

Pretraining with Artificial Language: Studying Transferable Knowledge in Language Models

Ryokan Ri and Yoshimasa Tsuruoka

The University of Tokyo

7-3-1 Hongo, Bunkyo-ku, Tokyo, Japan

{li0123,tsuruoka}@logos.t.u-tokyo.ac.jp

Abstract

We investigate what kind of structural knowledge learned in neural network encoders is transferable to processing natural language. We design *artificial languages* with structural properties that mimic natural language, pretrain encoders on the data, and see how much performance the encoder exhibits on downstream tasks in natural language. Our experimental results show that pretraining with an artificial language with a nesting dependency structure provides some knowledge transferable to natural language. A follow-up probing analysis indicates that its success in the transfer is related to the amount of encoded contextual information and what is transferred is the knowledge of *position-aware context dependence* of language. Our results provide insights into how neural network encoders process human languages and the source of cross-lingual transferability of recent multilingual language models.

1 Introduction

Pretrained language models (Devlin et al., 2019; Yang et al., 2019; Raffel et al., 2020) have demonstrated strong empirical performance not only within a language but also *across* languages. Language models pretrained with a mix of monolingual corpora, such as multilingual BERT, exhibit a decent zero-shot cross-lingual transfer capability, *i.e.*, a model fine-tuned in a single source language (L1) can solve the task in another language (L2) (Conneau et al., 2020a; Xue et al., 2021). Surprisingly, the transfer happens without lexical overlaps between L1 and L2 (Karthikeyan K and Roth, 2020; Conneau et al., 2020b) or even without joint pretraining (Artetxe et al., 2020): an encoder only pretrained on L1 can be transferred to L2 without any parameter updates. These results suggest that, whether the encoder is trained on single or multiple languages, it learns some *transferable knowledge* about language.

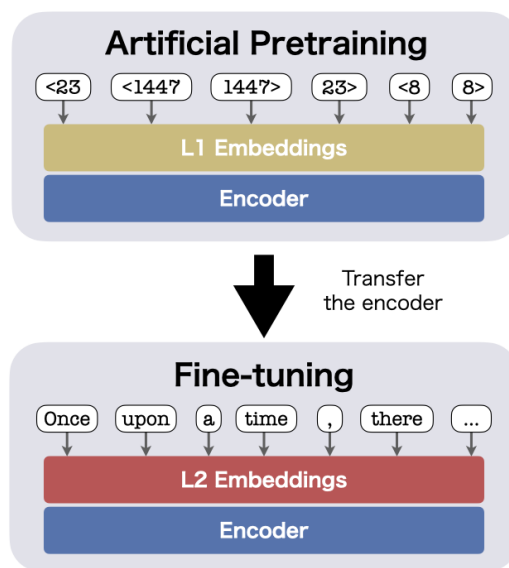


Figure 1: Transfer from artificial language to natural language. The artificial language encodes some structural properties (*e.g.*, token distributions, dependency structures) and we study how the learning of such properties can be transferred to natural language.

However, the characteristics of such transferable knowledge are still underexplored. Recent studies with the probing methodology (Hupkes and Zuidema, 2018; Conneau et al., 2018) have revealed that multilingual BERT captures language-independent linguistic structures such as universal dependency relations (Chi et al., 2020) and subjecthood (Papadimitriou et al., 2021), but it remains unknown whether learning such linguistic properties actually contributes to the performance, and whether there exists more abstract knowledge transferred across languages.

In this study, we try to shed light on these questions with the framework of the Test for Inductive Bias via Language Model Transfer (Papadimitriou and Jurafsky, 2020), focusing on designing *artificial languages* with natural-language-like structural properties (Figure 1). We pretrain encoders

with artificial languages and transfer the encoders to natural language tasks with their parameters frozen. This enables us to see how learning the specific structural properties of the artificial language affects the downstream performance.

Specifically, we explore whether it is beneficial for the encoder to know the following two characteristics of natural language: word distributions and latent dependency structures. We design artificial languages that represent such characteristics and perform an extensive study with different encoder architectures (LSTM and Transformer) pre-training objectives (causal and masked language modelings).

The contribution is summarized as follows:

- We first start by complementing the study in [Papadimitriou and Jurafsky \(2020\)](#). We train LSTM and Transformer encoders with the sentence-level causal language modeling task and evaluate the encoders in English. We show that an artificial language that models simple statistical dependency within a sentence provides decent transferable knowledge on natural language modeling. Furthermore, we find that the inductive bias of a nesting head-to-tail dependency structure is more useful than a flat one.
- We then proceed to investigate transfer learning in masked language modeling ([Devlin et al., 2019](#)), one of the current dominant pre-training paradigms. We evaluate pretrained Transformer encoders with dependency parsing and confirm that the nesting dependency structure is important to learn the structure of natural language.
- We hypothesize that the transfer performance of pretrained encoders is related to the way the encoder preserves the input contextual information in the output vectors. We perform a probing experiment and find that the artificial language with the nesting dependency structure trains encoders to encode the information on adjacent tokens into the output vector of each token. We conclude this paper with the hypothesis that a part of transferable knowledge in language models could be explained by the knowledge of *position-aware context dependence* of language.

2 Related Work

2.1 Transferable Structural Knowledge in Pretrained Encoders

Multilingual language models trained with masked language modeling objective ([Devlin et al., 2019](#); [Doddapaneni et al., 2021](#)) have demonstrated a surprisingly strong cross-lingual transfer capability ([Liu et al., 2020](#)), given the model is only trained with a mix of monolingual corpora. This leads to several studies investigating the source of the cross-lingual capability of multilingual models.

