSNLI	92.0	92.2	92.1	92.1±0.03

Table 7: Baseline results with Japanese and multilingual pre-trained language models for the NLI task with JSICK and JSNLI (%).

with multilingual Wikipedia and the XLM-RoBERTa-base⁸ and XLM-RoBERTa-large⁹ models (Conneau et al., 2020) pre-trained on CC-100 containing 100 languages. For mBERT, we used a multilingual cased model,¹⁰ as is recommended for languages with non-Latin alphabets, like Japanese. For each setting, we used learning rates $2e^{-5}$, $3e^{-5}$, and $5e^{-5}$ and 3, 4, and 5 training epochs to tune for the best parameters.

For the NLI task, to investigate whether the size and quality of fine-tuned data affect performance, we fine-tuned pre-trained models on three types of training data: (i) JSICK training data (5K), (ii) JSNLI training data (533K), and (iii) both JSICK and JSNLI training data (538K). As mentioned in Section 2, JSNLI is a machine-translated Japanese SNLI dataset. Since both SICK and SNLI are derived from image captions, we hypothesized that JSNLI might improve model performance on the JSICK test set. We used four standard evaluation metrics for NLI tasks: precision (Prec), recall (Rec), macro F1-score (F1), and accuracy (Acc). To analyze whether entailment labels are learned and predicted by referring only to hypothesis sentences, we investigated the performance of models trained on JSICK without the premise sentences. We performed five runs and present the averages below. We also report standard deviations for the accuracy of baseline results in the NLI task. As the baseline for the STS task, we used the Pearson correlation coefficient, Spearman correlation coefficient, and mean square error (MSE) between the prediction results for BERTScore (Zhang et al., 2020), a recent BERT-based model for unsupervised STS, and the gold similarity score.

4.2 Baseline Results

Table 7 shows the evaluation results for NLI models. For all models, the accuracy on JSICK is lower than that on JSNLI, indicating that JSICK poses more challenges than does JSNLI. Since performance under the hypothesis-only setting was low, JSICK does not allow model predictions from hypotheses alone.

In the standard train/test split setting for JSICK, accuracy with the jaRoBERTa-large

⁸<https://huggingface.co/xlm-roberta-base>.

⁹<https://huggingface.co/xlm-roberta-large>.

¹⁰<https://huggingface.co/bert-base-multilingual-cased>.

model had the best performance (acc. 90.3%). Surprisingly, multilingual models such as XLM-RoBERTa-large and mBERT achieved comparable accuracy (89.1% and 89.2%, respectively). Among the multilingual models, the mBERT model had the best performance. For the jaRoBERTa, jaBERT, and XLM-RoBERTa models, those trained on larger texts achieved higher accuracies on NLI tasks. Among the tokenization settings for jaBERT, whole-word masking (WHOLE) provided the highest accuracy (82.4%). Regarding fine-tuning data, mixing the training data with the JSICK and JSNLI training sets improved model performance for the JSICK test set for all models except jaRoBERTa-large. Since the jaRoBERTa-large model trained with a single training set (JSICK or JSNLI) already demonstrated high performance, additional training data did not improve performance.

Table 8 shows the results from the unsupervised STS model. Interestingly, mBERT achieved nearly the same high performance as did jaBERT on the STS task. Among different tokenization settings for jaBERT, the character-based tokenization (CHAR) produced the highest performance. This is due to the difference between NLI and STS tasks. Similarity scores are affected by the token overlap between two sentences, as suggested by the fact that the *contradiction* cases tended to have higher similarity scores. Character-based tokenization allows more precise calculations of the token overlap, and thus might be suitable for STS tasks.

Relevance Between Entailment and Similarity

We next analyze relations between entailment labels and similarity scores in cases where model predictions are difficult. Table 9 shows a distribution of accuracies from the jaRoBERTa-large and mBERT NLI models fine-tuned with JSICK for each similarity score. These results show that both Japanese and multilingual models struggled to predict entailment labels with low similarity scores, but their gold labels are *contradiction*. Both models also failed to predict cases where premise sentences are very similar to their hypothesis sentences but their gold labels are *neutral*. Table 10 shows a distribution of Pearson correlations on the JSICK STS test set for each entailment label. These results show that the STS models have the same trend for *contradiction* examples as do the NLI models; the STS models failed to predict

Score	Japanese models																	
	jaRoBERTa-large			jaRoBERTa-base			jaBERT-large			jaBERT-base								
	γ	ρ	MSE	γ	ρ	MSE	γ	ρ	MSE	(WHOLE)			(CHAR)			(SUBWORD)		
									γ	ρ	MSE	γ	ρ	MSE	γ	ρ	MSE	
1-2	32.9	32.8	1.63	12.8	9.3	1.93	30.0	30.3	1.40	33.3	34.8	1.62	38.1	39.0	1.48	31.9	32.8	1.61
2-3	29.6	29.4	1.34	20.5	21.9	1.93	27.8	27.7	1.12	28.9	29.3	1.33	30.9	30.9	1.22	27.5	27.7	1.32
3-4	34.4	34.2	1.07	20.5	21.9	1.62	32.8	32.0	91.8	32.9	32.4	1.07	34.7	34.9	0.99	32.0	31.7	1.06
4-5	16.7	21.8	72.4	12.5	15.5	0.88	25.0	30.6	70.0	20.0	24.9	0.72	24.5	26.2	0.70	22.0	26.2	0.72
All	74.6	75.3	1.22	65.3	69.1	1.47	71.6	72.1	1.05	73.8	74.1	1.22	77.1	77.1	1.12	72.3	72.6	1.21

Score	Multilingual models								
	XLM-RoBERTa-large			XLM-RoBERTa-base			mBERT		
	γ	ρ	MSE	γ	ρ	MSE	γ	ρ	MSE
1-2	12.8	9.3	1.93	33.0	32.4	1.93	38.1	39.9	1.46
2-3	20.5	21.9	1.62	28.5	28.7	1.62	30.0	29.7	1.20
3-4	28.8	31.8	1.30	33.8	34.8	1.30	35.6	35.8	0.98
4-5	12.5	15.5	0.88	15.6	19.2	0.88	22.2	25.3	0.69
All	65.3	69.1	1.47	75.5	75.7	1.47	77.3	77.4	1.10

Table 8: Baseline results from Japanese and multilingual BERTscore models on the STS task with JSICK (%). γ : Pearson correlation $\times 100$, ρ : Spearman correlation $\times 100$.

