

On Using Distribution-Based Compositionality Assessment to Evaluate Compositional Generalisation in Machine Translation

Anssi Moio^{1,2}, Mathias Creutz², and Mikko Kurimo¹

¹Department of Information and Communications Engineering, Aalto University, Finland

²Department of Digital Humanities, University of Helsinki, Finland

anssi.moio@aalto.fi, mathias.creutz@helsinki.fi, mikko.kurimo@aalto.fi

Abstract

Compositional generalisation (CG), in NLP and in machine learning more generally, has been assessed mostly using artificial datasets. It is important to develop benchmarks to assess CG also in real-world natural language tasks in order to understand the abilities and limitations of systems deployed in the wild. To this end, our GenBench Collaborative Benchmarking Task submission utilises the *distribution-based compositionality assessment* (DBCA) framework to split the Europarl translation corpus into a training and a test set in such a way that the test set requires compositional generalisation capacity. Specifically, the training and test sets have divergent distributions of dependency relations, testing NMT systems’ capability of translating dependencies that they have not been trained on. This is a fully-automated procedure to create natural language compositionality benchmarks, making it simple and inexpensive to apply it further to other datasets and languages. The code and data for the experiments is available at <https://github.com/aalto-speech/dbca>.

1 Introduction

An often-used definition, by Partee (1995), of the concept of *compositionality* of language is that the “meaning of a whole is a function of the meanings of the parts and of the way they are syntactically combined”. A more stringent definition adds that the composition of meaning is *systematic*: each part’s meaning is the same in all the different sequences it appears in, the syntactical rules work the same way for different parts, and the same function determines meaning for different wholes (Fodor and Pylyshyn, 1988; Pavlick, 2022, 2023). Compositionality enables generalising to new meanings by combining familiar parts, and to understand each other language users need to employ common systematic rules.

A number of benchmarks have been developed to assess systematic generalisation from different

perspectives and in different NLP tasks. Many of these consist of artificial data, such as the popular SCAN (Lake and Baroni, 2018) and COGS (Kim and Linzen, 2020) benchmarks. These artificial datasets are typically constructed to be highly systematic, to include straightforward syntactical rules. Natural languages, however, have varied irregularities, idiomatic expressions, and other exceptions to the rules, which make composition of meaning much more complicated than in the case of the highly-regular artificial datasets. It’s therefore important to assess whether NLP systems are able to generalise systematically also in the case of natural language, where systematic rules are obscured by exceptions (Dankers et al., 2022).

Works that aim to assess systematic generalisation in natural language tasks often still synthesise some part of the dataset in order to create test examples where systematic generalisation is needed (for example Li et al. (2021) and Dankers et al. (2022)). This enables precise testing of specific systematic rules, but comprehensive test suites that would assess the use of, for example, numerous syntactical rules can be arduous to synthesise. To sidestep the need to synthesise examples, a natural language dataset can be partitioned into training and test sets in such a way that the test set includes examples whose processing requires some systematic generalisation ability. A framework for partitioning data in this way was developed by Keysers et al. (2020), called *distribution-based compositionality assessment*, or DBCA for short. The main idea of DBCA is to control the distributions of *atoms* (primitive elements) and *compounds* (combinations of the atoms) to get approximately the same atom distributions but divergent compound distributions in the training and test sets.

We utilise the DBCA framework in our GenBench Collaborative Benchmarking Task submission, which consists of train-test splits of the Europarl parallel corpus with divergent distributions

Sentence	Atoms	Compounds
“Our vigilance is not partisan.”	nsubj, poss, our, vigilance, partisan	(vigilance, poss, our), (partisan, nsubj, vigilance)
“We shall now hear Mr Wurtz speaking against this request.”	hear, aux, shall, speak, nsubj, wurtz, hear, ccomp, speak	(hear, aux, shall), (speak, nsubj, wurtz), (hear, ccomp, speak)
“This seems to me to be a workable solution.”	solution, amod, workable, seem, xcomp, solution	(solution, amod, workable), (seem, xcomp, solution)

Table 1: Examples of what we call “atoms” and “compounds”. Atoms are the lemmas and dependency relations, and compounds the three-element tuples of the head lemma, the relation, and the dependant lemma.

of dependency relations. These data splits can be used to assess the ability of NMT systems to translate novel dependency relations. In the terminology of the DBCA framework, we define the atoms as lemmas and dependency relations, and the compounds as the three-element tuples of the head lemma, the dependant lemma, and the relation between them (see Table 1 for examples). This method to create compositional generalisation benchmarks does not require manual test suite construction, making it easy to extend it to other datasets and other languages.