An early common hypothesis was that the models take advantage of a common word-piece vocabulary across languages ([Wu and Dredze, 2019](#); [Pires et al., 2019](#)), which provides cross-lingual alignment signals to learn useful multilingual representations. However, this hypothesis has been questioned by recent studies ([Karthikeyan K and Roth, 2020](#); [Conneau et al., 2020b](#)) which show that shared word-pieces only play a minor role in the performance. These studies suggest that the model can exploit abstract structures of languages to learn shared multilingual representations.

Another line of research suggests that the learning of transferable knowledge happens even in monolingual pretraining. [Artetxe et al. \(2020\)](#) showed that a Transformer encoder pretrained only on L1 exhibits strong cross-lingual transfer performance simply by aligning the L2 embeddings to the encoder. [Papadimitriou and Jurafsky \(2020\)](#) pretrained LSTM encoders with natural languages and non-linguistic data (*e.g.*, code, music, and artificial data) to demonstrate that the encoders achieve reasonable performance in Spanish language modeling. These studies provide additional evidence for the existence of transferable linguistic knowledge learned in the model.

Then what is such knowledge? Probing studies ([Hupkes and Zuidema, 2018](#); [Conneau et al., 2018](#)) have revealed that the model captures language-independent structures such as universal dependency relations ([Chi et al., 2020](#)) and subjecthood ([Papadimitriou et al., 2021](#)). However, the probing methodology does not answer whether such linguistic knowledge contributes to the performance in cross-lingual transfer.

In this study, we shed light on this question by studying transfer learning from artificial language with the Test for Inductive Bias via Language Model Transfer (TILT) ([Papadimitriou and Jurafsky, 2020](#)). This framework enables us to

assess if abstract features generalizable to L2 (natural language) are encoded in L1. Here we explicitly design artificial languages with some structural properties as L1 to investigate their transferability.

2.2 Studying Language Models with Artificial Language

To study the behavior of language models, several studies have employed a specific type of *artificial language*: artificial variants of natural languages. A typical experimental framework is as follows: (1) create an artificial language that differs from a natural language in one linguistic property, such as word orders (Sinha et al., 2021b; Dufter and Schütze, 2020; Sinha et al., 2021a), scripts (Karthikeyan K and Roth, 2020; Dufter and Schütze, 2020; Conneau et al., 2020b), or morphology (Ravfogel et al., 2019); (2) train or evaluate the natural/artificial language models and compare the performance to analyze the model’s sensitivity to the linguistic property.

However, this methodology is limited to studying linguistic properties that are easily editable to create artificial variants and also offers limited control over the experiments. To overcome this problem, White and Cotterell (2021) created artificial languages by defining their own probabilistic context-free grammars (PCFG). As the concurrent work, Chiang and yi Lee (2022) trained Transformer encoders on artificial data with token dependencies in the sequences and showed that they perform reasonably well on the GLUE benchmark (Wang et al., 2019). In this research, we design artificial languages with certain structural properties from scratch to study knowledge transferable to natural language.

3 Approach

3.1 Experimental Framework

We first describe the experimental framework used throughout this paper, the Test for Inductive Bias via Language Model Transfer (TILT) introduced by Papadimitriou and Jurafsky (2020). TILT consists of pretraining and transfer steps:

1. Pretrain an encoder with a pretraining task in the source language (L1). We explore pretraining with causal language modeling in §4 and masked language modeling in §5.
2. Transfer the encoder to the target language (L2) in a downstream task. As we are inter-

ested in structural prior knowledge learned in the encoder, we discard the learned L1 word embeddings and initialize the embedding layer with the L2 vocabulary. We then train the model with the encoder parameters frozen and evaluate the task performance.

TILT reveals how transferrable the computation induced to solve the L1 pretraining task is to processing L2. In this study, we are interested in the transferability of certain types of structures to natural language, and thus we primarily use hand-designed artificial languages with the structural properties as L1 and natural language as L2.

3.2 Designing Artificial Languages

Artificial languages are designed to mimic a certain property of natural language. After providing a formal definition of artificial language, we introduce several languages used in this paper.

3.2.1 Formulation of Artificial Language

A artificial language refers to a set of a vocabulary and algorithms to generate sequential data for pretraining. Each language has a sentence-length distribution $p_{len}(l)$, token vocabulary $\{w|w \in \mathcal{V}\}$, and sentence-sampling function $f(l) : l \mapsto \mathcal{V}^l$. The training data is generated sentence by sentence as follows: we first sample a sentence length ($l \sim p_{len}(l)$) and then sample a sequence of tokens of that length ($[w_1, \dots, w_l] \sim f(l)$).

In this study, the token vocabulary V simply consists of integers (or integers with a special symbol) and is not intended to correspond to a vocabulary of any natural language. Also the sentence-length distribution $p_{len}(l)$ is fitted with a baseline dataset in each experiment. The focus is how to design the sentence-sampling function $f(l)$. This determines what kind of characteristics we want to encode in the artificial dataset.

3.2.2 Modeling Word Distribution

Words in natural language are distributed in non-trivial fashions. We will study whether prior knowledge of token distribution facilitates learning from natural language. We first present the simplest artificial language that serves as a baseline.

