

Probing for Multilingual Numerical Understanding in Transformer-Based Language Models

Devin Johnson, Denise Mak, Drew Barker, Lexi Loessberg-Zahl

Department of Linguistics

University of Washington

{dj1121, dpm3, barkand, lexilz}@uw.edu

Abstract

Natural language numbers are an example of compositional structures, where larger numbers are composed of operations on smaller numbers. Given that compositional reasoning is a key to natural language understanding, we propose novel multilingual probing tasks tested on DistilBERT, XLM, and BERT to investigate for evidence of compositional reasoning over numerical data in various natural language number systems. By using both grammaticality judgment and value comparison classification tasks in English, Japanese, Danish, and French, we find evidence that the information encoded in these pretrained models' embeddings is sufficient for grammaticality judgments but generally not for value comparisons. We analyze possible reasons for this and discuss how our tasks could be extended in further studies.

1 Introduction

In recent years, transformer-based language models such as BERT (Devlin et al., 2018), XLM (Lample and Conneau, 2019), and DistilBERT (Sanh et al., 2020) have achieved unprecedented results on a wide array of natural language understanding tasks, even when such models are trained on other tasks (i.e. transfer learning). In light of this success, there has been increased interest in investigating what particular information transformer-based language models encode in their word embeddings during pretraining that allows them to perform well in transfer learning experiments. Put into the context of this paper, we may ask: do such models gain certain linguistic/compositional understanding from pretraining? Attempts to assess such phenomena are commonly referred to as probing experiments. In this paper, we introduce a novel probing task using targeted datasets and classification tasks aimed at evaluating models' compositional reasoning capabilities

in relation to numbers¹ in multiple natural languages.

We choose both grammaticality judgment and value comparison tasks (Section 3) to assess multilingual DistilBERT, XLM, and BERT² over various number systems. We argue that high performance on these tasks indicates some ability of reasoning over compositional structures, particularly over the rules generating valid compositional structures (task 1) and the resultant meanings of the structures (task 2). After probing the selected models on our tasks, we discuss explanations for the performance of models, as well as possible future extensions to this probing task schema. Additionally, studies in cognitive psychology such as Miller et al. (1995) assert that children learning more transparent number systems (i.e. those exhibiting more regularity in their surface forms such as Japanese) have a greater counting proficiency in several tasks compared to those learning less transparent (opaque) systems, such as English or French. Although it is not the focus of our work, given the multilingual setting, we will refer to the idea of number system transparency when analyzing possible explanations of results.³

2 Related Work

Our approach is informed by previous linguistically-motivated probing studies such

¹For our purposes, numbers are spelled out, i.e. written out as words such as “ninety”.

²Models were chosen for their varied sizes (num. parameters) as well as our access to computing resources.

³Number system complexity could be the subject of its own paper. However, as a small example, we can look at “thirty” in English and “三十” in Japanese. In English, there is no previous number such as “three” that appears (unchanged) in the word “thirty”. In Japanese, however, the word consists of the kanji for 3 (“三”) and the kanji for 10 (“十”). If we continue comparing in this way, we would see compositionality more clearly and regularly in Japanese's surface forms, thus forming our intuitions.

at those discussed in [Belinkov and Glass \(2019\)](#) and [Ettinger \(2020\)](#). Though [Ettinger \(2020\)](#) discusses important findings on psycholinguistic probing experiments of BERT, we find [Ettinger et al. \(2016\)](#) particularly useful for our study due to its clear explanations of linguistic probing experiment setup. In their study, Ettinger et al. present methods for constructing linguistically-targeted datasets (including example sentences, as we use) and classification tasks for probing word embeddings for semantic knowledge. As one of our tasks also seeks to probe for semantic knowledge, we were able to use this setup as a rough guideline. In addition, the authors are careful to create linguistic data with sufficient diversity as not to give potentially helpful cues to their classifier which are not related to the knowledge they wish to probe for. We thus carefully create our task data in a similar manner by limiting the distribution of our data (described more later) and forbidding duplicates.