2 Related work

2.1 Compositional generalisation in machine translation

Compositional generalisation has been assessed in machine translation in a few works in recent years. Raunak et al. (2019) partitioned a dataset into short training sentences and longer test sentences in order to assess generalisation from short sentences to longer ones, a subtype of compositional generalisation sometimes called *productivity* (Hupkes et al., 2020). Li et al. (2021) synthesised sentences for the test set with novel constituents, such as noun and verb phrases to create the CoGnition benchmark. Dankers et al. (2022) assessed three aspects of compositionality in NMT, which they called systematicity, the ability to combine familiar parts into novel combinations; substitutivity, the consistency of translations when a word is replaced with its synonym; and over-generalisation, the tendency to follow a compositional rule even when the case is actually an exception to the rule.

Perhaps the most similar benchmark to ours is ReaCT by Zheng and Lapata (2023). In ReaCT, the IWSLT 2014 German-English corpus is used as the training set, and a test set is created by selecting sentences that have a high *compositionality*

degree from the WMT 2014 corpus. The compositionality degree of a test set sentence is defined as the number of training set n-grams needed to create the sentence, divided by the length of the sentence. The reasoning is that a test set sentence that includes long n-grams from the training set can be composed of fewer n-grams, having a low compositionality degree, whereas if we need to back off to shorter n-grams to create a sentence, the compositionality degree is high. This train-test data split has similar general characteristics as our data splits: it is completely natural data (no synthetic sentences) partitioned so that the test set has novel combinations of familiar primitives. In contrast to our work, the primitives in this scheme can be sequences of multiple words, whereas we assess the ability to translate novel combinations of just two words and their dependency relation. We also pay attention to the relative frequency distributions of the primitives and the combinations, following the DBCA framework, whereas in ReaCT the compositionality score is a function of the unique n-gram types.

In addition to word-based compositionality, *morphological* compositional generalisation in translation has been assessed by Meyer and Buys (2023) and Moisiso et al. (2023). In these works, the atoms are defined as the morphemes (or surface-level morphs), and compounds as the inflected word forms that consist of multiple morphemes. The test set therefore includes novel combinations of familiar morphemes, assessing NMT systems’ capacity for morphological generalisation.

2.2 Other related work

Aside from NMT tasks, Shaw et al. (2021) utilised the DBCA framework to assess compositional generalisation in a natural language semantic parsing task. Sogaard et al. (2021) provide a more general discussion and review of non-random training-test

Motivation			
<i>Practical</i>	<i>Cognitive</i>	<i>Intrinsic</i>	<i>Fairness</i>
		□	
Generalisation type			
<i>Compositional</i>	<i>Structural</i>	<i>Cross Task</i>	<i>Cross Language</i>
□			<i>Cross Domain</i>
			<i>Robustness</i>
Shift type			
<i>Covariate</i>	<i>Label</i>	<i>Full</i>	<i>Assumed</i>
□			
Shift source			
<i>Naturally occurring</i>	<i>Partitioned natural</i>	<i>Generated shift</i>	<i>Fully generated</i>
	□		
Shift locus			
<i>Train-test</i>	<i>Finetune train-test</i>	<i>Pretrain-train</i>	<i>Pretrain-test</i>
□			

Figure 1: The categorisation of our benchmark in the taxonomy by Hupkes et al. (2023).

data splitting.

Besides creating artificial train-test splits, another option to test systematic generalisation in NLP systems, without the need for manual test suite design, is to leverage the fact that systematicity can be seen as an inherent symmetry in the data (Manino et al., 2022), which has also been utilised to generate new training examples (Akyurek and Andreas, 2023).

A study not on compositional generalisation, but in other ways related to our work, is that by McCoy et al. (2023), who evaluated the degree of novelty of the text generated by language models, using both n-grams and dependency relations. They found that language-model-generated text tends to be less novel than the baseline of human-generated text in local structure (small n-grams), but more novel than the human baseline in more global structure (large n-grams). This finding provides a good backdrop for our research question of how NMT models handle a test set that includes novel *local* structure, such as dependency relations.