Model	Score	Entailment	Contradiction	Neutral	All
jaRoBERTa-large	1-2	-	25.0(1/4)	100.0(647/647)	99.5(648/651)
	2-3	33.3(1/3)	44.6(37/83)	97.8(1136/1162)	94.1(1174/1248)
	3-4	62.4(98/157)	72.5(237/327)	89.8(991/1103)	83.6(1326/1587)
	4-5	87.1(560/643)	97.9(375/383)	67.7(88/130)	88.5(1023/1156)
	All	86.0(936/1088)	81.6(650/797)	94.1(2862/3042)	-
mBERT	1-2	-	25.0 (1/4)	100.0 (647/647)	99.5 (648/651)
	2-3	100.0 (3/3)	45.8 (38/83)	96.8 (1125/1162)	93.4 (1166/1248)
	3-4	61.8 (97/157)	73.4 (240/327)	55.8 (615/1103)	81.9 (1299/1587)
	4-5	84.6 (544/643)	96.6 (370/383)	39.2 (51/130)	86.9 (1005/1156)
	All	84.7 (922/1088)	81.4 (649/797)	92.8 (2825/3042)	-

Table 9: Distribution of accuracies for the JSICK NLI test set for each similarity score.

Model	Label	1-2	2-3	3-4	4-5	All
jaRoBERTa-large	Entailment	-	-62.5	25.0	21.7	36.6
	Contradiction	-83.0	13.3	15.4	13.5	39.5
	Neutral	32.8	27.1	39.6	28.0	68.0
	All	32.9	29.6	34.4	24.5	-
mBERT	Entailment	-	-46.1	28.6	24.8	45.5
	Contradiction	-82.5	18.7	20.9	13.6	45.3
	Neutral	38.0	27.5	38.9	30.9	70.6
	All	38.1	30.0	35.6	23.2	-

Table 10: Distribution of Pearson correlations for the JSICK STS test set for each entailment label.

low similarity scores in cases where a premise sentence contradicts a hypothesis sentence.

Linguistic Phenomena Table 11 shows evaluation results for the jaRoBERTa-large model and the mBERT model for each linguistic tag. Regarding the jaRoBERTa-large model performance, there is little difference for each linguistic phenomenon, and accuracy for examples involving anaphora, disjunction, and additive particle was comparatively low in the NLI task. For

the STS task, Pearson correlations for examples involving quantification and passive were slightly low.

Regarding differences between training data, models fine-tuned with JSICK performed better for almost all linguistic phenomena than did those fine-tuned with JSNLI. Furthermore, adding the JSNLI training set to the JSICK training set did not improve the model performance on most linguistic phenomena. This suggests that training data quality is more critical for learning linguistic phenomena than is quantity.

4.3 Comparison with Other Languages

We next compare the performance of mBERT on the JSICK dataset with that on the original English SICK dataset (SICK-EN), the Portuguese SICK-BR dataset (Real et al., 2018), and the Dutch SICK-NL dataset (Wijnholds and Moortgat, 2021). Since gold labels for the JSICK datasets differ from those for the SICK-EN and SICK-NL datasets (SICK-NL uses the same labels as SICK-EN), we

Phenomena	jaRoBERTa-large				mBERT			
	Similarity	Entailment			Similarity	Entailment		
		I	N	N+I		I	N	N+I
Numeral (1513)	76.4	92.7	50.2	89.4	80.0	91.9	45.8	91.2
Negation (1140)	80.4	93.9	49.5	93.5	82.8	94.6	45.6	93.4
Quantification (744)	73.9	91.7	51.3	89.6	78.1	90.9	49.6	90.7
Passive (695)	73.1	91.1	60.0	90.3	77.1	90.4	57.7	89.4
Anaphora (700)	73.7	89.7	53.1	86.4	77.0	89.3	53.1	89.4
Conjunction (640)	78.8	91.6	56.2	84.1	79.5	90.8	55.2	90.9
Disjunction (428)	73.5	90.0	49.3	84.5	79.2	89.0	46.7	89.0
Modal (69)	90.4	91.3	60.9	88.3	92.7	94.2	53.6	92.8
Additive particle (13)	77.1	84.6	53.8	79.2	64.9	76.9	30.8	76.9
All	74.6	90.3	53.8	87.9	77.3	89.2	51.9	89.3

Table 11: Results on JSICK for each linguistic tag (%). We evaluated NLI models for accuracy and STS models with the Pearson correlation $\times 100$. N: JSNLI-train, I: JSICK-train, N+I: JSNLI+JSICK-train. Accuracies lower than the overall model accuracy are indicated in red.

Language	NLI task				STS task		
	Prec	Rec	macro-F1	Acc	γ	ρ	MSE
En	87.2	84.9	86.0	86.6	59.9	56.1	0.98
Nl	86.5	83.4	85.3	86.2	57.8	54.6	0.95
Br	85.4	83.1	84.2	85.0	61.3	57.2	0.97
Ja (L-En)	84.1	82.9	83.0	85.2	62.3	60.7	1.08
Ja (L-Ja)	88.2	86.4	87.3	89.2	77.3	77.4	1.11

Table 12: Baseline results from mBERT on different languages in the SICK test set (%). γ : Pearson correlation $\times 100$, ρ : Spearman correlation $\times 100$. Ja (L-En) indicates evaluation results with gold labels from the original SICK test set.

also evaluated mBERT while assuming JSICK gold labels to be the same as those for SICK-EN. Table 12 shows the baseline results for mBERT with different languages in the SICK test set. For both STS and NLI tasks, mBERT performance was relatively higher for Japanese SICK than that for the other datasets.