To our knowledge, there have been few studies conducted on investigating numerical understanding specifically in transformer-based language models with a multilingual approach. However, one particularly relevant study on English comes from [Wallace et al. \(2019\)](#). Wallace et al. probe the embeddings of various models (BERT, ELMo, word2vec, etc.) using three tasks: find the maximum of a list, decode a number word to its numerical form, and add number words to produce a numerical form. In their results, the authors note that all embeddings contain some understanding, though standard embeddings perform particularly well and character-level models perform best. Although we investigate similar phenomena as Wallace et al., our methodology includes several key differences:

- Our focus is first and foremost to present a novel probing task schema/data and test it on a selection of transformer-based models - not to compare performance of differing language model architectures on previously-made tasks.
- We seek to draw conclusions from our task performance about model weaknesses over varied languages and suggest ways in which further probing experiments in this area can be designed in the future.
- We assert that the inclusion of other lan-

guages besides English is an important addition to probing experiments, as variation in language structure may help point to previously-unseen weaknesses in pretrained models.

- We include spelled-out numbers above 100 (up to 1000), which were not used in Wallace et al. We believe having a larger range of numbers might highlight weaknesses of models in handling multiple identical tokens in one word.
- We use only spelled-out number words, and do not include tasks where both Arabic numerals and spelled-out words might be used. Our reasoning for this choice is our desire to leave out the possibility of models merely learning a mapping from number words to numerals in order to perform well on tasks. In this way, we hope to make our tasks/data as restrictive as possible in order that they *require* a certain compositional/linguistic understanding.

3 Methods

We propose and perform two classification tasks in English, Danish, Japanese, and French. Task 1 is a probe for underlying syntactic information encoded in pretrained word embeddings, while task 2 is a probe of underlying semantic information. We run our tasks on all three models over two different datasets which we have generated: one where number words are inserted into sentences (e.g. “There are seven hundred books in the library.”) and one with numbers alone (e.g. “seven-hundred”). Our probing model features a multilayer perceptron (a NN with a single hidden layer) classifier on top of the existing transformer language model architecture. In this manner, pretrained word embeddings from the language model (BERT, DistilBERT, XLM) are fed as input to the MLP classifier which itself is then trained on our tasks. A depiction of the probing model structure is shown in Figure 1.

3.1 Task 1: Grammaticality Judgment

We specify the first task as follows:

- Let $v \in \{bare, sentence\}$ specify the variant of our task. If $v = bare$, then only training examples with numbers *not* inserted

into sentences will be used for grammaticality judgments. Otherwise, only training examples with number inserted into sentences are used. A mixture of two input data types is never used.

- Let the training set of task 1, T , be defined by pairs $t_0 \dots t_n$ where $t_i = (x_{t_i}, y_{t_i})$
- Let x_{t_i} be the input of the i 'th training example t_i such that $x_{t_i} = s$. s is a string consisting of a number word such as “thirty-two” or a sentence containing a number word such as “He could eat thirty-two oranges” (depending on the value of v)
- Let $y_{t_i} \in \{0, 1\}$ be the corresponding label of input x_{t_i} of training sample t_i . $y_{t_i} = 1$ if the input string s of x_{t_i} is ungrammatical, otherwise $y_{t_i} = 0$.

We argue high accuracy on this task is evidence of some understanding of the underlying compositional/syntactic rules for the process which generates natural language numbers.

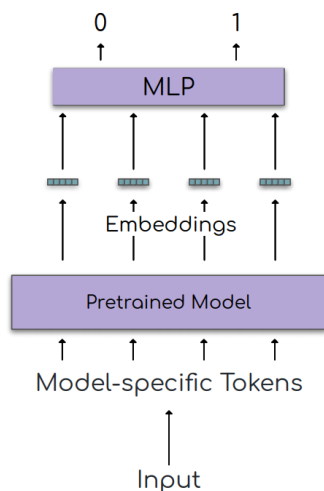


Figure 1: Probing model; Inputs can vary from sentences with numbers to standalone numbers. All weights are fixed except the connections to the MLP layer

As an example, “two hundred three and fifty” is not a number in the English language, while “five hundred” is. We also assert that the importance of using spelled-out numbers comes into play since no string of Arabic numerals is ungrammatical except a small amount of strings such as

“0112”. Often in written text, Arabic numerals are used in place of number words; However, given human ability to generalize compositional rules from other structures in order to learn new structures, even if it were the case that fewer number words were seen in pretraining, we would hope to see a similar compositional generalization capability present in pretrained language models’ embeddings which would allow them to perform well on this task.