2.3 The taxonomy

The eval card in Figure 1 shows how our benchmark can be categorised in the taxonomy of Hupkes et al. (2023). The motivation is primarily intrinsic: it is important to assess if translation models learn the systematic rules that characterise natural language, in order to get some understanding how the models work. Another motivation is practical; testing compositional generalisation is important for the practical reason of knowing how robustly the models generalise to novel dependency relations. The type of the generalisation is compositional, and

the shift type is covariate, since the input data distribution changes but the task remains otherwise the same. Shift source is partitioned natural data, since we do not use any artificial data, but the train-test split is artificial. Lastly, the shift locus in our experiments is train-test, but the method and benchmark could also possibly be used as a finetune train-test benchmark, by finetuning a pretrained model on the training set.

3 Europarl data splits

3.1 Data partitioning process

We use the Europarl corpus (Koehn, 2005) of transcribed European parliament proceedings, with the multilingual sentence alignments from the OPUS corpus (Tiedemann, 2012). We chose the Europarl corpus because of the good quality of the translations and because it includes parallel sentences for multiple languages. For our benchmark submission, we select English as the source language, and as the target languages we use four languages that represent different (branches of) language families: German, French, Greek, and Finnish.

As pre-processing, duplicate sentences are removed and maximum sentence length is restricted to 30 words before tokenisation. We take a random subsample of 300k sentences from which we extract the data splits. This relatively small size was chosen for convenient use as well as to allow comparison with previously published similar benchmarks: CoGnition Li et al. (2021) and ReaCT Zheng and Lapata (2023), which are similar in size.

The dependency parsing for the English source corpus is done using the LAL-parser (Mrini et al.,

2020). To calculate the divergences in data splitting, we only consider the English side. Therefore, the benchmark primarily assesses the encoder’s capacity to represent novel syntactic relations. However, presumably a high compound divergence of the source side sentences means that the target side sentences also include an increased number of novel syntactical, or possibly morphological, structures, assessing at the same time the decoder’s capacity to generate these. We define the atoms as the lemmas and dependency relations and compounds as the three-element-tuples of the dependant lemma, the head lemma and their relation. To make the number of atoms manageable, we exclude from the distribution calculations lemmas that appear either very frequently (the 200 most frequent lemmas, which each appear from 387k times (most frequent word “the”) to 3576 times (200th most frequent word “then”)), or fewer than 10 times in total in the corpus. After this filtering, about 8000 lemmas remain in the atom set, which includes also the dependency relation tags.

Following [Keysers et al. \(2020\)](#), we calculate a weight for each compound so that those sub-compounds that appear predominantly only in one super-compound get lower weight than those appearing in many different super-compounds. In our case, this means that if a (dependant lemma, relation type) pair occurs for example 8/10 times with just one head lemma, this pair gets a score of $1 - (8/10) = 0.2$. The idea is somewhat similar to that behind the Kneser-Ney smoothing ([Kneser and Ney, 1995](#)), where the number of different bigrams a word appears in correlates with the unigram back-off probability. We filter out compounds that get a weight less than 0.5, after which there are about 8000 atom types and 400k compound types whose distributions are controlled in the data splits.

Atom and compound divergences are calculated similarly to [Keysers et al. \(2020\)](#): divergence \mathcal{D} between distributions P and Q is calculated using the Chernoff coefficient $C_\alpha(P\|Q) = \sum_k p_k^\alpha q_k^{1-\alpha} \in [0, 1]$ ([Chung et al., 1989](#)), with $\alpha = 0.5$ for the atom divergence and $\alpha = 0.1$ for the compound divergence. [Keysers et al. \(2020\)](#) notes about the α values that $\alpha = 0.5$ for atom divergence “reflects the desire of making the atom distributions in train and test as similar as possible”, and $\alpha = 0.1$ for compound divergence “reflects the intuition that it is more important whether a certain compound occurs in P (train) than whether the probabilities

in P (train) and Q (test) match exactly”. The divergence is the complement of the Chernoff coefficient, since the Chernoff coefficient measures similarity between two vectors. The atom and compound divergences for training set V and test set W are:

$$\begin{aligned} \mathcal{D}_A(V\|W) &= 1 - C_{0.5}(\mathcal{F}_A(V) \parallel \mathcal{F}_A(W)) \\ \mathcal{D}_C(V\|W) &= 1 - C_{0.1}(\mathcal{F}_C(V) \parallel \mathcal{F}_C(W)). \end{aligned}$$

where \mathcal{F}_A is the atom distribution and \mathcal{F}_C is the compound distribution of a data set.

