

Towards Understanding the Relationship between In-context Learning and Compositional Generalization

Sungjun Han*, Sebastian Padó†

*42dot Inc., sungjun.han@42dot.ai

†IMS, University of Stuttgart, pado@ims.uni-stuttgart.de

Abstract

According to the principle of *compositional generalization*, the meaning of a complex expression can be understood as a function of the meaning of its parts and of how they are combined. This principle is crucial for human language processing and also, arguably, for NLP models in the face of out-of-distribution data. However, many neural network models, including Transformers, have been shown to struggle with compositional generalization. In this paper, we hypothesize that forcing models to *in-context learn* can provide an inductive bias to promote compositional generalization. To test this hypothesis, we train a causal Transformer in a setting that renders 'ordinary' learning very difficult: we present it with different orderings of the training instance and shuffle instance labels. This corresponds to training the model on all possible few-shot learning problems attainable from the dataset. The model can solve the task, however, by utilizing earlier examples to generalize to later ones – i.e., in-context learning. In evaluations on the datasets, SCAN, COGS, and GeoQuery, models trained in this manner indeed show improved compositional generalization. This indicates the usefulness of in-context learning problems as an inductive bias for generalization.

Keywords: neural language representation models, statistics and machine learning methods, semantics

1. Introduction

As humans, we have the ability to combine atomic parts in reoccurring structures in novel manners (Fodor and Pylyshyn, 1988). This ability, known as *compositional generalization*, is an important aspect of human language processing, affording us with an "infinite use of finite means" (Chomsky, 1965). For example, when we understand the meaning of the predicate *dax* in phrases such as "I can dax" and "dax twice", we can also understand phrases such as "dax voluntarily" or "must dax".

In contrast, many modern deep neural architectures struggle with compositional generalization (Baroni, 2020; Lake and Baroni, 2017; Hupkes et al., 2020; Kim and Linzen, 2020; Keysers et al., 2020). While they excel at making predictions for test sets similarly distributed to the training set (i.e., *in-distribution*), their performances significantly decrease when generalizing to test distributions that are differently structured (i.e., *out-of-distribution*) even if they contain the same set of atoms.

We believe that standard models lack an inductive bias towards acquiring compositional representation, which arises from the independent parallel processing of examples in mini-batches. In most mini-batches, the models do not have *explicit* access to a sufficient number of instances of the atoms to make it worthwhile to learn compositionally generalizable representations for the atoms. Contrast this with symbolic accounts of compositional generalization, e.g., in the shape of *case-based reasoning* (Leake, 1996), where prediction can always rely on the availability of a sufficient number of relevant examples in memory. Thus, the

ability to understand "dax thrice" from "dax twice" can be thought of as a generalization of relevant past uses of "thrice" in memory, such as "eat thrice", combined with the use of "dax" in "dax twice".

If this is true, then compositional generalization should be encouraged by forcing models to *in-context learn* (Brown et al., 2020; Chowdhery et al., 2022) – that is, forcing them to generalize to new examples conditioned on a few demonstrations of input-output mappings provided in the model's context (or memory) without parameter updates. In-context learning forces the model to compute in the forward pass how the past examples provided the context can be utilized in a novel manner for the later examples. We observe that it is the same mechanism that supports the learning of compositionally generalizable input-output mappings.

The intuition behind our study is aligned with theoretical studies that explain in-context learning (Ortega et al., 2019; Xie et al., 2022) as an implicit Bayesian inference, where the model learns to approximate the latent parameters. However, it is yet unclear empirically how compositional generalization and in-context learning are related. On the one hand, the reported improvement in compositional generalization for the large Transformer-based LLMs (Zhou et al., 2023; Hosseini et al., 2022) with *emergent* in-context learning ability seem to point to an underlying relationship. On the other hand, our understanding is limited by (a) the lack of control over the training data in these studies and (b) the uncertainty regarding how much of the inductive biases implicit in the prompting methods contribute to the improvement. Indeed, Hosseini et al. (2022) and Qiu et al. (2022b) report that only

some in-context learning LLMs can compositionally generalize and only as they scale up.