3.2 Task 2: Value Comparison

We specify the second task as follows:

- Let $v \in \{bare, sentence\}$ specify the variant of our task. If $v = bare$, then only training examples with numbers *not* inserted into sentences will be used for value comparison. Otherwise, only training examples with number inserted into sentences are used. A mixture of two input data types is never used.
- Let the training set of task 2, U , be defined by pairs $u_0 \dots u_n$ where $u_i = (x_{u_i}, y_{u_i})$
- Let x_{u_i} be the input of the i 'th training example u_i such that $x_{u_i} = (s_0, s_1)$. s_0 and s_1 represent bare number words or number words inserted into sentences (depending on the value of v).
- Let $y_{u_i} \in \{0, 1\}$ be the corresponding label of input x_{u_i} of training sample u_i . $y_{u_i} = 0$ if for s_0 and s_1 of x_{u_i} , s_0 refers to a value larger than that of s_1 . $y_{u_i} = 1$ if for s_0 and s_1 of x_{u_i} , s_0 refers to a value smaller than that of s_1 .

With this task, we take high accuracy as evidence of some understanding of the compositional semantic information carried by the number, i.e. its magnitude. For example, given the pair ($s_0 =$ “twelve”, $s_1 =$ “fifteen”) the correct output should be 1, since the first number in the pair is less than the second. This task is similar in form to the list maximum task of Wallace et al. (2019); However, notable differences include our usage of number words in sentences, our inclusion of number words above 100, and languages other than English.

Task (v)	Input	Output
Task 1 (sentence)	“He could eat three hundred and two oranges”	0 (s grammatical)
Task 1 (bare)	“seventy-four six hundred and thirty-eight”	1 (s ungrammatical)
Task 2 (sentence)	“There are seven hundred and eighty-six books in the library” “There are thirty-eight books in the library”	0 (s_0 greater)
Task 2 (bare)	“five hundred” “six hundred”	1 (s_0 less)

Table 1: Sample of generated English data for our tasks

4 Data

Separate datasets for each variant of each task were made, each with a 60-20-20 train-validation-test split. To generate number words to create training example inputs, the python package `num2words` (Ogawa) was used for text conversion from numerical to standard spelled-out numbers. For task 1, it was necessary to create ungrammatical numbers in each language. These ungrammatical number words were created by randomly appending grammatical number words (or parts of them) together and controlling for length. Lastly, it was also necessary to create custom sentences to insert our number words into. Eleven sentence templates were used, translated into each of our languages and verified by native speakers. A sample of our data can be seen in Table 1. Detailed information on dataset statistics and data generation techniques can be found in appendix B.

5 Results

The following sections show probing results on both tasks. Before probing and as a precaution, we fine-tuned our models on both tasks. As we expected, we find that both tasks are learnable to accuracies above 95%. A further discussion of fine-tuning results is left for appendix A.

5.1 Task 1 Results

Our results on task 1 (Figure 2) show that the pre-trained embeddings of multilingual BERT, DistilBERT and XLM seem to have sufficient information to be able to determine grammaticality of number words in Japanese, English, Danish, and French at better-than-chance performance. We thus argue that these results suggest that the pre-trained embeddings of these models contain some understanding of the compositional rules for generating number words in the languages we’ve selected.

There are a few patterns worth noting in the results. Firstly, accuracy on bare numbers was always better than accuracy on numbers in sentences. Though accuracy on numbers in sentences was not extremely poor, our initial prediction was that they would perform better as they would resemble the type of data which the models were pretrained on (Wikipedia). Secondly, in terms of overall model performance, DistilBERT performs best (sometimes even at 100% accuracy) in the majority of cases, followed by XLM and BERT. A further analysis of these patterns is left for our discussion (6).