Splitting the data is done using a greedy algorithm similar to that by [Moisio et al. \(2023\)](#). This algorithm places one sentence at each iteration into either the training or test set, such that the atom and compound divergences are as close to the respective desired values as possible. Specifically, at each iteration we try to maximise a score that is the negated linear combination of the differences between the desired and actual divergence values:

$$\text{score}(V, W) = -|c - \mathcal{D}_C(V\|W)| - \mathcal{D}_A(V\|W),$$

where c is the desired compound divergence, and the desired atom divergence is 0 (minimum).

3.2 Comparison to random splits and previous benchmarks

Table 2 lists the sizes, and atom and compound divergences (\mathcal{D}_A and \mathcal{D}_C) for the Europarl data splits, as well as for random splits, and for the CoGnition and ReaCT benchmarks for comparison. All the divergences are calculated after similar filtering of the atoms and compounds as described for the Europarl splits above. From the divergences of the random train-test splits we can notice that the size of the test set correlates inversely with both atom and compound distributions; when the training and test sets are closer in size, the distributions are also naturally closer to each other.

From the table, we can also see that CoGnition includes relatively short sentences and, importantly, a relatively small number of unique lemmas. Creating CoGnition, [Li et al. \(2021\)](#) removed some of the complexity of natural language, such as polysemous words, and deliberately kept the vocabulary small and excluded all low-frequency words. This follows the design choices made by [Keysers et al. \(2020\)](#) of aiming to have only few meaningful atoms from which a large number of compounds can be created, which was motivated by practical concerns: this way it is easier to have a large range

	# sentences		# words		# unique lemmas	\mathcal{D}_A	\mathcal{D}_C
	train	test	train	test			
Europarl random split 3k	200k	3k	3.6M	54k	29k	0.28	0.60
Europarl random split 10k	200k	10k	3.6M	180k	30k	0.18	0.55
Europarl random split 30k	200k	30k	3.6M	540k	31k	0.13	0.52
CoGnition (Li et al., 2021)	196k	10k	1.9M	96k	1.7k	0.13	0.47
ReaCT (Zheng and Lapata, 2023)	160k	3k	3.3M	45k	53k	0.32	0.90
Europarl min \mathcal{D}_C split #1	203k	37k	3.9M	650k	34k	0.01	0.10
Europarl min \mathcal{D}_C split #2	194k	36k	3.7M	625k	34k	0.01	0.10
Europarl min \mathcal{D}_C split #3	195k	35k	3.7M	625k	34k	0.01	0.10
Europarl max \mathcal{D}_C split #1	197k	23k	3.8M	390k	34k	0.001	1.0
Europarl max \mathcal{D}_C split #2	198k	22k	3.8M	390k	34k	0.002	1.0
Europarl max \mathcal{D}_C split #3	198k	22k	3.8M	390k	34k	0.001	1.0

Table 2: Comparison of the Europarl splits to other translation benchmarks that aim at assessing compositional generalisation. The number of unique lemmas includes both training and test set lemmas.

of compound divergences while keeping the atom divergences same. However, this contrasts with the distribution of primitives in natural language, an issue we discuss more in Section 5. ReaCT, on the other hand, has similarly sized vocabulary as natural language data; in fact the vocabulary is significantly larger than the Europarl random sample vocabulary, possibly because the WMT and IWSLT corpora are lexically more diverse than Europarl.

Table 2 also lists, for reference, the atom and compound divergences, calculated for the dependency relations as described in Section 3, for CoGnition and ReaCT, even though neither of these data splits are designed to minimise or maximise these values. In both of these data sets, the test set is designed to contain novel combinations of familiar parts, as is our test sets, but in these works the parts are normally sequences of multiple words. In spite of this, the ReaCT data split has a relatively high dependency relation compound divergence.

The last rows of Table 2 show the sizes and divergences of the minimum- and maximum-compound-divergence Europarl data splits. The atom divergences are significantly lower for these splits than they are for the random splits. As noted above, the low atom divergence follows the principles of the DBCA framework: the compositional generalisation test set should be difficult not because of novel primitive elements (in our case mostly lemmas) but because of *novel combinations* of the known elements. The maximum-compound-divergence splits get a \mathcal{D}_C of exactly 1, but here we note that, as

described in Section 3, the most frequent and most infrequent lemmas are left out from the divergence calculations, which means that only a subset of the dependency relation compounds are novel in the test set, even though $\mathcal{D}_C = 1.0$.