Table 13 shows confusion matrices for multilingual BERT models on different languages of the SICK NLI test set. Comparing across languages, mBERT performance for contradiction cases was lower in Japanese. Table 14 compares mBERT performance on different languages of the SICK test set for each Japanese linguistic tag. Note that since linguistic phenomena manifest differently by language, Table 14 shows only an approximated comparison of linguistic phenomena. For the STS task, there was little difference among languages, but performance tended to be lower for problems involving additive particles and anaphora. Performance for problems involving additive particles was also low in the NLI task. Performance for problems involving disjunction was low for all languages.

The results of our experiments suggest that multilingual BERT models achieved high performance on SICK across languages. However, the results related to multilingual SICK for each linguistic tag indicate room for improvement regarding the use of multilingual models to capture anaphora, disjunction, and additive particles. Moreover, it remains to be investigated whether pre-trained language models are sensitive to compositional aspects of inference, such as word order and case marking in Japanese. In the next section, we describe the extent to which language models capture word order and case particles, phenomena that are characteristic of Japanese.

5 JSICK Stress Test

5.1 Evaluation Setting

Japanese grammar allows both subject–object–verb and object–subject–verb orders, with the former usually taken as the basic word order and the latter derived by a *scrambling* operation (Hoji, 1985; Saito, 1985). Instead of word order, postpositional case particles function as case markers. For example, the case particles *ga*, *ni*, *o* represent the nominative, dative, and accusative cases, respectively. The JSICK test set contains 1666, 797, and 1006 premise–hypothesis sentence pairs (A, B) whose premise sentences A include basic word orders involving *ga-o* (nominative–accusative), *ga-ni* (nominative–dative), and *ga-de* (nominative–instrumental/locative) relations, respectively. By transforming the syntactic structures of these pairs, we created a JSICK stress-test dataset involving word scrambling and particle-swapping to analyze

	Prediction En				Prediction NI				Prediction Br			
	E	C	N	Rec	E	C	N	Rec	E	C	N	Rec
Gold												
E	1151	2	261	81.4%	1083	5	316	77.1%	1094	5	305	77.9%
C	16	601	103	83.5%	17	592	103	83.1%	23	586	103	82.3%
N	229	51	2513	90.0%	217	67	2506	89.8%	238	63	2489	89.2%
Prec	82.4%	91.9%	87.3%		82.2%	89.2%	85.7%		80.7%	89.6%	85.9%	

	Prediction Ja (L-En)				Prediction Ja (L-Ja)			
	E	C	N	Rec	E	C	N	Rec
Gold								
E	972	1	115	89.3%	922	12	154	84.7%
C	16	574	207	72.0%	11	649	137	81.4%
N	315	73	2654	87.2%	144	73	2825	92.9%
Prec	74.6%	88.6%	89.2%		85.6%	88.4%	90.7%	

Table 13: Confusion matrices for mBERT on different languages in the SICK NLI test set. **Rec** and **Prec** indicate Recall and Precision, respectively.

Phenomenon	NLI task					STS task				
	En	NI	Br	Ja (L-En)	Ja (L-Ja)	En	NI	Br	Ja (L-En)	Ja (L-Ja)
Numeral (1513)	87.9	85.8	85.1	88.2	91.9	66.2	63.3	65.9	65.3	80.0
Negation (1140)	90.3	89.3	88.8	90.7	94.6	61.3	59.0	60.9	58.0	82.8
Quantification (744)	86.0	85.4	83.8	87.9	90.9	62.3	58.7	60.3	58.7	78.1
Passive (695)	87.6	84.4	86.2	87.3	90.4	62.7	58.7	64.2	68.6	77.1
Anaphora (700)	85.3	87.4	85.4	87.0	89.3	59.8	55.7	60.8	60.4	77.0
Conjunction (640)	84.7	85.9	85.8	86.2	90.8	64.7	63.3	65.4	64.9	79.5
Disjunction (428)	85.5	83.9	83.8	87.6	89.0	65.3	59.9	62.9	65.1	79.2
Modal (69)	92.8	95.7	88.4	89.9	94.2	85.5	81.3	85.7	84.3	92.7
Additive particle (13)	61.5	69.2	53.8	53.8	76.9	44.2	24.5	58.7	39.9	64.9
All	86.6	86.2	85.0	85.2	89.2	55.9	57.8	61.3	62.3	77.3

Table 14: Comparison of mBERT performance for different languages in the SICK test set for each linguistic tag (%). Accuracies lower than the overall model accuracy are indicated in red. Ja (L-En) indicates evaluation results with gold labels from the original SICK test set.

A:	小さな	女の子	が	女性	を	見て	いる	
	little	girl	Nom	woman	Acc	looking	is	
	'A little girl is looking at a woman'							(A, B): entailment
A _{ord} :	女性	を	小さな	女の子	が	見て	いる	
	woman	Acc	little	girl	Nom	looking	is	
	'A little girl is looking at a woman'							(A _{ord} , B): entailment
A _{case} :	小さな	女の子	を	女性	が	見て	いる	
	little	girl	Acc	woman	Nom	looking	is	
	'A woman is looking at a little girl'							(A _{case} , B): neutral
A _{del} :	小さな	女の子	女性	見て	いる			
	little	girl	woman	looking	is			
	'A woman is looking at a little girl'							(A _{del} , B): neutral
B:	女の子	が	女性	を	見て	いる		
	girl	Nom	woman	Acc	looking	is		
	'A girl is looking at a woman'							

Table 15: Evaluation settings for the JSICK stress-test dataset.

whether models correctly capture the free-order property of Japanese.