As implementation, we utilize a *meta-learning* (Schmidhuber et al., 1996; Bengio et al., 1991; Hochreiter et al., 2001; Duan et al., 2017; Ortega et al., 2019) regime to explicitly incentivize in-context learning for a causal Transformer (Vaswani et al., 2017; Radford et al., 2019) with a language modelling objective. We train *from scratch* to eliminate the possible confounders introduced from pre-training in studying the relationship. Each task of our *meta-task* distribution is one possible linear ordering of input-output pairs of the training dataset formed into a single sequence via concatenation. This trains the model on all possible few-shot in-context learning problems attainable from the dataset. In order to discourage the model from relying on memorization, we also shuffle the labels. At prediction time, we condition the inference on the test examples on randomly sampled training mappings, maintaining the *zero-shot* prediction setting. We evaluate our approach on three widely used datasets targeting specifically compositional generalization, namely SCAN (Lake and Baroni, 2017; Keysers et al., 2020), COGS (Kim and Linzen, 2020), and GeoQuery (Zelle and Mooney, 1996; Shaw et al., 2021). Our contributions are:

1. We empirically study the relationship between in-context learning and compositional generalization through a novel meta-learning training regime that incentivizes in-context learning on established compositional generalization datasets along with a corresponding evaluation regime that maintains a zero-shot prediction setting.
2. We show that a causal Transformer trained through meta-in-context learning from scratch without any pretraining exhibits a significant improvement in performance on compositional generalization compared to the models without meta-learning.
3. We demonstrate several connections between in-context learning and compositional generalization through ablations: More in-context learning problems lead to better compositional generalization (Exp. 2); trained models are indeed generalizing through in-context learning in informative contexts (Exp. 3); the success of in-context learning depends on the absence of memorization (Exp. 4); pre-trained models have a better prior for in-context learning and can also benefit from meta-learning (Exp. 5).

Plan of the paper §2 introduces important background concepts and reviews notable related works. §3 presents our meta-learning regime in detail. §4

provides information on experimental setup, followed by the experimental results in §5. §6 concludes the paper along with future directions.

2. Related Work

2.1. Compositional Generalization

The difficulties of neural networks in compositional generalization have been identified by many studies. In the following, we focus on studies on unimodal language data.

Notable text-to-text benchmarks include SCAN (Lake and Baroni, 2017), PCFG (Hupkes et al., 2020), COGS (Kim and Linzen, 2020), and CFQ (Keysers et al., 2020). These datasets generate a single data distribution which are split into a train and test set in a systematic manner, attempting to capture the notion of *systematicity* or/and *productivity* (Fodor and Pylyshyn, 1988; Hupkes et al., 2020). The former refers to the ability to recombine parts in a novel manner, and the latter to the ability to recursively combine known structures. Hence, if a model learning from the train set can find a compositional solution, it can be successful on the test as the same data generative process underlie the two.

Many studies have proposed different inductive biases to promote compositionality. They include new deep learning architectures structurally constraining the ways the inputs are processed and represented (Li et al., 2019; Russin et al., 2019; Gordon et al., 2020; Bergen et al., 2021), providing additional supervisory signals (Jiang and Bansal, 2021), data augmentation (Andreas, 2020; Guo et al., 2020b; Akyürek et al., 2021; Qiu et al., 2022a; Li et al., 2023), and hybrid symbolic reasoning approaches (Nye et al., 2020; Liu et al., 2020b; Guo et al., 2020a). These approaches have shown to improve compositional generalization. However, they often require prior knowledge of the dataset, and their scalability to bigger and more general datasets is uncertain.

Following these concerns, some studies have constrained their investigations to the popular neural sequences model such as Transformers (Ontanon et al., 2022; Csordás et al., 2021), finding that their compositional generalization capacity can be improved with the available variants (e.g. relative positional encoding (Dai et al., 2019) or tying the layers (Dehghani et al., 2019)). Patel et al. (2022) showed that popular architectures including Transformers can be improved by increasing diversity in the data distribution. Our work follows this line of research by studying how a better inductive bias can be provided without a major change in the Transformer architecture.