4 Experiments

4.1 Transformer NMT system results

We train Transformer (Vaswani et al., 2017) models on the Europarl data splits using the OpenNMT python library (Klein et al., 2017) and the same hyperparameters as in the example OpenNMT-py Transformer configuration¹, including the standard 6 transformer layers with 8 heads, a hidden layer size of 512 and feed-forward layer size of 2048. The configuration includes around 60M parameters. We create a BPE (Sennrich et al., 2016) vocabulary for each language with 10k token types. We didn’t tune any of the hyperparameters for these datasets, and didn’t use a validation set. We trained the models for 6000 training steps, evaluated the test set translations at intervals of 1000 steps, and report the best of these 6 results for each model. For more details on training the models, see the Github repository linked on the first page.

Figure 2 displays the chrF2++ (Popović, 2017) scores for the minimum- and maximum-compound-divergence splits and the four target languages: German, French, Greek, and Finnish. We run the data split algorithm three times using different random

¹<https://github.com/OpenNMT/OpenNMT-py/blob/9d617b8b/config/config-transformer-base-1GPU.yml>




Figure 2: The chrF2++ scores for the Transformer NMT systems trained on the minimum and maximum compound divergence splits. Source language is English and target languages are German, French, Greek, and Finnish. The middle scores are for the random data splits, which happen to have a compound divergence of about 0.5. Each data split setup ($\min\mathcal{D}_C$, $\max\mathcal{D}_C$, and random) is run 3 times with different random seeds, creating 9 different data splits for which NMT models are trained.

seeds. Between the $\min\mathcal{D}_C$ and $\max\mathcal{D}_C$ split results are the results for the random splits, of which there are also 3 random runs (with 200k sentences in training and 30k sentences in test set). As shown in Table 2 the random splits with 30k-sentence test set happen to have a compound divergence of about 0.5.

Figure 2 shows a modest but statistically significant and consistent decrease in performance from the random data split to the $\max\mathcal{D}_C$ split, for all four target languages. This is expected, as the $\max\mathcal{D}_C$ split includes more novel dependency relations than the random split. However, as shown in Table 2, the $\max\mathcal{D}_C$ split has a significantly lower *atom* divergence than the random split, as this is deliberately minimised in the artificial data splits, which follows the principle of the DBCA framework of having the same *atom* distribution in training and test sets. This means that the $\max\mathcal{D}_C$ split should be easier than the random split in this regard, but it still gets worse results because of the high compound divergence.

The $\min\mathcal{D}_C$ split, on the other hand, has a similar *atom* divergence as the $\max\mathcal{D}_C$ split, so comparison between these two results is in that sense fairer. There is a larger difference in the results between these two data splits; depending on the target language the chrF2++ drops from about 4% to 8%. Since we use relatively large test corpora (from about 20k to 40k sentences), even small differences in chrF2++ are statistically significant.

4.2 Generalisation score

To assess whether one NMT system is more capable in (this dependency-relation-related type of) compositional generalisation than some other sys-

tem, one option is simply to compare their translation performances on the $\max\mathcal{D}_C$ split. However, to get a sense of the generalisation capacity as a part of the system’s translation capacity in general, it may be more meaningful to assess how the performance deteriorates between the $\min\mathcal{D}_C$ and $\max\mathcal{D}_C$ splits. To get a *generalisation score*, we propose to take the ratio of the results on these two data splits. This way the generalisation scores, using the average of the three chrF2++ scores for each experiment setup listed in Figure 2, are: $50.12/54.23 = 0.92$ for the German-target Transformer system, $57.13/60.11 = 0.95$ for French, $53.77/57.05 = 0.94$ for Greek, and $48.30/50.48 = 0.96$ for Finnish. Since this a relative score, the absolute chrF2++ results should be reported in addition to this generalisation score.

5 Discussion: handling the long tail

The core of each natural language is a set of words (a lexicon), and a set of grammar rules that define how to combine the words into meaningful sequences. The lexicon is divided into content words, which possess semantic content, and function words, which denote the grammatical relationships between content words. The function words belong to closed classes, for example prepositions, that normally do not accept new words, while content words belong to open classes; new nouns, for example, are coined regularly. The long tail of the Zipfian distribution of word frequencies consists of content words while the grammar-enforcing function words are mostly in the head of the distribution.