Uniform language samples each token in a sentence independently and uniformly. Specifically, the probability of a token w being sampled is

$$p(w) = \frac{1}{|\mathcal{V}|}. \quad (1)$$

However, this deviates from the token distribution of natural language. Natural language is empirically known to follow the Zipf’s law (Zipf, 1949), *i.e.*, the relation between the frequency of a word and its rank is given by $\text{frequency}(w) \propto \text{rank}(w)^{-\alpha}$. The coefficient α is typically around 1, although the coefficient shows some variation according to the corpus domain (Zanette and Montemurro, 2005).

Zipf language captures this property and samples each token w from the following probability distribution assuming $\alpha = 1$:

$$p(w) \propto \frac{1}{\text{rank}(w)}. \quad (2)$$

The two languages introduced so far generate tokens in a sentence independently. However, words within a sentence of natural language are known to have statistical dependencies, *i.e.*, specific cooccurrence patterns (Church and Hanks, 1989). Consider the sentence “*The cat and dog are fighting over food.*” The words *the* and *cat* would cooccur much more often than by chance because *cat* (noun) is dependent on *the* (determinant); so would *dog* and *cat* because they are topically related. The words in a sentence are usually coherent according to some syntactic and semantic dependencies.

Log-linear language is designed to capture this property. Inspired by the log-linear model in Arora et al. (2016), tokens in a sentence s are drawn from the following probability distribution:

$$p(w|s) \propto \exp(\vec{c}_s \cdot \vec{v}_w), \quad (3)$$

where \vec{c}_s is the discourse vector of the sentence and \vec{v}_w is the word vector of the token w . Intuitively, we can imagine that the discourse vector represents the *topic* of the sentence and determines the unigram distribution over the vocabulary (Blei et al., 2003). Sampling tokens this way, non-trivial cooccurrence patterns within sentences emerge in the language.

We speculate that pretraining with the Log-linear language will endow the model with an inductive bias to aggregate the context in a sentence to predict the identity or property of tokens, which is likely to benefit natural language processing.

In the experiments, the word vectors \vec{v}_w are initialized with the normal distribution, and the discourse vector \vec{c}_s is also drawn from the normal distribution each time we generate a sentence. We

set the dimension of the word and discourse vector to 10 as we empirically find that this makes the entire token distribution close to the Zipfian distribution.

3.2.3 Modeling Latent Dependency Structure

Sentences in natural language are known to have latent structures, which are often described in the form of trees (Chomsky, 1957) or dependency graphs (Mel’čuk, 1988). Now we consider how to endow the sampled tokens with such structures.

In this study, we adopt a dependency-based latent structure. Words in sentences of natural language often have dependency relations and the existence of a certain word can be predictive of another word (*e.g.*, the verb *am* always cooccurs with *I*). We hypothesize that, pretrained on such data, language models may acquire inductive bias towards finding relations between tokens in the input, which is presumably important in processing natural language.

Inspired by Papadimitriou and Jurafsky (2020), we design algorithms that generate structured sentences given a set of tokens sampled with any of the strategies described in §3.2.2. The general idea is that half of the tokens (heads) in the vocabulary are all paired with another half of tokens (tails). A pair of head and tail can be represented in right and left brackets with the same integer (*e.g.*, “<123”, “123>”). The pairs always appear together in a sentence and express simple dependency relations. After determining the sentence length $l \sim f(l)$, we first sample $\frac{l}{2}$ (rounded to an integer) pairs of head and tail and then arrange them with one of the following structures.

Flat Dependency structure simply arranges the tokens randomly while keeping the right order of the brackets (*e.g.*, [“<5”, “<84”, “5>”, “<123”, “123>”, “84>”]). The dependency arcs are allowed to be crossed and thus often result in a non-projective dependency structure.

Nesting Dependency language, by contrast, does not allow any dependency arcs to be crossed, and the brackets are nested hierarchically (*e.g.*, [“<5”, “<84”, “84>”, “5>”, “<123”, “123>”]). The sentences are generated from the stack-based algorithm described in Appendix A.

These structures are similar to the **Parenthesis** languages used to study the inductive bias of language models in Papadimitriou and Jurafsky (2020). However, our Dependency languages differ from them in how to represent the head and tail tokens. In the Parenthesis language, the head and

tail are represented with the same token (e.g., [“5”, “84”, “84”, “5”, “123”, “123”]), which we argue deviates from the dependency structure in natural language, because in natural language, dependency relations usually hold between different words (e.g., *I* and *am*). We will show that this difference is in fact crucial and draw a different conclusion from Papadimitriou and Jurafsky (2020) on the importance of the nested structure (§4.2).

4 Causal Language Model Pretraining with Artificial Language

In this section, we complement the study of Papadimitriou and Jurafsky (2020). While they studied the inductive bias learned in LSTM encoders with some artificial languages, here we provide additional studies with the newly introduced Log-linear and Dependency artificial languages, and the Transformer encoder.

4.1 Experimental Setups

Task. We study sentence-level causal (left-to-right) language modeling (CLM), where the model needs to predict the next word given the previous context in the sentence. Note that, Papadimitriou and Jurafsky (2020) experiment with language modeling across sentences, but we adopt sentence-level modeling because we would like to focus on the learning of sentence structures here. As we will see in §4.2, we observe the same tendency in regard to the effect of artificial pretraining where we share the setups. The task performance is measured by the average perplexity scores for each token.