2.2. Meta-learning

Meta-learning (Bengio et al., 1991; Schmidhuber et al., 1996) aims to enable machine learning models to learn how to learn by exposing them to a *distribution of tasks* where a model can improve from experience. The tasks are selected to be similarly structured but differ in details such that it is profitable for the model to find a generalizable solution rather than memorizing individual answers. Our work follows the line of work known as *memory-based meta-learning* or *meta-in-context learning* (Hochreiter et al., 2001; Santoro et al., 2016; Duan et al., 2017; Wang et al., 2017; Ortega et al., 2019), which incentivizes the model to learn to in-context learn by training on a task distribution of *sequences* of input-output mappings.

Meta-learning was applied to various tasks in language processing such as cross-lingual transfer (Gu et al., 2018), question answering (Nooralahzadeh et al., 2020), and domain adaption (Qian and Yu, 2019). However, it has rarely found application in semantic processing. The challenge arises from the difficulty of not knowing beforehand the relevance of specific examples, which makes it difficult to construct the task distribution with the right inductive bias for compositional generalization. Lake (2019) evaded this problem by using the ground truth grammar of the data distribution. This allowed them to permute only the input-output mappings of the primitives, which was shown to improve compositional generalization. Conklin et al. (2021) used model-agnostic meta learning (MAML) (Finn et al., 2017) as an auxiliary loss for supervised learning. In this approach, a single gradient step is taken on one set of support examples and the auxiliary loss is accrued by how well the updated model performs on another structurally similar set. This loss is back-propagated through the gradient optimization step all the way back to the model weights. The proposed method alleviated the problem of selecting support examples during evaluation, but the approach still relied on ground truth structural knowledge.

2.3. In-context Learning

A long line of work attempts to understand the property of in-context learning, especially related to their ability to generalize to out-of-distribution data. A number of studies has shown that in-context learning in LLMs can be utilized for compositional generalization using specific prompting methods (Zhou et al., 2023; Wei et al., 2022; Fu et al., 2023) or by scaling up the model (Hosseini et al., 2022; Qiu et al., 2022b). One recent work by An et al. (2023) has also investigated how the in-context learning ability of LLMs can be improved by selecting better demonstrations with relevant linguistic structure.

As explained above, the in-context learning ability in these models was also analyzed theoretically, and the driving force was found to be latent text properties that heavily affects token distributions (Xie et al., 2022).

Some works have studied the effects of further meta-training LLMs for in-context learning (Chen et al., 2022; Min et al., 2022) and found the tuned models to perform better than the base LLM. Our work is closer to the studies that train Transformers from *scratch* instead of looking at LLMs. Chan et al. (2022) showed that the emergence of in-context learning to depend on the informativeness of contexts. Garg et al. (2022) showed that Transformers are able to in-context learn simple functions and generalize to out-of-distribution samples, and Kirsch et al. (2022) extended its study to in-context learning arbitrary image-label mappings.

3. Methods

We now introduce a meta-learning regime that can be generally applied to a sequence to sequence dataset consisting of input-output sequence pairs. The main goal is how to construct a meta task-distribution with the right inductive bias for compositional generalization. The key idea is the inductive bias created by *online* learning an entire dataset: The model observes each example in the dataset only once and sequentially one after the other. When learning on such a linear ordering of examples (i.e., trajectory), the model cannot memorize and needs to successfully store and represent the past examples to generalize to the future examples.

Since there is no inherent order between the examples in a sequence to sequence dataset, different linear orderings of the dataset pose different generalization problems for a model. However, no matter which ordering we choose, the structure behind each trajectory remains invariant as it is governed by the same set of latent parameters. Hence, when meta-learning on such a task distribution, a model has a chance of approximating the underlying structure of the dataset. Note that this way of constructing the task distribution do not require any prior knowledge of the dataset, in contrast to the earlier approaches (Lake and Baroni, 2017; Conklin et al., 2021).