Recent studies suggest that the distinction between content and function words as well as the

Zipfian distribution are important for (compositional) generalisation to arise: [Steinert-Threlkeld \(2020\)](#) provides empirical evidence that function words enable robust forms of non-trivial compositional communication, and [Chan et al. \(2022\)](#) found that the Zipfian distribution helps language models strike a balance between memorisation and (in context) generalisation. As shown also by [Feldman \(2020\)](#), memorising some of the long tail is not in conflict with generalisation; on the contrary, it *enables* optimal generalisation.

At the same time, some studies have shown that neural NLP systems have some difficulties with handling the long tail, or at least handling the tail in a way that we as users of the models would expect and want. [Wei et al. \(2021\)](#) and [Czarnowska et al. \(2019\)](#) found that NLP systems’ performance is heavily influenced by word frequency in training. [LeBrun et al. \(2022\)](#) compared language-model-generated text to a reference and found that the models underestimate the long tail of well-formed sequences; furthermore, this probability mass didn’t go to the head of the distribution but rather the models overestimate the probability of ill-formed sequences.

As noted in previous sections, the desired distribution of atoms in the DBCA framework contrasts with the Zipfian distribution found in natural language. As [Keyzers et al. \(2020\)](#) explain, they design their CFQ benchmark “so as to have few and meaningful atoms” which means there is no long tail of infrequent primitives. This is related to having a close-to-zero atom divergence: if the atom distribution had a long tail, it would not be easy to have the same relative atom distributions in training and test sets since at least those atoms that appear only once would make the distributions diverge. This harks back to the question of how compositional generalisation could be assessed in purely natural tasks, where every rule has an exception, and where idioms and irregularities muddle the systematicity (as discussed by [Dankers et al. \(2022\)](#)).

Our benchmark provides one answer to this question. Although we use the principles of the DBCA framework regarding distribution divergence, we don’t make the corpus less natural by artificially shrinking the vocabulary size of the corpus. Instead, to make the divergence calculations manageable in practice, we leave some of the vocabulary out of the calculations. A downside of this choice is

that the test sentences in the $\max \mathcal{D}_C$ splits don’t contain exclusively novel dependencies. The advantage is that the vocabulary is similar to that in the original real-world natural language dataset, while the test set includes an increased number of novel dependencies to test generalisation.

6 Conclusion

Our GenBench Collaborative Benchmarking Task submission consists of train-test splits of the Europarl parallel corpus with divergent distributions of dependency relations. These data splits can be used to assess the ability of NMT systems to translate novel dependency relations in a purely natural language translation task. We derive the data partitioning method from the distribution-based compositionality assessment framework, which provides generalisable principles of how to assess compositionality. Our application of the DBCA framework is straightforward to extend further to new datasets and languages, and to any other NLP task where the training and test sets consist of sentences, such as paraphrase detection. The DBCA framework is useful for real-world natural language tasks too, even though it was originally designed for a more artificial data setting. It should be kept in mind, however, how the principles of the DBCA framework diverge from the reality of natural language data.

7 Limitations

7.1 Applicability of the method

Nowadays, the state-of-the-art methods in many NLP tasks are based on pretrained language models. However, our data partitioning method, as used in this work, requires controlling the training data, as well as the test data, to have partitions with specific atom and compound divergence values. Therefore, the method is not directly applicable to a pretrain-finetune training scheme, if the pretraining dataset is not modifiable. However, there are no limitations, in principle, on having a fixed training corpus and compiling only a new test set to have specific divergence values, although the divergence values might not be so easy to minimise and maximise in this case.

7.2 Validity of the method

We have not conclusively assessed whether the benchmark actually tests what we assume it tests, that is, compositional generalisation ability. To

rule out the most obvious potential confounding variable, we checked the sentence lengths to see if the $\max\mathcal{D}_C$ test sets for some reason included longer sentences, making the test set more difficult this way. We did not find large differences: in the $\min\mathcal{D}_C$ splits the average sentence lengths in train/test sets are 19.3/17.6 words, and in the $\max\mathcal{D}_C$ splits 19.1/17.3 words. Although there is a difference (of unknown origin) between train and test set sentence lengths, there is no significant difference between $\min\mathcal{D}_C$ and $\max\mathcal{D}_C$ splits that could confound the results. From Table 2 we can also see that the training sets are similar in size.