Model. We study two encoder architectures: LSTM (Hochreiter and Schmidhuber, 1997) and Transformer (Vaswani et al., 2017). These architectures are known to exhibit different abilities in capturing the underlying hierarchical structure of sequential data (Tran et al., 2018).

The size of word embeddings is set to 300. For both LSTM and Transformer encoders, the number of layers is set to 3, and the number of parameters is configured to be the same (6.9M parameters) to enable a fair comparison between architectures (for further details, see Appendix B).

Pretraining Data. We generate artificial corpora with three unstructured languages, which randomly arrange the tokens sampled from Uniform, Zipf, and Log-linear languages, and four structured languages which combine the Zipf sampling strategy with the structures of Flat Parenthesis, Nesting

Parenthesis, Flat Dependency, and Nesting Dependency.

We also experiment with natural language corpora. We create training corpora from Wikipedia dumps of English, Japanese, and Spanish. The sentences are tokenized with the Moses tokenizer¹ for English and Spanish and MeCab² for Japanese.

The sentence lengths of artificial data were sampled from the empirical distribution of the English Wikipedia corpus. The size of the vocabulary $|V|$ is set to 32,000 for both artificial and natural corpora, and out-of-vocabulary words in natural language are replaced with the OOV token. For each corpus, we sample 12.8 M sentences and train the model with one iteration over the corpus.

Evaluation Data. We evaluate the pretrained encoders on the Penn Treebank (PTB) corpus (Marcus et al., 1993) with preprocessing from Mikolov et al. (2010). Note that, when we train language models with the pretrained encoders, the parameters of the encoder are not updated and only the English word embeddings are learned from scratch (optimization details in Appendix B.2).


4.2 Results

We provide two baseline models trained on the L2 training corpus from scratch and trained with frozen random weights in the encoder to compare with pretrained encoders. For each configuration, we pretrain three encoders with different random seeds, and for each encoder fine-tuned three models, which results in nine models in total. We summarize the average scores and standard deviations in Figure 2.


The Transformer encoder is more flexible than LSTM. We start by discussing overall trends. We observe that the Transformer encoders give lower perplexity scores compared to LSTM regardless of pretraining language. This tendency is in line with the observations on the surprisingly good transferability of pretrained Transformer encoders to other languages (Conneau et al., 2020a), or even other modalities (Lu et al., 2021; Reid et al., 2022). We think that this is because Transformer encoders are better at aggregating and preserving the context information at each time step, as we will see in §6, presumably because the Transformer architecture has self-attention and residual connections.

¹<https://github.com/moses-smt/mosesdecoder>


²<http://taku910.github.io/mecab/>



(a) Comparison of token distributions.



(b) Comparison of dependency structures.



(c) Comparison of natural languages.

Figure 2: The perplexity scores (the lower the better) on the sentence-level causal language modeling task with the English Penn Treebank dataset. The two baselines (From scratch and Random weights) are not pretrained, and the others are the results of pretrained encoders.

Natural languages are better than the artificial languages. As expected, pretraining with natural languages (English, Spanish and Japanese) provides better encoders for language modeling than the artificial languages both with LSTM and Transformer. However, the performance differences between natural languages seem to be negligible, indicating that there is not much difference in the way the encoders process these different languages, conforming with the observation of cross-lingual transferability of pretrained encoders (Artetxe et al., 2020).

The Uniform and Zipf languages degrade the encoders. Looking at the difference among unstructured languages (Figure 2a), Uniform and Zipf languages give higher perplexities than the Random weights baseline particularly with LSTM. In hindsight, it is natural that encoders would be degraded even from random weights when trained with sequences where tokens are drawn independently from each other because the encoders are not incentivized to use contextual information and will even learn to discard the input information. We will demonstrate this with a follow-up probing experiment in §6.

The Log-linear language provides a useful inductive bias to language modeling. On the

contrary, the Log-linear language gives reasonably lower perplexities compared to Random weights (Figure 2a). This indicates that knowing the existence of statistical dependency within a sentence, or learning to predict tokens from the cooccurrence information, is a useful inductive bias even though the cooccurrence statistics is not necessarily in line with L2.

We do not observe the importance of the nested structure in the Parenthesis languages. Papadimitriou and Jurafsky (2020) showed that LSTM encoders trained on the Flat Parenthesis and Nesting Parenthesis structures do not provide a significant difference in perplexity, and concluded that simple non-hierarchical head-dependent-type relations are important in LSTM language processing. A similar observation can be made in Figure 2b: although the Nesting Parenthesis exhibits the lower average score, there is no significant difference between Flat Parenthesis and Nesting Parenthesis (232.9 ± 30.0 vs. 203.8 ± 7.7 , $p > 0.01$ in Welch’s t-test) with the unstable results of Flat Parenthesis. Also, the trend of the average scores is reversed in Transformer: the Nesting Parenthesis exhibits the higher average score (212.4 ± 8.8) than Flat Parenthesis (191.9 ± 11.8), which makes it difficult to draw a consistent conclusion from here.

However, **the Dependency languages suggest that the nested structure is actually important in language modeling.** While the Parenthesis language represents dependency relations with two identical tokens (e.g., “4543” and “4543”), our Dependency language represents relations with two different tokens (e.g., “<4543” and “4543>”). We expect that expressing dependency relations with two different tokens is closer to natural language and thus provides more viable insights into natural language. When we compare the scores of the Dependency languages, Nesting Dependency provides the lower and more stable perplexity than Flat Dependency with LSTM (175.7 ± 4.3 vs. 187.2 ± 10.7) and the significantly lower score with Transformer (160.6 ± 1.6 vs. 175.7 ± 4.3 , $p > 0.01$ in Welch’s t-test). Overall, Nesting Dependency performs best among other artificial languages, indicating our Dependency language is closer to natural language and the nested structure is useful for language modeling.