Consider the examples of (A, B) pairs in Table 15, whose gold labels are *entailment*. We first provide a scrambled pair (A_{ord}, B) , where the word order of the premise sentence A is scrambled into *o-ga*, *ni-ga*, or *de-ga* order. Since the meaning of the sentence A_{ord} is the same as that of the original sentence A , the semantic relatedness of the scrambled pair (A_{ord}, B) should be

the same as those of (A, B) . To analyze whether models consider Japanese case particles when predicting entailment labels and similarity scores, we use a rephrased pair (A_{case}, B) , where the only case particles in the premise A are swapped, and a rephrased pair (A_{del}, B) , where the case particles in A are deleted. Since these transformations affect case relations in premise sentences, the meanings of sentences A_{case} and A_{del} should differ from those of the original sentence A . The semantic

Type	Japanese models				Human	Natural
	jaRoBERTa-large	jaRoBERTa-base	jaBERT-large	jaBERT-base (WHOLE)		
Case-scrambling	98.9	97.5	92.4	97.3	93.3	97.3
Part-swapping	99.0	97.0	92.5	98.3	66.7	23.0
Part-deleting	98.3	95.4	92.3	91.7	85.3	6.0

Type	Multilingual models		
	XLM-RoBERTa-large	XLM-RoBERTa-base	mBERT
Case-scrambling	98.4	94.5	98.4
Part-swapping	98.7	94.7	99.4
Part-deleting	98.0	92.0	97.2

Table 16: Comparison between model and human predictions for entailment labels in the JSICK stress set that are the same as those for the original test set (%). **Natural** indicates human-rated results for naturalness (acceptability).

relatedness of the rephrased pairs (A_{case}, B) and (A_{del}, B) should thus also be changed. If a model has generalized word order and case particles, it should consistently predict the same labels for both (A, B) and (A_{ord}, B). Moreover, the model should change the labels for (A_{case}, B) and (A_{del}, B) to *neutral*. We therefore checked the extent to which models changed predictions for (A_{ord}, B), (A_{case}, B), and (A_{del}, B) pairs as compared with those for (A, B).

To rephrase premise sentences, we first parse the sentences using the Japanese constituency parser depccg (Yoshikawa et al., 2017), transform the parse trees using Tsurgeon (Levy and Andrew 2006), and produce their surface strings as rephrased premise sentences. For the JSICK stress set, we evaluated four Japanese pre-trained language models (jaRoBERTa-large, jaRoBERTa-base, jaBERT-large, and jaBERT-base (WHOLE)) and three multilingual models (XLM-RoBERTa-large, XLM-RoBERTa-base, and mBERT) fine-tuned with the JSICK training set. For the NLI task, we also used crowdsourcing to collect human judgments on a subset of the JSICK stress set. We asked the same annotators as those assigning entailment labels for the JSICK dataset to also annotate entailment labels for the JSICK stress set. We also asked the annotators whether each sentence pair was natural, meaning both premise and hypothesis sentences were grammatically correct. Note that we asked them to judge entailment labels even for unnatural sentences. We selected 100 examples for each of three rephrase types and three case particle types, for 900 inference problems in total.

5.2 Results

Table 16 compares percentages of model and human predictions for entailment labels for the

JSICK stress-test dataset that are the same as those for the original JSICK test set. We also show the results for the human naturalness (acceptability) rating task. While humans predicted the same labels for the scrambled examples, they changed their labels for examples where only the case particles were swapped or deleted. Interestingly, humans tend to change more predictions for those examples where only the case particles are swapped than for those where the case particles are deleted, but the acceptability rate for the former was much higher than that for the latter. One reason for this is that Japanese case particles can be dropped under adjacency to become a verb (Saito, 1985), in which case humans can complement the dropped case particles. On the other hand, the models predicted nearly the same labels for A_{ord}, B pairs as for (A, B) pairs. In addition, the predicted labels for A_{case}, B pairs and A_{del}, B remained almost the same as those for the original (A, B) pairs.

To investigate details of the model performance, we confirmed the percentage of predictions for each case particle in the JSICK stress-test dataset that are the same as those predicted for the original test set, as shown in Table 17. In the NLI task, all pre-trained language models predicted nearly the same labels even when the case particle is swapped or deleted, regardless of model type and the kind of case particle. These results indicate that the models predict entailment labels without considering word order and case particles. Similarly, for the STS task, Pearson correlations for A_{ord}, B and A_{case}, B pairs were nearly the same as those for the original (A, B) pairs. Pearson correlations for A_{del}, B pairs were a little lower than those for the original (A, B) pairs, because the deletion of case particles in A_{del} decreases word overlap between A_{del} and B .