5.2 Task 2 Results

On our second task, we find that overall, the information in the pretrained embeddings of BERT, DistilBERT, and XLM is mostly insufficient for comparing number magnitudes in our tested languages with high accuracy. This suggests that these pretrained embeddings may struggle with understanding the compositional semantics of number words (i.e. how compositional elements in number words form to create meaning).

As for patterns in these results, we can again see that bare number performance is always equal to or better than when numbers are inserted into sentences. Looking at bare results, DistilBERT again performs well, but this time XLM performs best on Japanese and English. A further analysis of these patterns is also left for our discussion (6).




Figure 2: Task test accuracy on probing per language/model.

6 Discussion

6.1 Pretraining Data

Since this probing experiment is using pretrained transformer-based language models (Wolf et al., 2019), we can briefly discuss pretraining methods in order to ascertain potential effects on results. Each model was pretrained using masked language modelling and all except XLM (unspecified) were pretrained on multilingual Wikipedia data. With this, one may ask whether this method of training should be expected to encode the information we probe for. Particularly, one may point to the usage of Arabic numerals in Wikipedia text and less-frequent usage of spelled-out numbers as a cause for poor performance on task 2. In imagining alternatives to Wikipedia, a consideration of Arabic numerals must still be present, which serves as a reminder that there is no *perfect* pretraining set.

However, we do not believe this would prevent language models from capturing the information necessary to complete our task 2. We assert that even if these models have seen fewer spelled-out numbers, they could still learn compositional rules from other linguistic structures and generalize to our probing task. If Arabic numerals proved to be an issue, we would expect to see our results be worse across the board, not only on task 2. Thus, from our results, it seems that pretraining on Wikipedia was certainly not sufficient for encoding a highly accurate sense of number magnitude for any of the models/languages, but this was likely not due to pretraining methods.

6.2 Worse Performance on Task 2

Why might models have more difficulty in ascertaining magnitude of number words? For one, we believe this task is naturally more difficult than the first because of the deep semantic information necessary to succeed on it. In the first task, it may have been possible to leverage at least *some* of the surface level characteristics of grammatical and ungrammatical words, whereas in the second task there is no such leveraging possible. That is to say, in the first task, a model can learn syntactic information more directly from surface level patterns. Instead, in task 2, the models need to have encoded some semantic information about the magnitude of number words, where the surface forms of these words gives less indication of their underlying meaning (except for the possibility of longer words having larger quantities, though this is not always the case). Given this is true, this may point to a weakness in models to make fine-grained semantic distinctions regarding quantities, especially when quantities are used in sentences and not left bare.

6.3 Language Transparency

In terms of number system transparency (as mentioned in our introduction), we loosely presumed that accuracies might follow the order of Japanese > English/Danish > French, with Japanese performing best given its higher transparency and French the worst due to its vigesimal number system. Again, we choose not to formally define transparency, as such a formal definition is an in-depth topic of its own. To our slight surprise, our

results did not match these predictions, with rankings of performance by language varying by task and task variation.

We argue there could be many reasons (such as pretraining data) for the unexpected results pattern. However, since there is more of a consistent performance pattern across models than there is across languages, we believe it is far more likely that differences in performance are not necessarily due to language transparency, but rather model architecture and therefore that the structural differences between these languages is not a significant contributing factor to performance patterns. If this were true (which we think is probable), this is a good sign for these models, since language structure differences are not proving to be a challenge to performance, but rather some other factor.

6.4 Model Architecture and Performance

One clear pattern we can see is that multilingual BERT’s embeddings consistently perform much worse than both XLM and DistilBERT. This is especially clear in French results, where the gap between BERT and the other two models is sometimes more than 20 points. A relatively simple explanation for this is the size of each model in terms of number of parameters. Indeed, as model size increases, performance on both tasks increases (BERT -> DistilBERT -> XLM). Of course, correlation is no evidence of causation; However, if this were true, it is quite consistent with other trends in recent NLP studies. In the case of this study, we can say that bigger is (almost) always better, with XLM mostly performing best, followed by DistilBERT and BERT. Though, this is somewhat undesirable on a larger scale since, ideally, we would hope that it would not require such a large model to encode the information we probe for.