5 Masked Language Model Pretraining with Artificial Language

We proceed to investigate transfer learning from artificial languages in one of the most successful pretraining paradigms, masked language modeling (MLM) (Devlin et al., 2019) to see if we can observe similar trends to what we see in the CLM experiment (§4).

5.1 Experimental Setups

Pretraining. To allow for fast experimentation, we train small Transformer encoders. The size of word embeddings is set to 300 and the encoders have three layers (further details in Appendix C). The pretraining datasets are the same as in §4.1.

Downstream Task. We evaluate the pretrained encoders with dependency parsing to see if the structural knowledge learned with artificial language is beneficial to predict the structure of natural language. We use the English EWT dataset from Universal Dependencies (UD) v2.8 (Nivre et al., 2020)³.

Model. We adopt the biaffine graph-based parser (Dozat and Manning, 2017) with the Transformer encoder. The input word representations are the concatenation of word embeddings and character features computed by a character-level bi-directional LSTM encoder (Ling et al., 2015). For

³<https://universaldependencies.org/>




Figure 3: The downstream performance on two syntactic tasks with the English EWT dataset. The two baselines (From scratch and Random weights) are not pretrained, and the others are the results of encoders pretrained with masked language modeling.

the details on fine-tuning these models, please refer to Appendix C.

5.2 Results

We provide two baseline models trained from scratch and trained with random encoder weights. For each pretraining language, we again train three encoders and fine-tune three models for each, and take the mean and standard deviation of the nine models. Figure 3 shows the results.

The unstructured languages do not provide useful transferable knowledge for dependency parsing. The Uniform, Zipf, and Log-linear encoders perform comparably to or worse than the Random weights baseline. This is in contrast with the causal language modeling task, where the Log-linear language at least outperforms the Random weights baseline (§4.2).

On the other hand, **learning from structured languages seems to be important in dependency parsing.** The Dependency encoders outperform the Random weights baseline, and also we can observe that learning from the nesting structure is more effective than the flat structure, and Dependency languages outperform Parenthesis languages, as observed in the CLM in §4.

6 How much contextual information do the pretrained encoders capture?

In the previous sections, we have seen that the encoders pretrained with different artificial languages exhibit various degrees of transferability to natural

language. In this section, we try to explain why pretraining with some artificial languages is better or worse for the transfer to natural language from the perspective of the amount of contextual information in the encoder outputs.

The intuition is, for example, if a pretrained encoder has learned to discard the input information, we cannot expect the encoder to perform well when transferred to any tasks. Also, existing studies show that neural language models assign more importance to local context when they make predictions (Khandelwal et al., 2018; Lai et al., 2020). Can we observe that encoders pretrained with artificial languages exhibit similar patterns to natural languages regarding how they encode the contextual information?

6.1 Experimental Setups

We investigate how much contextual information can be extracted from the outputs of the pretrained encoders by setting up a simple probing task. In this task, the encoder is asked to recover the identity of the contextual words given the contextualized vector of a target word.

Specifically, we first randomly generate 100K sequences of integers with the length of $15 \sim 25$ (close to most frequent sequence lengths in the pretrained corpus) with the vocabulary size 100 and split them into training (90K sequences), validation (5K) and test (5K) sets.

Then we simultaneously train several linear classifiers, each of which predicts the ID of the context word at a fixed relative position to the target word in the sequence, on top of a frozen pretrained encoder. For the encoders pretrained with CLM in §4, the target word is the last word in sequences and the classifiers predict the words at the positions of $[-9, -4, -3, -2, -1, 0]$; for the encoders pretrained with MLM in §5, the target word is the middle word and the classifiers predict the words at $[-6, -3, -2, -1, 0, 1, 2, 3, 6]$.

After training, we measure the accuracy of predicting the words at each position on the test set and interpret this as how much information on each contextual word the encoder preserves.

6.2 Results

Figure 4 summarizes the results of the encoders trained in §4 and §5.

The amount of the encoded contextual information can explain the transfer performance in some obvious cases. In the experiment of CLM

(Figure 2a), we observed that the Uniform and Zipf encoders tend to perform worse even than Random weights. Figure 4a and 4d demonstrate that their poor performance is because the encoders are trained to discard the input information. The Uniform and Zipf encoders tend to preserve less contextual information even than Random weights because capturing the contextual information does not lead to solving the pretraining task in these languages.

On the other hand, if words are predictable from the context, encoders are encouraged to learn to preserve the contextual information. The Log-linear encoders trained with CLM encode a decent amount of the contextual information (Figure 4a and 4d) and also performed best among the unstructured artificial languages in CLM (Figure 2a). Moreover, encoders trained with natural languages (Figure 4c, 4f and 4i) capture not only the local context well (at distance $0 \sim 2$) but also a modest amount of the farther context (at distance $3 \sim$), which is consistent with the existing observation that LSTM encoders trained with natural language are better at memorizing the inputs than ones trained with randomly sampled data (Liu et al., 2018). In these cases, the downstream performance and the amount of the encoded contextual information seem to be correlated.

However, this trend is not as clear when comparing the structured artificial languages. For example, the Nesting Dependency encoders perform the best for the downstream tasks among the structured artificial languages but do not necessarily in the probing task (Figure 4b and 4e).