Order	Model	Type	NLI task				STS task
			Yes	No	Unk	All	
ga-de	jaRoBERTa-large	Case-scrambling	98.8	100.0	99.5	99.4	98.8
		Part-swapping	97.2	99.1	99.8	99.1	97.0
		Part-deleting	95.2	99.1	99.1	98.1	92.6
	jaRoBERTa-base	Case-scrambling	98.4	96.6	98.0	97.9	95.0
		Part-swapping	97.6	97.4	96.9	97.1	99.1
		Part-deleting	95.6	95.7	95.0	95.2	91.3
	jaBERT-large	Case-scrambling	84.3	87.2	94.2	90.9	98.6
		Part-swapping	85.1	88.0	94.1	91.1	95.8
		Part-deleting	83.5	86.3	94.7	90.9	94.1
	jaBERT-base (WHOLE)	Case-scrambling	95.5	100.0	97.2	97.1	99.2
		Part-swapping	98.8	100.0	99.4	99.3	96.8
		Part-deleting	91.5	96.6	95.2	94.4	92.4
	XLM-RoBERTa-large	Case-scrambling	98.8	98.3	99.2	99.0	89.4
		Part-swapping	100.0	98.3	99.2	99.3	94.2
		Part-deleting	97.2	97.4	98.4	98.0	86.8
	XLM-RoBERTa-base	Case-scrambling	89.1	93.2	96.9	94.5	97.7
		Part-swapping	85.5	90.6	97.3	93.6	96.9
		Part-deleting	79.4	88.9	95.5	90.7	88.1
mBERT	Case-scrambling	97.2	97.4	98.6	98.1	98.0	
	Part-swapping	100.0	99.1	100.0	99.9	98.3	
	Part-deleting	94.8	99.1	98.4	97.6	96.4	
ga-ni	jaRoBERTa-large	Case-scrambling	97.2	96.0	99.2	98.4	98.6
		Part-swapping	97.7	96.0	99.2	98.5	97.0
		Part-deleting	96.0	96.0	98.7	97.7	93.6
	jaRoBERTa-base	Case-scrambling	98.3	97.0	97.5	97.6	96.7
		Part-swapping	98.9	97.0	97.1	97.5	94.6
		Part-deleting	97.2	95.0	94.9	95.4	91.3
	jaBERT-large	Case-scrambling	90.3	87.0	93.7	92.1	97.7
		Part-swapping	91.5	89.0	93.8	92.7	95.2
		Part-deleting	91.5	86.0	93.5	92.1	94.1
	jaBERT-base (WHOLE)	Case-scrambling	97.2	97.0	96.4	96.6	98.2
		Part-swapping	99.4	99.0	98.5	98.7	96.0
		Part-deleting	93.2	94.0	92.3	92.7	92.0
	XLM-RoBERTa-large	Case-scrambling	99.4	96.0	99.2	98.9	84.8
		Part-swapping	98.3	95.0	99.0	98.4	93.5
		Part-deleting	98.3	95.0	99.0	98.4	86.9
	XLM-RoBERTa-base	Case-scrambling	94.9	82.0	96.0	94.0	93.0
		Part-swapping	92.0	84.0	97.9	94.8	91.9
		Part-deleting	91.5	81.0	94.6	92.2	88.3
mBERT	Case-scrambling	98.9	94.0	99.6	98.7	95.7	
	Part-swapping	98.9	97.0	97.0	99.6	98.4	
	Part-deleting	97.7	94.0	98.8	98.0	96.0	
ga-o	jaRoBERTa-large	Case-scrambling	98.3	97.7	99.5	99.0	98.5
		Part-swapping	98.6	98.9	99.6	99.3	95.6
		Part-deleting	98.3	98.1	99.1	98.8	92.6
	jaRoBERTa-base	Case-scrambling	96.3	96.6	97.8	97.3	91.4
		Part-swapping	95.8	95.5	97.3	96.7	97.2
		Part-deleting	95.5	95.1	95.5	95.4	95.9
	jaBERT-large	Case-scrambling	89.3	92.4	95.2	93.5	96.9
		Part-swapping	88.4	92.8	94.9	93.2	93.5
		Part-deleting	89.8	92.0	94.8	93.3	90.6
	jaBERT-base (WHOLE)	Case-scrambling	96.1	98.1	96.1	96.4	97.8
		Part-swapping	98.3	98.9	99.1	98.9	94.8
		Part-deleting	92.4	93.6	92.5	92.6	90.0
	XLM-RoBERTa-large	Case-scrambling	97.2	96.6	98.2	97.7	87.5
		Part-swapping	98.3	97.7	98.6	98.4	94.5
		Part-deleting	96.6	96.2	98.7	97.8	86.8
	XLM-RoBERTa-base	Case-scrambling	90.4	90.5	97.3	94.8	96.5
		Part-swapping	92.4	92.8	96.8	95.3	94.9
		Part-deleting	86.4	89.0	95.5	92.6	85.1
mBERT	Case-scrambling	99.2	97.7	98.3	98.4	95.5	
	Part-swapping	99.2	98.5	99.4	99.2	97.9	
	Part-deleting	95.5	94.7	97.5	96.6	94.1	

Table 17: Percentage of predictions for the JSICK stress-test dataset that are the same as those for the original test set for each case particle (%). Yes, No, and Unk indicate accuracies on *entailment*, *contradiction*, and *neutral* examples, respectively. For the STS task, we calculated the Pearson correlation between predictions for the original pairs and those for the rephrased pairs.

Type	jaRoBERTa-l	jaRoBERTa-l+aug
Case-scrambling	98.9	97.7
Part-swapping	99.0	69.2
Part-deleting	98.3	69.1

Table 18: Percentages of model predictions that are the same as those for the original JSICK when a subset of rephrased examples is added to the training set to learn word order and case particles. **jaRoBERTa-l+aug** shows the result with data augmentation.

Augmenting Training Data for Sensitivity to Case Particles To analyze whether data augmentation improves the model behavior for word order and case particles, we rephrased a small subset of the training set in three ways, creating (i) data where the NP argument is scrambled but its gold label is the same as the original, (ii) data where the case particle is deleted, and its gold label is set randomly, and (iii) data where only the case particle is scrambled and its gold label is set randomly.

We added 300 examples of each data type. These additional training data play a role in exposing models to three cues: (i) the order of NP arguments does not change entailment labels for sentence pairs, (ii) the existence of a case particle, and (iii) that its position can change entailment labels. Table 18 shows the percentage of predictions by the jaRoBERTa-large NLI model that are the same as those for the original JSICK when a subset of rephrased examples is added to the training set. As that table shows, data augmentation changed model predictions on examples where the case particle is swapped or deleted. This indicates that although the NLI model does not implicitly learn case particles during pre-training and fine-tuning, a small amount of data augmentation to learn word order and case particles can improve the model sensitivity to case particles.