3.1. Meta-training

Given a sequence to sequence dataset $D = \{(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})\}_{i=1}^N$ with a vocabulary V , we form the task distribution $P(\tau)$ for meta-learning, where each task τ is one possible linear ordering of the dataset $(\mathbf{x}^{(1)}, \mathbf{y}^{(1)}), \dots, (\mathbf{x}^{(N)}, \mathbf{y}^{(N)})$. We feed this to the model as a concatenation $\tau =$

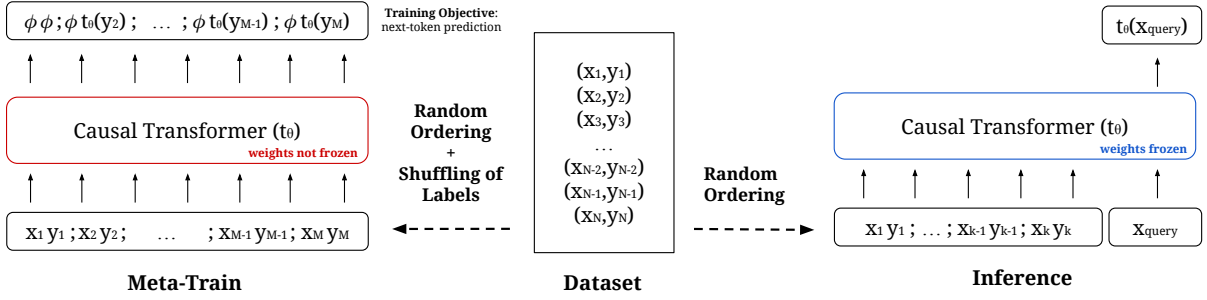


Figure 1: Illustration of our meta-in-context learning framework. (Left) We build our meta-task distribution by sampling random linear orderings from a sequence to sequence dataset and concatenating the input-output mappings (i.e., (x_i, y_i)). We optionally shuffle the labels to eliminate memorization and keep only M examples. A causal Transformer (t_θ) is trained with these concatenated results for next-token prediction, only predicting for the outputs. ϕ refers to the pad-token. (Right) At inference, we freeze the weights and randomly sample $k < M$ train examples to use as a context in predicting the test query x_{query} .

$[\mathbf{x}^{(1)}; \mathbf{y}^{(1)}; \dots; \mathbf{x}^{(N)}; \mathbf{y}^{(N)}]$ using two delimiter tokens, one to distinguish the inputs from the outputs and the other to separate the sequence elements. We assume a uniform distribution for $P(\tau)$.

A limitation of this approach is the possibility of memorization as each example occurs many times *across* different trajectories, although it occurs only once *within* each trajectory. Hence, a model might learn to ignore the context and memorize the examples, which is especially true for small datasets. To counteract this danger, we randomly shuffle the labels of the vocabulary V . For example, given a dataset $\{(jump, J), (run\ twice, R\ R)\}$, we can create an alternate version of the dataset $\{(jump, R), (run\ twice, J\ J)\}$ by the shuffling all instances of J with R . This results in more than one output label to be assigned to each input token throughout training and can prevent memorization.

Formally, we train a model M_θ given a linear ordering sequence of tokens upto and including the i -th token $\tau[i]$ to predict the next token $f(\tau[i]) = \tau[i + 1]$ if the i -th token belongs to the output and a pad token $f(\tau[i]) = \phi$ if it belongs to the input or the first output of the sequence. The objective is to minimize the expected loss over all possible orderings, optionally shuffling the labels¹:

$$\min_{\theta} \mathbb{E}_{\tau \sim P(\tau)} \left[\sum_{i=1}^{|\tau|} \frac{1}{|\tau|} \ell(M_\theta(\tau[i]), f(\tau[i])) \right] \quad (1)$$

where $\ell(\cdot, \cdot)$ is the cross-entropy loss function. Figure 1 (left half) illustrates the training procedure.