6.5 Bare Number Words Perform Better

On both tasks, our results also show that bare numbers performed equal to or better than numbers in sentences. We propose that this is due to the sentences creating noisiness, thus creating more difficulty for a model to know exactly where it should be looking for the necessary information to complete the tasks. This is very much the case for task 2, where we believe it would be harder to know the magnitude a sentence is referring to than merely if the sentence is grammatical. We argue that, besides adding noise, this method of probing exploits

a possible weakness in masked language modeling as a pretraining method. That is, given that masked language modeling’s task is to predict appropriate (grammatical) words, there may be less emphasis on learning the underlying semantics of those words, thus the better performance on task 1 sentences and worse performance on task 2 sentences.

7 Further Work

As this work is an exploration of a new probing method for state-of-the-art language model architecture, there are surely a number of ways to extend from it.

Though we discussed it briefly here, exploring the architectural reasons for the shortcomings of these pretrained embeddings, especially in the case of task 2 and with sentences is an important area for future work. Indeed, in a similar task from Wallace et al. (2019), BERT was also found to have poor performance. In the future, several more specific probes could be designed to test for understanding of magnitude in various linguistic contexts to find strengths and weaknesses of transformer-based models. A particularly interesting case would be in testing magnitude comprehension in sentences of varying structures. Our sentence templates used in this study are few, and experimenting with other varieties could prove to be insightful.

Our experiment also made use of the idea of language transparency. We also find this to be a topic for possible further work. Namely, is there a method to reliably measure transparency of languages to predict performance on numerical understanding tasks such as these? We believe this may be possible through measuring complexities of grammars which generate number words in each language. Overall, in future extensions of this study, there is room for more languages, sentence types, task renditions, and models.

8 Conclusion

In this paper, we introduced methods for probing the multilingual compositional reasoning capabilities of transformer-based models’ pretrained embeddings over natural language numbers. From our experiments, we’ve shown that these pretrained embeddings show some capabilities in making grammatical judgments of number words, though they are less capable of making value com-

parisons. In addition, we find that results generally follow a trend based upon model size. Our results are in accord with previous work such as Wallace et al. (2019); However, we have also highlighted further model weaknesses through our probing methods. Therefore, the opportunities for future work, especially with a multilingual focus, are plenty.

Acknowledgements

Thank you to all anonymous reviewers for your helpful comments. Thank you to Professor Shane Steinert-Threlkeld for your guidance and throughout all stages of this paper. Thank you to Professor Luke Zettlemoyer for your advice on our paper in its earlier stages. Thank you to the several volunteer native speakers for grammaticality judgments. –This work was facilitated through the use of advanced computational, storage, and networking infrastructure provided by the Hyak supercomputer system and funded by the STF at the University of Washington.–

References

- Yonatan Belinkov and James R. Glass. 2019. Analysis methods in neural language processing: A survey. *Transactions of the Association for Computational Linguistics*, 7:49–72.
- J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding.
- Allyson Ettinger. 2020. What bert is not: Lessons from a new suite of psycholinguistic diagnostics for language models. *Transactions of the Association for Computational Linguistics*, 8:34–48.
- Allyson Ettinger, Ahmed Elgohary, and Philip Resnik. 2016. Probing for semantic evidence of composition by means of simple classification tasks. In *Proceedings of the 1st Workshop on Evaluating Vector-Space Representations for NLP*, pages 134–139, Berlin, Germany. Association for Computational Linguistics.
- Guillaume Lample and Alexis Conneau. 2019. Cross-lingual language model pretraining. *CoRR*, abs/1901.07291.
- K. F. Miller, C. M. Smith, J. Zhu, and H. Zhang. 1995. Preschool origins of cross-national differences in mathematical competence: The role of number-naming systems. *Psychological Science*, 6:56–60.
- T. Ogawa. num2words.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2020. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter.
- Eric Wallace, Yizhong Wang, Sujian Li, Sameer Singh, and Matt Gardner. 2019. Do nlp models know numbers? probing numeracy in embeddings. *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R’emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.