The nesting structure seems to facilitate encoders to remember the local context with MLM. The difference between the Nesting and Flat languages is striking in Figure 4f. The Nesting encoders are consistently better at capturing the local contextual information (at positions $-2 \sim 2$) than their flat counterparts, which may explain the better performance of the Nesting encoders in dependency parsing (Figure 3), given that the local contextual information is particularly important to predict the syntactic characteristics of words (Levy and Goldberg, 2014; Ri and Tsuruoka, 2020).

7 Discussion and Future Work

In this paper, we studied what kind of structural properties in pretraining data is useful to train encoders for natural language tasks. We have found

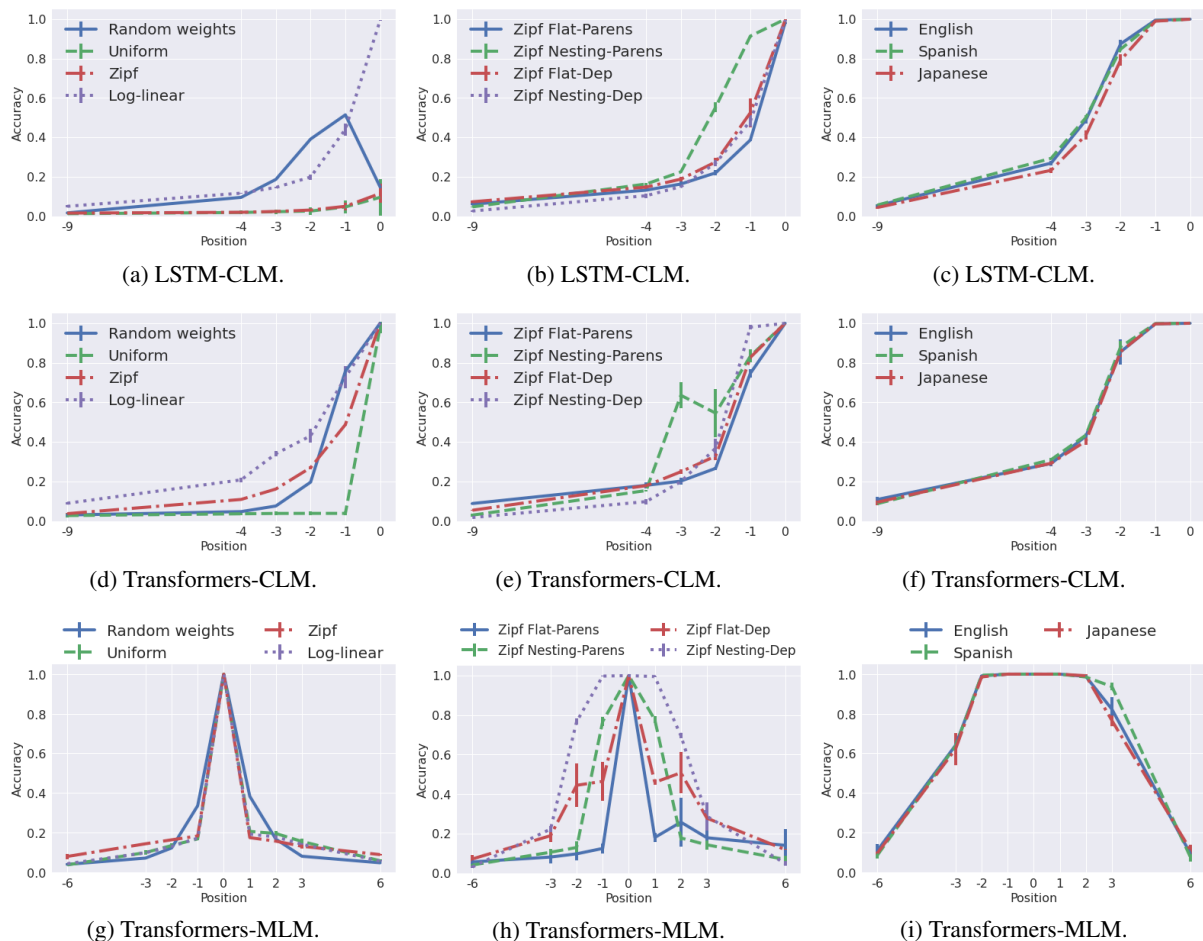


Figure 4: The accuracy of the task of recovering the contextual words from the encoder output of target words.

that to achieve decent results, L1 needs at least statistical dependency in a sentence (§4), and having the head-to-tail dependency with the nesting structure is further beneficial (§4 and §5). The probing experiment in §6 suggests that the encoders trained with languages with the above characteristics are good at capturing the positions and identities of the context words.

From these observations, we suggest a tentative answer to the initial research question: what knowledge in pretrained encoders are transferred across different languages? That is *position-aware context dependence* of language, in other words, **“tokens in a sequence can be characterized by its neighbor tokens at specific positions”**.

We think that it can explain the success of transferring the encoder across languages to some extent. To solve natural language tasks, it is often useful to characterize words in a sentence by the words around them. For example, to understand the semantics of a sentence, it would be useful to look for the subject by looking for a noun that precedes

the word *is*; to parse a sentence, a word can be identified as a noun because it follows the article *the*. If the encoder computes the output representation of a word in a sentence by aggregating the information from its surrounding words, that should be a useful inductive bias to solve most NLP tasks in any language. Also, it is easy to imagine that the knowledge of position-aware context dependence gives a reasonable prior for solving sequence modeling problems in other domains, which may explain the success of cross-modality transfer of language models (Lu et al., 2021; Reid et al., 2022).