This objective can be interpreted as training a model on all possible few-shot learning problems attainable from the dataset. Hence, any specific way of defining the meta task-distribution is a subset of our distribution. The strength of the injected inductive bias for compositional generalization is limited

¹For clarity, we do not formalize the label shuffling in the following equation.

to the kinds of generalization problems inherent in each dataset. A final practical problem is that for most datasets, the set of all shuffled variants does not fit into GPU memory. Hence, we fix a certain roll-out length $M < N$ to limit each sequence τ to consist of M input-output pairs. Note that as we make M smaller, the number of distinct tasks in the task distribution decreases. We investigate the impact of M in Exp. 2 below.

Underlying Neural Network. This meta-training can be applied to any neural network model with memory. However, the use of an autoregressive model is very advantageous: In such a model, a single trajectory consisting of N concatenated input-output mappings can be combined with causal masking to yield $k - 1$ few-shot learning problems in one go. In a bidirectional model, in contrast, one needs to provide $k - 1$ different problems separately. Hence, we adopt the causal Transformer (Vaswani et al., 2017; Radford et al., 2019).

3.2. Inference

Compositional generalization datasets are designed as a zero-shot generalization task. This means that a model is required to generalize to the test examples only by using the train examples. We randomly sample a *training* trajectory of length $k < M$ to condition the inference for each test input x_q . Although we cannot guarantee the relevance of every sample, the model can still *choose* among these samples through the attention mechanism, analogously to case-based reasoning (Leake, 1996). See Figure 1 (right half) for the illustration of our inference method. We investigate the implications of the choice of k in Exp. 3 below.

4. Experimental Setup

4.1. Datasets

SCAN (Lake and Baroni, 2017) consists of natural language commands that need to be mapped to sequences of actions (e.g. jump twice \rightarrow JUMP JUMP). Among various compositional generalization *splits* of SCAN, we use the Maximum Compound Divergence (MCD) splits introduced by Keysers et al. (2020). These splits capture a general notion of compositional generalization by capturing both systematicity and productivity. They maximize the divergence between the compounds while maintaining the closeness of the atom frequency distribution. There are three SCAN-MCD splits with increasing difficulty (i.e., MCD1 being the easiest), each with 8,365 train and 1,045 test examples.

COGS (Kim and Linzen, 2020) is a semantic parsing dataset with a diverse set of natural language sentences. The compositional generalization split called "Gen(eralization)" was constructed based on various kinds of linguistic generalizations found in English, such as generalizing the subject role to the object role (systematicity) or generalizing to sentences with more depth (productivity). The training set consists of 24,155 examples while 21,000 examples make up the test.

GeoQuery / GEO (Zelle and Mooney, 1996; Dong and Lapata, 2016) is a semantic parsing dataset consisting of natural language database queries. We use the TMCD compositional split (Shaw et al., 2021) which adapts the MCD principle to a non-synthetic dataset. It is a fairly small dataset consisting of 440 examples for both train and test.

4.2. Preprocessing

For SCAN and COGS, we preprocess the output sequences to reduce their lengths in order to be able to fit longer trajectories into memory (i.e., increase M). For SCAN, we represent the action sequences in Python syntax as was done in Zhou et al. (2023) for evaluating LLMs. For example, "LOOK LOOK" is represented as LOOK * 2. For COGS, we omit the brackets and represent the variables x_n as n . Both are intermediate representations that can be fully mapped back to the original form. For GEO, we follow the same preprocessing step introduced in the original paper for TMCD (Shaw et al., 2021) which replaces the entities with a placeholder. This further brings down the unique number of train examples to 262. See Appendix C for more details.

4.3. Model Configuration and Training

We use an 8-layer 8-head causal Transformer with a model dimension of 512 and a feedforward dimension of 2,048 with absolute sinusoidal positional encoding (Vaswani et al., 2017), using the basic implementation available in PyTorch (Paszke et al., 2019). We do not use any pre-trained weights, and initialize the model and the word embeddings from scratch. Appendices A and B provide details on the used hyperparameters and method of checkpoint selection.