7.3 Limitations of the experiments

There are multiple levels of compositionality in language, from morphemes to words to phrases to clauses. Our experiments focus on just one intermediate level of compositionality, since we define compounds as dependencies between two words. This choice was based primarily on convenience: we could define compounds as constructions of more than just two words, but the large number of these constructions would make the data partitioning heavier computationally. Focusing on just one level of compositional constructions means that our experiments are not exhaustive in this regard, and we hope to assess other levels of compositionality in future work.

Our goal was to create a benchmark that tests generalisation to novel dependency relations in as comprehensively as possible, not selecting some specific types of dependency relations and leaving out other types. However, memory requirements of the data splitting algorithm do not permit us to use all of the atoms and compounds in the distribution divergence calculations, so we opted to leave out the most frequent and the most infrequent lemmas, and the dependency relations that include them. This means that our set of atoms represents a middle section of the distribution, where the head turns into a tail. Therefore, the controlled dependency relations include lemmas both from the head of the distribution and the tail, although neither of the extremes. We have not been able to assess how this particular choice affects the results.

8 Acknowledgements

We thank the anonymous reviewers for their insightful comments and feedback, which significantly improved this paper. The work was supported by the

Academy of Finland grant 337073. The computational resources were provided by Aalto ScienceIT.

References

- Ekin Akyurek and Jacob Andreas. 2023. [LexSym: Compositionality as lexical symmetry](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 639–657, Toronto, Canada. Association for Computational Linguistics.
- Stephanie Chan, Adam Santoro, Andrew Lampinen, Jane Wang, Aaditya Singh, Pierre Richemond, James McClelland, and Felix Hill. 2022. [Data distributional properties drive emergent in-context learning in transformers](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 18878–18891. Curran Associates, Inc.
- JK Chung, PL Kannappan, CT Ng, and PK Sahoo. 1989. [Measures of distance between probability distributions](#). *Journal of mathematical analysis and applications*, 138(1):280–292.
- Paula Czarnowska, Sebastian Ruder, Edouard Grave, Ryan Cotterell, and Ann Copestake. 2019. [Don’t forget the long tail! A comprehensive analysis of morphological generalization in bilingual lexicon induction](#). 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 974–983, Hong Kong, China. Association for Computational Linguistics.
- Verna Dankers, Elia Bruni, and Dieuwke Hupkes. 2022. [The paradox of the compositionality of natural language: A neural machine translation case study](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4154–4175, Dublin, Ireland. Association for Computational Linguistics.
- Vitaly Feldman. 2020. [Does learning require memorization? A short tale about a long tail](#). In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 954–959.
- Jerry A Fodor and Zenon W Pylyshyn. 1988. [Connectionism and cognitive architecture: A critical analysis](#). *Cognition*, 28(1-2):3–71.
- Dieuwke Hupkes, Verna Dankers, Mathijs Mul, and Elia Bruni. 2020. [Compositionality decomposed: How do neural networks generalise?](#) *Journal of Artificial Intelligence Research*, 67:757–795.
- Dieuwke Hupkes, Mario Giulianelli, Verna Dankers, Mikel Artetxe, Yanai Elazar, Tiago Pimentel, Christos Christodoulopoulos, Karim Lasri, Naomi Saphra, Arabella Sinclair, Dennis Ulmer, Florian Schottmann, Khuyagbaatar Batsuren, Kaiser Sun, Koustuv Sinha, Leila Khalatbari, Maria Ryskina, Rita Frieske, Ryan