6 Conclusion

We introduced JSICK, a Japanese standard NLI/STS dataset, by manually translating English SICK into Japanese and re-annotating its gold labels. In baseline experiments, we compared the performance of various pre-trained Japanese language models on JSICK. While the Japanese RoBERTa-large model achieved state-of-the-art performance, the performance of multilingual pre-trained language models achieved comparable

results. Experiments with multilingual models on SICK datasets in different languages, including JSICK, showed that the performance of multilingual models was relatively low on inference problems involving anaphora, disjunction, and additive particles.

Furthermore, to investigate the extent to which Japanese and multilingual pre-trained language models are sensitive to word order and case particles, we provided a JSICK stress-test dataset involving word scrambling and particle-swapping from JSICK. The results from that dataset suggest that both Japanese and multilingual models do not consider word order and case particles when making predictions for Japanese NLI/STS tasks. These are novel findings that are not obtainable from other datasets, including SICK in English and other languages. Overall, the results suggest large room for improvement of both Japanese and multilingual pre-trained language models regarding their sensitivity to flexible word order and the representations of case particles. Further improvements might be obtained by more adequate representations of Japanese vocabularies in multilingual pre-trained language models (Chung et al., 2020; Rust et al., 2021). We believe our dataset will be useful in future research for realizing more advanced models capable of appropriately performing multilingual compositional inference.

Acknowledgments

We thank the anonymous reviewers and the Action Editor for helpful comments and suggestions that improved this paper. We also thank Daisuke Kawahara and Tomohide Shibata for helpful advice on the experimental settings of Japanese RoBERTa and BERT models. This work was supported by JSPS KAKENHI grant number JP20K19868, JST, PRESTO grant number JPMJPR21C8, and JST, CREST grant number JPMJCR2114, Japan.

References

Eneko Agirre, Carmen Banea, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, Rada Mihalcea, German Rigau, and Janyce Wiebe. 2016. SemEval-2016 task 1: Semantic textual similarity, monolingual and cross-lingual

- evaluation. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 497–511. <https://doi.org/10.18653/v1/S16-1081>
- Hossein Amirkhani, Mohammad Azari Jafari, Azadeh Amirak, Zohreh Pourjafari, Soroush Faridan Jahromi, and Zeinab Kouhkan. 2020. FarsTail: A Persian natural language inference dataset. *CoRR*, cs.CL/2009.08820. Version 1.
- Masayuki Asahara and Yuji Matsumoto. 2003. ipadic version 2.7.0 User’s Manual. Nara Institute of Science and Technology.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642. <https://doi.org/10.18653/v1/D15-1075>
- Daniel Cer, Mona Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. 2017. SemEval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14.
- Hyung Won Chung, Dan Garrette, Kiat Chuan Tan, and Jason Riesa. 2020. Improving multilingual models with language-clustered vocabularies. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4536–4546. <https://doi.org/10.18653/v1/2020.emnlp-main.367>
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451. <https://doi.org/10.18653/v1/2020.acl-main.747>
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. XNLI: Evaluating cross-lingual sentence representations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485. <https://doi.org/10.18653/v1/D18-1269>
- Robin Cooper, Richard Crouch, Jan van Eijck, Chris Fox, Josef van Genabith, Jan Jaspers, Hans Kamp, Manfred Pinkal, Massimo Poesio, Stephen Pulman. 1994. FraCaS—a framework for computational semantics. *Deliverable*, D6.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2006. The pascal recognising textual entailment challenge. In *Machine Learning Challenges. Evaluating Predictive Uncertainty, Visual Object Classification, and Recognising Tectual Entailment*, pages 177–190. https://doi.org/10.1007/11736790_9
- 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.
- Gottlob Frege. 1963. Compound thoughts. *Mind*, 72(285):1–17. <https://doi.org/10.1093/mind/LXXII.285.1>
- William Gantt, Benjamin Kane, and Aaron Steven White. 2020. Natural language inference with mixed effects. In *Proceedings of the Ninth Joint Conference on Lexical and Computational Semantics*, pages 81–87.
- Max Glockner, Vered Shwartz, and Yoav Goldberg. 2018. Breaking NLI systems with sentences that require simple lexical inferences. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 650–655. <https://doi.org/10.18653/v1/P18-2103>
- Emily Goodwin, Koustuv Sinha, and Timothy J. O’Donnell. 2020. Probing linguistic systematicity. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1958–1969. <https://doi.org/10.18653/v1/2020.acl-main.177>
- Ashim Gupta, Giorgi Kvernadze, and Vivek Srikumar. 2021. BERT & family eat word salad: Experiments with text understanding. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021*, pages 12946–12954.