A Fine-Tuning

The fine-tuning results shown in (Figure 3) are validation accuracies after one epoch of training. When trained for 20 epochs (Figure 4), the models reach over 99% accuracy on all languages and tasks except for Danish which reaches 95% on sentence tasks.




Figure 3: Fine-tuned task validation accuracy per language/model, when trained for 1 epoch.




Figure 4: Fine-tuned task validation accuracy per language/model, when trained for 20 epochs.

B Data Generation

B.1 Dataset Parameters

The parameters below were used to generate data for each language per each model per each task variation:

- Data Gen. Seed: 1
- Data Gen. Number Range: [0-999]
- Train Set Size: 30,000
- Validation Set Size: 10,000
- Test Set Size: 10,000
- Shuffle = True

B.2 Task 1 Data

For both variants of task 1 (sentences/bare), grammatical data are generated by creating random numbers then converting them to text through the

num2words (Ogawa) package. For the ungrammatical data, two grammatical numbers are randomly generated, both converted to text, then appended together to create an ungrammatical number. For example the ungrammatical number “fifty-five two hundred” is the combination of “fifty-five” and “two hundred”. Another example made from the same original elements could be: “two fifty-five hundred”. Grammatical numbers used in splits, however, were only split such that the resulting elements were grammatical words themselves. So, for example, an non-continuous number word string like “fi-tfy te nnine” would never occur.

Since generating numbers with appendage can naturally occur in ungrammatical number words being longer than grammatical, we control for length by limiting our set of ungrammatical number words to words that are *at most* as long as the longest grammatical number. Lastly, we ensure no grammatical numbers are accidentally created in this process by keeping a list of known grammatical numbers in text form (generated by num2words) that ranges from number sufficiently higher than our generation range. For example, if our number word generation ranges from 0-1000, we would make this list of known grammatical numbers from 1-100,000,000. These number words are finally either left bare or inserted into sentences to form our x inputs and are labeled 0 if grammatical and 1 if ungrammatical.

B.3 Task 2 Data

The data for the semantic task are generated by creating pairs of random (grammatical) number words and labeling the pair with one of two categories: 0 if the first numbers is larger than the second and 1 if the second is larger than the first. Through our process of data generation, we ensure that there are never two pairs using the same number. They are then converted to text form for input to a model. These number words are finally either left bare or inserted into sentences to form our x inputs. When numbers are used in sentence templates, it is ensured that the numbers are used in the same template. For example, given the number pair “five” and “six”, we could compare the sentences: “There are five apples.” and “There are six apples.”.

C Modeling

C.1 Pytorch Hugging Face Transformers

We use the configurations below of transformers from Hugging Face Transformers (Wolf et al., 2019) in Pytorch on all of our reported experimental runs. Average runtimes were all around 1 hour or less.

- **DistilBERT:**
 - Class: DistilBertForSequenceClassification
 - Config: distilbert-base-multilingual-cased
 - Tokenizer: DistilBertTokenizer
 - Num. Parameters: 134 million total
- **BERT**
 - Class: BertForSequenceClassification
 - Config: base-multilingual-cased
 - Tokenizer: BertTokenizer
 - Num. Parameters: 110 million total
- **XLNet**
 - Class: XLNetForSequenceClassification
 - Config: xlm-mlm-100-1280
 - Tokenizer: XLNetTokenizer
 - Num. Parameters: ~550 million total (inexact)

C.2 Hyperparameters

All experiments which produced our final results shown in the paper were run with the following hyperparameters which were selected manually by tuning for accuracy over a validation set:

- Epochs: 20 (Range: 10-20)
- Learning Rate: 0.00001 (Range: 1e-5 - 1e-4)
- Minibatch size: 32

C.3 Infrastructure

- GPU: Nvidia Tesla P100
- CUDA Version: 10.1
- Python Version: 3.7

C.4 Code Repository

Our Github repository can be found [here](#). Code is subject to change after publishing of this paper. Refer to the Github README for latest information.