- Cotterell, and Zhijing Jin. 2023. [A taxonomy and review of generalization research in NLP](#). *Nature Machine Intelligence*, 5(10):1161–1174.
- Daniel Keysers, Nathanael Schärli, Nathan Scales, Hylke Buisman, Daniel Furrer, Sergii Kashubin, Nikola Momchev, Danila Sinopalnikov, Lukasz Stafiniak, Tibor Tihon, Dmitry Tsarkov, Xiao Wang, Marc van Zee, and Olivier Bousquet. 2020. [Measuring compositional generalization: A comprehensive method on realistic data](#). In *International Conference on Learning Representations*.
- Najoung Kim and Tal Linzen. 2020. [COGS: A compositional generalization challenge based on semantic interpretation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9087–9105, Online. Association for Computational Linguistics.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander Rush. 2017. [OpenNMT: Open-source toolkit for neural machine translation](#). In *Proceedings of ACL 2017, System Demonstrations*, pages 67–72, Vancouver, Canada. Association for Computational Linguistics.
- R Kneser and H Ney. 1995. [Improved backing-off for m-gram language modeling](#). In *Acoustics, Speech, and Signal Processing, IEEE International Conference on*, volume 1, pages 181–184. IEEE Computer Society.
- Philipp Koehn. 2005. [Europarl: A parallel corpus for statistical machine translation](#). In *Proceedings of Machine Translation Summit X: Papers*, pages 79–86, Phuket, Thailand.
- Brenden Lake and Marco Baroni. 2018. [Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks](#). In *International Conference on Machine Learning*, pages 2873–2882. PMLR.
- Benjamin LeBrun, Alessandro Sordoni, and Timothy J O’Donnell. 2022. [Evaluating distributional distortion in neural language modeling](#). In *International Conference on Learning Representations*.
- Yafu Li, Yongjing Yin, Yulong Chen, and Yue Zhang. 2021. [On compositional generalization of neural machine translation](#). 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 4767–4780, Online. Association for Computational Linguistics.
- Edoardo Manino, Julia Rozanova, Danilo Carvalho, Andre Freitas, and Lucas Cordeiro. 2022. [Systematicity, compositionality and transitivity of deep NLP models: a metamorphic testing perspective](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2355–2366, Dublin, Ireland. Association for Computational Linguistics.
- R. Thomas McCoy, Paul Smolensky, Tal Linzen, Jianfeng Gao, and Asli Celikyilmaz. 2023. [How much do language models copy from their training data? evaluating linguistic novelty in text generation using RAVEN](#). *Transactions of the Association for Computational Linguistics*, 11:652–670.
- Francois Meyer and Jan Buys. 2023. [Subword segmental machine translation: Unifying segmentation and target sentence generation](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 2795–2809, Toronto, Canada. Association for Computational Linguistics.
- Anssi Moio, Mathias Creutz, and Mikko Kurimo. 2023. [Evaluating morphological generalisation in machine translation by distribution-based compositionality assessment](#). In *Proceedings of the 24th Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 738–751, Tórshavn, Faroe Islands. University of Tartu Library.
- Khalil Mrini, Franck Dernoncourt, Quan Hung Tran, Trung Bui, Walter Chang, and Ndapa Nakashole. 2020. [Rethinking self-attention: Towards interpretability in neural parsing](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 731–742, Online. Association for Computational Linguistics.
- Barbara Partee. 1995. [Lexical semantics and compositionality](#). *An invitation to cognitive science*, 1:311–360.
- Ellie Pavlick. 2022. [No one metric is enough! combining evaluation techniques to uncover latent structure](#). In *Workshop on The Challenge of Compositionality for AI*.
- Ellie Pavlick. 2023. [Symbols and grounding in large language models](#). *Philosophical Transactions of the Royal Society A*, 381(2251):20220041.
- Maja Popović. 2017. [chrF++: words helping character n-grams](#). In *Proceedings of the Second Conference on Machine Translation*, pages 612–618, Copenhagen, Denmark. Association for Computational Linguistics.
- Vikas Raunak, Vaibhav Kumar, and Florian Metze. 2019. [On compositionality in neural machine translation](#). In *NeurIPS Workshop on Context and Compositionality in Biological and Artificial Neural Systems*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Peter Shaw, Ming-Wei Chang, Panupong Pasupat, and Kristina Toutanova. 2021. [Compositional generalization and natural language variation: Can a semantic parsing approach handle both?](#) 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 922–938, Online. Association for Computational Linguistics.

Anders Søgaard, Sebastian Ebert, Jasmijn Bastings, and Katja Filippova. 2021. [We need to talk about random splits](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1823–1832, Online. Association for Computational Linguistics.

Shane Steinert-Threlkeld. 2020. [Toward the emergence of nontrivial compositionality](#). *Philosophy of Science*, 87(5):897–909.

Jörg Tiedemann. 2012. [Parallel data, tools and interfaces in OPUS](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 2214–2218, Istanbul, Turkey. European Language Resources Association (ELRA).

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. Curran Associates, Inc.

Jason Wei, Dan Garrette, Tal Linzen, and Ellie Pavlick. 2021. [Frequency effects on syntactic rule learning in transformers](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 932–948, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Hao Zheng and Mirella Lapata. 2023. [Real-world compositional generalization with disentangled sequence-to-sequence learning](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1711–1725, Toronto, Canada. Association for Computational Linguistics.