- <https://doi.org/10.1609/aaai.v35i14.17531>
- Jiyeon Ham, Yo Joong Choe, Kyubyong Park, Ilji Choi, and Hyungjoon Soh. 2020. KorNLI and KorSTS: New benchmark datasets for Korean natural language understanding. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 422–430. <https://doi.org/10.18653/v1/2020.findings-emnlp.39>
- John Hawkins. 1978. *Definiteness and Indefiniteness. A Study in Reference and Grammaticality Prediction*, Routledge.
- Yuta Hayashibe. 2020. Japanese realistic textual entailment corpus. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 6827–6834.
- Irene Heim. 1982. The Semantics of Definite and Indefinite Noun Phrases. Ph.D. thesis, UMass Amherst.
- Jack Hessel and Alexandra Schofield. 2021. How effective is BERT without word ordering? implications for language understanding and data privacy. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 204–211. <https://doi.org/10.18653/v1/2021.acl-short.27>
- John Hinds. 1986. *Japanese: Descriptive Grammar*. Croom Helm.
- Hajime Hoji. 1985. *Logical Form Constraints and Configurational Structures in Japanese*. Ph.D. Thesis. University of Washington.
- Hai Hu, Kyle Richardson, Liang Xu, Lu Li, Sandra Kübler, and Lawrence Moss. 2020. OCNLI: Original Chinese Natural Language Inference. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3512–3526. <https://doi.org/10.18653/v1/2020.findings-emnlp.314>
- Theo Janssen and Barbara Partee. 1997. Compositionality. Johan van Benthem and Alice ter Meulen, editors, *Handbook of Logic and Language*, pages 417–473. Elsevier. <https://doi.org/10.1016/B978-044481714-3/50011-4>
- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. The state and fate of linguistic diversity and inclusion in the NLP world. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293. <https://doi.org/10.18653/v1/2020.acl-main.560>
- Aikaterini-Lida Kalouli, Livy Real, and Valeria de Paiva. 2017. Textual inference: Getting logic from humans. In *Proceedings of the 12th International Conference on Computational Semantics (IWCS) — Short papers*.
- Jerrold Katz and Jerry Fodor. 1963. The structure of a semantic theory. *Language*, 39(2):170–210. <https://doi.org/10.2307/411200>
- Ai Kawazoe, Ribeka Tanaka, Koji Mineshima, and Daisuke Bekki. 2017. An inference problem set for evaluating semantic theories and semantic processing systems for Japanese. In *New Frontiers in Artificial Intelligence*, pages 58–65. https://doi.org/10.1007/978-3-319-50953-2_5
- Tatsuki Kuribayashi, Yohei Oseki, Takumi Ito, Ryo Yoshida, Masayuki Asahara, and Kentaro Inui. 2021. Lower perplexity is not always human-like. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5203–5217. <https://doi.org/10.18653/v1/2021.acl-long.405>
- Hang Le, Loïc Vial, Jibril Frej, Vincent Segonne, Maximin Coavoux, Benjamin Lecouteux, Alexandre Allauzen, Benoit Crabbé, Laurent Besacier, and Didier Schwab. 2020. FlauBERT: Unsupervised language model pre-training for French. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 2479–2490.
- Roger Levy and Galen Andrew. 2006. Tregex and tsurgeon: Tools for querying and manipulating tree data structures. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC'06)*.
- Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fenfei Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, Xiaodong

- Fan, Ruofei Zhang, Rahul Agrawal, Edward Cui, Sining Wei, Taroon Bharti, Ying Qiao, Jiun-Hung Chen, Winnie Wu, Shuguang Liu, Fan Yang, Daniel Campos, Rangan Majumder, and Ming Zhou. 2020. XGLUE: A new benchmark dataset for cross-lingual pre-training, understanding and generation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6008–6018. <https://doi.org/10.18653/v1/2020.emnlp-main.484>
- Tal Linzen. 2020. How can we accelerate progress towards human-like linguistic generalization? In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5210–5217. <https://doi.org/10.18653/v1/2020.acl-main.465>
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized bert pretraining approach. *CoRR*, cs.CL/1907.11692. Version 1.
- Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, and Roberto Zamparelli. 2014. A SICK cure for the evaluation of compositional distributional semantic models. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 216–223.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448. <https://doi.org/10.18653/v1/P19-1334>
- Richard Montague. 1973. The proper treatment of quantification in ordinary English, K. J. J. Hintikka, J. Moravcsic, and P. Suppes, editors, *Approaches to Natural Language*, pages 221–242. Reidel, Dordrecht. https://doi.org/10.1007/978-94-010-2506-5_10
- Hajime Morita, Daisuke Kawahara, and Sadao Kurohashi. 2015. Morphological analysis for unsegmented languages using recurrent neural network language model. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2292–2297. <https://doi.org/10.18653/v1/D15-1276>
- Aakanksha Naik, Abhilasha Ravichander, Norman Sadeh, Carolyn Rose, and Graham Neubig. 2018. Stress test evaluation for natural language inference. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2340–2353.
- Kimiko Nakanishi and Satoshi Tomioka. 2004. Japanese plurals are exceptional. *Journal of East Asian Linguistics*, 13(2):113–140. <https://doi.org/10.1023/B:JEAL.0000019058.46668.c1>
- Sungjoon Park, Jihyung Moon, Sungdong Kim, Won Ik Cho, Jiyeon Han, Jangwon Park, Chisung Song, Junseong Kim, Yongsook Song, Taehwan Oh, Joohong Lee, Juhyun Oh, Sungwon Lyu, Younghoon Jeong, Inkwon Lee, Sangwoo Seo, Dongjun Lee, Hyunwoo Kim, Myeonghwa Lee, Seongbo Jang, Seungwon Do, Sunkyoung Kim, Kyungtae Lim, Jongwon Lee, Kyumin Park, Jamin Shin, Seonghyun Kim, Lucy Park, Alice Oh, Jung-Woo Ha, and Kyunghyun Cho. 2021. KLUE: Korean language understanding evaluation. *CoRR*, cs.CL/2105.09680. Version 1.
- Ellie Pavlick and Tom Kwiatkowski. 2019. Inherent disagreements in human textual inferences. *Transactions of the Association for Computational Linguistics*, 7:677–694. https://doi.org/10.1162/tacl_a_00293
- Thang Pham, Trung Bui, Long Mai, and Anh Nguyen. 2021. Out of order: How important is the sequential order of words in a sentence in natural language understanding tasks? In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1145–1160. <https://doi.org/10.18653/v1/2021.findings-acl.98>
- Shauli Ravfogel, Yoav Goldberg, and Tal Linzen. 2019. Studying the inductive biases of RNNs with synthetic variations of natural languages. 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 3532–3542. <https://doi.org/10.18653/v1/N19-1356>
- Livy Real, Ana Rodrigues, Andressa Vieira e Silva, Beatriz Albiero, Bruna Thalenberg,