Of course, we do not expect that the knowledge of position-aware context dependence explains every aspect of the success of cross-lingual transfer. As future work, we need further investigation for a more fine-grained view of the transferred knowledge. Important questions include how much the model size affects the transferability of the encoder or if there is any difference in the knowledge transferred among different downstream tasks.

Acknowledgement

We thank the anonymous reviewers for their insightful comments and constructive suggestions to improve the paper.

References

- Sanjeev Arora, Yuanzhi Li, Yingyu Liang, Tengyu Ma, and Andrej Risteski. 2016. [A Latent Variable Model Approach to PMI-based Word Embeddings](#). *Transactions of the Association for Computational Linguistics*, 4:385–399.
- Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2020. [On the Cross-lingual Transferability of Monolingual Representations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- David M. Blei, A. Ng, and Michael I. Jordan. 2003. [Latent Dirichlet Allocation](#). *Journal of Machine Learning Research*, 3:993–1022.
- Ethan A. Chi, John Hewitt, and Christopher D. Manning. 2020. [Finding Universal Grammatical Relations in Multilingual BERT](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Cheng-Han Chiang and Hung yi Lee. 2022. [On the Transferability of Pre-trained Language Models: A Study from Artificial Datasets](#). In *Proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence*.
- Noam Chomsky. 1957. *Syntactic Structures*. De Gruyter Mouton.
- Kenneth Ward Church and Patrick Hanks. 1989. [Word Association Norms, Mutual Information, and Lexicography](#). In *27th Annual Meeting of the Association for Computational Linguistics*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020a. [Unsupervised Cross-lingual Representation Learning at Scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Alexis Conneau, German Kruszewski, Guillaume Lample, Loïc Barrault, and Marco Baroni. 2018. [What you can cram into a single vector: Probing sentence embeddings for linguistic properties](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*.
- Alexis Conneau, Shijie Wu, Haoran Li, Luke Zettlemoyer, and Veselin Stoyanov. 2020b. [Emerging Cross-lingual Structure in Pretrained Language Models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- 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*.
- Sumanth Doddapaneni, Gowtham Ramesh, Anoop Kunchukuttan, Pratyush Kumar, and Mitesh M. Khapra. 2021. [A Primer on Pretrained Multilingual Language Models](#). *ArXiv*, abs/2107.00676.
- Timothy Dozat and Christopher D. Manning. 2017. [Deep Biaffine Attention for Neural Dependency Parsing](#). In *International Conference on Learning Representations*.
- Philipp Dufter and Hinrich Schütze. 2020. [Identifying Elements Essential for BERT’s Multilinguality](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long Short-Term Memory](#). *Neural Computation*, 9:1735–1780.
- Dieuwke Hupkes and Willem Zuidema. 2018. [Visualisation and ‘Diagnostic Classifiers’ Reveal how Recurrent and Recursive Neural Networks Process Hierarchical Structure \(Extended Abstract\)](#). In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*.
- Stephen Mayhew Karthikeyan K, Zihan Wang and Dan Roth. 2020. [Cross-Lingual Ability of Multilingual BERT: An Empirical Study](#). In *International Conference on Learning Representations*.
- Urvashi Khandelwal, He He, Peng Qi, and Dan Jurafsky. 2018. [Sharp Nearby, Fuzzy Far Away: How Neural Language Models Use Context](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*.
- Yi-An Lai, Garima Lalwani, and Yi Zhang. 2020. [Context Analysis for Pre-trained Masked Language Models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*.
- Omer Levy and Yoav Goldberg. 2014. [Dependency-Based Word Embeddings](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*.
- Wang Ling, Chris Dyer, Alan W Black, Isabel Trancoso, Ramón Fernandez, Silvio Amir, Luís Marujo, and Tiago Luís. 2015. [Finding Function in Form: Compositional Character Models for Open Vocabulary Word Representation](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*.
- Nelson F. Liu, Omer Levy, Roy Schwartz, Chenhao Tan, and Noah A. Smith. 2018. [LSTMs Exploit Linguistic Attributes of Data](#). In *Proceedings of The Third Workshop on Representation Learning for NLP*.

- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. [Multilingual Denoising Pre-training for Neural Machine Translation](#). *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Kevin Lu, Aditya Grover, P. Abbeel, and Igor Mordatch. 2021. [Pretrained Transformers as Universal Computation Engines](#). *ArXiv*, abs/2103.05247.
- Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. [Building a Large Annotated Corpus of English: The Penn Treebank](#). *Comput. Linguistics*, 19:313–330.
- Igor Mel'čuk. 1988. *Dependency Syntax: Theory and Practice*. State University Press of New York.
- Tomas Mikolov, Martin Karafiát, Lukás Burget, Jan Honza Cernocký, and Sanjeev Khudanpur. 2010. [Recurrent neural network based language model](#). In *Proceedings of Annual Conference of the International Speech Communication Association*.
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Jan Hajič, Christopher D. Manning, Sampo Pyysalo, Sebastian Schuster, Francis Tyers, and Daniel Zeman. 2020. [Universal Dependencies v2: An Evergrowing Multilingual Treebank Collection](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*.
- Isabel Papadimitriou, Ethan A. Chi, Richard Futrell, and Kyle Mahowald. 2021. [Deep Subjecthood: Higher-Order Grammatical Features in Multilingual BERT](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*.
- Isabel Papadimitriou and Dan Jurafsky. 2020. [Learning Music Helps You Read: Using Transfer to Study Linguistic Structure in Language Models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. [How Multilingual is Multilingual BERT?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.
- Ofir Press and Lior Wolf. 2017. [Using the output embedding to improve language models](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- 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*.
- Machel Reid, Yutaro Yamada, and Shixiang Shane Gu. 2022. [Can Wikipedia Help Offline Reinforcement Learning?](#) *ArXiv*, abs/2201.12122.
- Ryokan Ri and Yoshimasa Tsuruoka. 2020. [Revisiting the Context Window for Cross-lingual Word Embeddings](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- 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](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*.
- Koustuv Sinha, Prasanna Parthasarathi, Joelle Pineau, and Adina Williams. 2021b. [UnNatural Language Inference](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*.
- Ke Tran, Arianna Bisazza, and Christof Monz. 2018. [The Importance of Being Recurrent for Modeling Hierarchical Structure](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention Is All You Need](#). In *Advances in Neural Information Processing Systems*, volume 30.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *International Conference on Learning Representations*.
- 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*.
- Shijie Wu and Mark Dredze. 2019. [Beto, Bentz, Becas: The Surprising Cross-Lingual Effectiveness of BERT](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A Massively Multilingual](#)

[Pre-trained Text-to-Text Transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.

Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. [XLNet: Generalized Autoregressive Pretraining for Language Understanding](#). In *Advances in Neural Information Processing Systems*, volume 32.

Damián H. Zanette and Marcelo A. Montemurro. 2005. [Dynamics of Text Generation with Realistic Zipf's Distribution](#). *Journal of Quantitative Linguistics*, 12:29 – 40.

George Kingsley Zipf. 1949. *Human behavior and the principle of least effort*. Addison-Wesley Press.

Appendix for “Pretraining with Artificial Language: Studying Transferable Knowledge in Language Models”

A Generating the Nesting Structure

In the Nesting languages introduced in §3.2.3, tokens are ordered in a way that any dependency arcs in a sequence are not crossed. This is realized by the stack-based algorithm in Algorithm 1. We set the probability of closing a dependency pair to 0.4 following Papadimitriou and Jurafsky (2020).

Algorithm 1 Generating a sentence from the Nesting Dependency language.

Input: input_pairs: Stack[(w, w)]
Output: sentence: List[w]

```
1: closing_stack = []
2: while not input_pairs.is_empty() do
3:   Uniform sampling  $p \sim [0, 1]$ 
4:   if closing_stack.is_empty() or  $p < 0.4$  then
5:     head, tail = input_pairs.pop()
6:     sentence.append(head)
7:     closing_stack.push(tail)
8:   else
9:     tail = closing_stack.pop()
10:    sentence.append(tail)
11:  end if
12: end while
13: while not closing_stack.is_empty() do
14:   tail = closing_stack.pop()
15:   sentence.append(tail)
16: end while
17: return sentence
```

B Details of Causal Language Modeling Task

B.1 Model configuration

For the experiment with causal language modeling (§4), we set the number of layers of the LSTM and Transformer encoders to 3 and configure them so that they have the same number of parameters (2.1 M parameters without the embedding and output projection layers). The details of configuration are shown in Table 1 and Table 2.

The weights of the output projection layer are tied with the word embedding layer (Press and Wolf, 2017). Note that, to enable this, the LSTM encoder has an additional linear layer to project the hidden vector (294 dim) to the input size (300 dim), which the Transformer encoder does not have.

# of layers	3
input size	300
hidden size	294

Table 1: Configuration of the LSTM encoder.

# of layers	3
size	300
feedforward size	600
# of attention heads	4

Table 2: Configuration of the Transformer encoder.

B.2 Optimization

We optimize the pretrained models for 10k steps with 12.8 M sentences and the batch size of 128 using AdamW (Loshchilov and Hutter, 2019). We use the the Noam Learning rate scheduler described in Vaswani et al. (2017) with the warmup steps of 4000, and the other hyper-parameter details are shown in Table 3. We use the same hyper-parameters for fine-tuning with the L2 language.

Name	Value
Pretraining minimum sentence length	6
Pretraining maximum sentence length	60
Dropout	0.1
Weight decay	0.01
Adam β_1	0.9
Adam β_2	0.98
Adam ϵ	1e-9
Gradient clipping	0.25

Table 3: Hyper-parameters for pretraining.

C Details of Masked Language Modeling Task

C.1 Model configuration

For the experiment with masked language modeling (§5), we set the number of layers of the Transformer encoders to 3. The details of configuration are shown in Table 4 (2.1 M parameters without the embedding and output projection layers).

The hyper-parameters for the masked language modeling task is shown in Table 5. For optimization, we used the same hyper-parameters as in Appendix B.2.

# of layers	3
size	300
feedforward size	600
# of attention heads	4

Table 4: Model configuration of the Transformer encoder.

Mask probability for words	15%
Random-word probability for words	10%
Unmasked probability for words	10%

Table 5: The hyper-parameters for masked language modeling.

D Computing Infrastructure

We ran the experiments on a server with a Intel(R) Xeon(R) CPU E5-2698 v4 @ 2.20GHz CPU and 10 NVIDIA TITAN Xp GPUs. Each pretraining and finetuning were run with a single GPU.