For SCAN, we study the effect of both applying and not applying shuffling of output labels, as its relatively small vocabulary ($|V| = 30$) affords us a full coverage of all words with a few samples. In contrast, COGS and GeoQuery have a much bigger vocabulary ($|V| = 871$ and $|V| = 154$, respectively). Therefore, we do not consider shuffling.

4.4. Evaluation

For evaluation, we report *sequence-level* accuracy, where a sequence is only deemed correct if it is predicted completely correctly. For each accuracy result, we also report the number of randomly sampled training examples k used for testing. If not mentioned, we set k to be one less than the maximum roll-out length M . All results report averages over five training runs for SCAN and GEO and three for COGS (for computational reasons).

4.5. Baselines and Points of Comparison

Herzig et al. (2021) showed that intermediate representations can lead to an improved compositional generalization. This especially applies to SCAN. Hence, we additionally train a 3-layer encoder-decoder Transformer (Vaswani et al., 2017) and Universal Transformer (Dehghani et al., 2019) with absolute positional encoding for SCAN. We also train the same baselines for GEO, as they are not reported in the literature. As for COGS, we find the impact of preprocessing on COGS to be minimal and different than the format (Ontanon et al., 2022) which appears to be optimal (see Appendix C for details). Hence, for COGS, we report the Transformer and Universal Transformer results from the literature.

For all datasets, we also report on a causal Transformer baseline trained using standard supervised learning, which is equivalent to training with the roll-out length of 1 (i.e., $M = 1, k = 0$). Finally, we compare our approach with the prior meta-learning work: the MAML-augmented Transformer with Tree-based search for COGS and string-based for SCAN (Conklin et al., 2021).

Method	Bidir	IntRep	SCAN			COGS	GEO
			MCD1	MCD2	MCD3	Gen	TMCD
Transformer (lit.)	+	-	0.4±0 [1]	1.8±0 [1]	0.5±0 [1]	35±6 [2]	NA
Transformer (ours)	+	+	41.7±4	20.3±5	17.1±6	80±0 [3]	36.8±1
Universal Transformer (ours)	+	+	36.4±9	34.1±6	25.5±10	78±0 [3]	37.3±1
Transformer + MAML (lit.)	+	-	2.6±0 [4]	5.6±1 [4]	6.7±1 [4]	66.7±4 [4]	NA
C-Transformer (ours)	-	+	21.8±3	25.6±2	19.7±2	51.9±4	37.4±1
C-Transformer + meta-ICL	-	+	60.4±13	53.3±2	50.7±7	75.7±1	40.8±1
C-Transformer + meta-ICL + LB	-	+	71.2±7	74.8±9	38.7±8	NA	NA

Table 1: Exp. 1: Mean sequence-level accuracies and standard deviations across runs. For our meta-ICL C(ausal)-Transformer, we present the results from best M with $k = M - 1$. "Bidir" stands for bidirectional (vs. causal). "IntRep" indicates the use of the optimized intermediate representation for SCAN. Best model on each dataset boldfaced. "ours": own experimental results, "lit.": results from literature. "LB": label-shuffling. References: [1] Furrer et al. (2021), [2] Kim and Linzen (2020), [3] Csordás et al. (2021), [4] Conklin et al. (2021).

5. Experiments and Results

We now present the results of five experiments to better understand the relationship between in-context learning and compositional generalization: (1) We compare the performance of our models using $k = M - 1$ supports for evaluation with the baselines. (2) We test the effect of training the model with longer trajectories (i.e., bigger M) which is equivalent to training with a larger number of unique in-context learning problems. (3) We test the effect of varying the number of support *training* examples used during evaluation (i.e., varying k) to test whether the models are generalizing through in-context learning. (4) We test how general the in-context learning extracts the latent parameters by testing its ability to learn from a new distribution, providing the model with *test* examples. (5) We test whether a pre-trained model can also benefit from additional training with meta-in-context learning.