- Bruno Guide, Cindy Silva, G. Lima, Igor C. S. Câmara, Miloš Stanojević, Rodrigo Souza, and Valeria de Paiva. 2018. SICK-BR: A portuguese corpus for inference. In *Computational Processing of the Portuguese Language*, pages 303–312. https://doi.org/10.1007/978-3-319-99722-3_31
- Ohad Rozen, Vered Shwartz, Roei Aharoni, and Ido Dagan. 2019. Diversify your datasets: Analyzing generalization via controlled variance in adversarial datasets. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 196–205. <https://doi.org/10.18653/v1/K19-1019>
- Phillip Rust, Jonas Pfeiffer, Ivan Vulić, Sebastian Ruder, and Iryna Gurevych. 2021. How good is your tokenizer? On the monolingual performance of multilingual language models. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3118–3135. <https://doi.org/10.18653/v1/2021.acl-long.243>
- Mamoru Saito. 1985. *Some Asymmetries in Japanese and their Theoretical Implications*. Ph.D. thesis, NA Cambridge.
- Mike Schuster and Kaisuke Nakajima. 2012. Japanese and Korean voice search. In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5149–5152. <https://doi.org/10.1109/ICASSP.2012.6289079>
- Haitham Seelawi, Ibraheem Tuffaha, Mahmoud Gzawi, Wael Farhan, Bashar Talafha, Riham Badawi, Zyad Sober, Oday Al-Dweik, Abed Alhakim Freihat, and Hussein Al-Natsheh. 2021. ALUE: Arabic language understanding evaluation. In *Proceedings of the Sixth Arabic Natural Language Processing Workshop*, pages 173–184.
- Tatiana Shavrina, Alena Fenogenova, Emelyanov Anton, Denis Shevelev, Ekaterina Artemova, Valentin Malykh, Vladislav Mikhailov, Maria Tikhonova, Andrey Chertok, and Andrey Evlampiev. 2020. RussianSuperGLUE: A Russian language understanding evaluation benchmark. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4717–4726. <https://doi.org/10.18653/v1/2020.emnlp-main.381>
- Masayoshi Shibatani. 1990. *The Languages of Japan*, Cambridge University Press.
- Koustuv Sinha, Robin Jia, Dieuwke Hupkes, Joelle Pineau, Adina Williams, and Douwe Kiela. 2021a. Masked language modeling and the distributional hypothesis: Order word matters pre-training for little. *CoRR*, cs.CL/2104.06644. Version 1. <https://doi.org/10.18653/v1/2021.emnlp-main.230>
- Koustuv Sinha, Prasanna Parthasarathi, Joelle Pineau, and Adina Williams. 2021b. UnNatural Language Inference. In *Proceedings of the Joint Conference of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL-IJCNLP2021)*. <https://doi.org/10.18653/v1/2021.acl-long.569>
- Arseny Tolmachev, Daisuke Kawahara, and Sadao Kurohashi. 2018. Juman++: A morphological analysis toolkit for scriptio continua. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 54–59. <https://doi.org/10.18653/v1/D18-2010>
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2019. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the International Conference on Learning Representations (ICLR)*. <https://doi.org/10.18653/v1/W18-5446>
- Jennifer C. White and Ryan Cotterell. 2021. Examining the inductive bias of neural language models with artificial languages. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 454–463.
- Gijs Wijnholds and Michael Moortgat. 2021. SICK-NL: A dataset for Dutch natural language inference. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*,

- pages 1474–1479. <https://doi.org/10.18653/v1/2021.eacl-main.126>
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. 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 1112–1122. <https://doi.org/10.18653/v1/N18-1101>
- Liang Xu, Hai Hu, Xuanwei Zhang, Lu Li, Chenjie Cao, Yudong Li, Yechen Xu, Kai Sun, Dian Yu, Cong Yu, Yin Tian, Qianqian Dong, Weitang Liu, Bo Shi, Yiming Cui, Junyi Li, Jun Zeng, Rongzhao Wang, Weijian Xie, Yanting Li, Yina Patterson, Zuoyu Tian, Yiwen Zhang, He Zhou, Shaowei Hua Liu, Zhe Zhao, Qipeng Zhao, Cong Yue, Xinrui Zhang, Zhengliang Yang, Kyle Richardson, and Zhenzhong Lan. 2020. CLUE: A Chinese language understanding evaluation benchmark. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4762–4772.
- Hitomi Yanaka, Koji Mineshima, and Kentaro Inui. 2021. Exploring transitivity in neural NLI models through veridicality. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 920–934. <https://doi.org/10.18653/v1/2021.eacl-main.78>
- Yinfei Yang, Yuan Zhang, Chris Tar, and Jason Baldridge. 2019. PAWS-X: A cross-lingual adversarial dataset for paraphrase identification. 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 3687–3692. <https://doi.org/10.18653/v1/D19-1382>
- Masashi Yoshikawa, Hiroshi Noji, and Yuji Matsumoto. 2017. A* CCG parsing with a supertag and dependency factored model. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 277–287. <https://doi.org/10.18653/v1/P17-1026>
- Takumi Yoshikoshi, Daisuke Kawahara, and Sadao Kurohashi. 2020. Multilingualization of natural language inference datasets using machine translation (in Japanese). In *Proceedings of the 244th Meeting of Natural Language Processing*.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. Bertscore: Evaluating text generation with bert. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Yuan Zhang, Jason Baldridge, and Luheng He. 2019. PAWS: Paraphrase adversaries from word scrambling. 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 1298–1308.