5.1. Exp. 1: Main Results

Table 1 summarizes the main results of a causal Transformer trained from scratch using our meta-training method, along with the baselines.

We first make the observation that our intermediate representation for SCAN leads to an improvement in compositional generalization (compare rows 1 and 2). Although the improvement is substantial, the datasets still remain difficult for the models and the relative difference of difficulty between the MCD splits are retained (see the decreasing performance for SCAN in columns 4-6). We also note that our causal Transformer baseline (row 5) performs mostly worse than the encoder-decoder counterparts, probably due to its unidirectionality. For all three MCD splits, our causal Transformer trained with meta-in-context learning

(row 6 and 7) substantially outperforms all other approaches with or with label shuffling. For COGS, it beats all models except the Transformer model of Csordás et al. (2021). For GEO, it seems to provide only a small gain in performance, probably due to its small size which makes it easier to memorize. Next, label shuffling seems to provide a positive boost for the first two splits of SCAN compared to their counterparts.

It is interesting to note that for the models without label shuffling, the improvement over the causal baseline simply comes from how the data was presented to the model. While it is unclear whether the improvement can be attributed to their in-context learning ability (explored in Exp. 3), we believe that training on linear orderings of examples at least resulted in some form of regularization, where the pressure to learn representations not only for the prediction but also for their use in the future contributed to the improvement. Since we did not use any informed strategy, such as a retriever, to construct the meta-training distribution, we see these results as evidence for a strong and general inductive bias in-context learning can provide for compositional generalization.

For the ablation studies (Exp. 2–5), we use only SCAN and COGS since GEO is too small for stable experimentation.

5.2. Exp. 2: More Learning Problems

Setup. Next, we investigate how the performance of our model changes when training on trajectories of different lengths. Constructing the task distribution with longer trajectories gives us more unique few-shot learning problems, which we expect to lead to better compositional generalization. However, longer trajectories can also lead to more overfitting, as the model needs to extrapolate less given

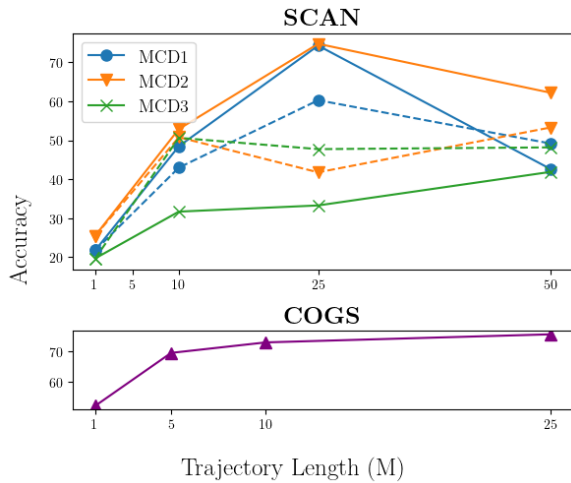


Figure 2: Exp. 2: Models trained on different lengths of trajectories (i.e., M), with $k = M - 1$. $M = 1$ is equivalent to the causal Transformer baseline. Dotted lines: models without label shuffling.

more support samples. Hence, we train three different values of $M = \{10, 25, 50\}$ for SCAN and $\{5, 10, 25\}$ for COGS, evaluating with $k = M - 1$.

Result. The results in Figure 2 show that an improvement for MCD1 and MCD2 of SCAN from increasing the trajectory length from 10 to 25, but we see signs of overfitting as we increase further, for the label shuffled models. We see a slightly different trend for the non-label-shuffled models for these splits, probably due to memorization, but still $M = 25$ and 50 perform better than 10. For MCD3, we see a monotonic improvement as we increase the trajectory length for both kinds. For COGS, we see a similar trend with improving performance as the length of the trajectories are increased. The gain diminishes after a certain point because the task becomes easier for the model. Where this turning point occurs depends on the available few-shot learning problems implicit in each dataset; this could be tuned using a standard hyperparameter search. In sum, the in-context learning ability of Transformers is sensitive to the kinds and number of few-shot generalization problems that it is exposed to during training, and having to solve more unique in-context learning problems can lead better compositional generalization.

5.3. Exp. 3: Number of Demonstrations

Setup. So far, we have only used $k = M - 1$ support examples for evaluation. In this experiment, we investigate how the model generalizes for different number of support examples. If the models are truly generalizing through in-context learning, then it should perform better with more demonstrations




Figure 3: Exp. 3: Models evaluated on different numbers of support examples k . Lines differ in M (max. roll-out length of meta-training trajectories). Dotted lines: models without label shuffling.

(i.e., larger k). If this is the case, then we also wish to rule out whether the improvement from Exp. 2 was simply due to the models with higher M having access to more support examples during evaluation by comparing with the same k . To do so, we evaluate the SCAN-MCD models using different values of $k = \{1, 3, 5, 9, 12, 16, 24, 49\}$ using the full test set. For COGS, we take 20% of the test set and vary $k = \{1, 4, 9, 24\}$ ². Note that we only evaluate when k does not exceed the maximum roll-out length for each given model (i.e., $k < M$).

Result. Figure 3 shows the results. For SCAN, all models improve as it receives more and more support examples, even for the models trained without label shuffling. The result also confirms the conclusion from Exp. 2: For all MCD splits of SCAN, even when the model is given the same number of support examples, the best model performs better than the rest (i.e., 25 for MCD1 and MCD2 and 50 for MCD3 with label shuffling), suggesting that these models did learn more generalizable in-context learning.

²The reason is that the test set of COGS is 21 times larger than SCAN. We found the results on our sample to be representative for evaluating on the entire test set.

Method		SCAN			COGS
		MCD1	MCD2	MCD3	Gen
Causal Transformer	baseline	21.8 ± 3	25.6 ± 2	19.7 ± 2	51.9 ± 4
	Meta-ICL	71.2 ± 7	74.8 ± 9	50.7 ± 7	75.7 ± 1
GPT-2	fine-tuning	42.8 ± 15	49.6 ± 16	39.6 ± 2	81.0 ± 0
	Meta-ICL	84.7 ± 3	71.5 ± 10	69.1 ± 7	82.4 ± 0

Table 2: Exp. 5: Comparison of causal transformer (Causal Tr., copied from Tab. 1) and GPT-2 with and without meta-ICL training, with best M for each dataset (mean sequence-level accuracies and standard deviations across runs).

For COGS we do not see a clear trend, except for $M = 25$. We believe that this can be attributed to the *informativeness* of support examples. We non-rigorously define *informativeness* as the set of underlying latent rules that governs the input-output mappings available in the given context relevant for answering the next (i.e., test) example. Since no shuffling was applied, the COGS models appear to have memorized the examples, though it was being regularized in doing so. The pressure to memorize seems to be stronger in COGS because the dataset is more diverse than SCAN, hence the support examples that make up the context are less likely to be informative. When trained with larger M (e.g., $M = 25$), where informative examples become more probable, the COGS model shows a similar pattern to SCAN, where the past examples are bound to be more informative due to its small vocabulary. This bolsters the importance of context informativeness for the emergence of in-context learning in Transformers. This result is also consistent with the findings of [Chan et al. \(2022\)](#) who demonstrated that informative context can drive the emergence of in-context learning, even despite the possibility of memorization.

5.4. Exp. 4: New Distributions

Setup. We now test how general our models can in-context learn by experimenting how well they can learn from a new distribution. For SCAN splits, we hold out 49 examples from the *test* set to sample our support examples during evaluation. We chose the value 49 as this is the maximum k for any model. We then evaluate our models on the rest of the test set by sampling from the held out *test* examples to form the contexts. We repeat Exp. 3 and report on the relative improvement (RI), $\frac{new-old}{old} \times 100$.

Result. As Figure 4 shows, for every label-shuffled model (left column), there exists a value of k that leads to an improved performance. This suggests that the learned in-context learning ability is *general* to a degree, being able to learn from *test* examples. This is especially true for MCD2

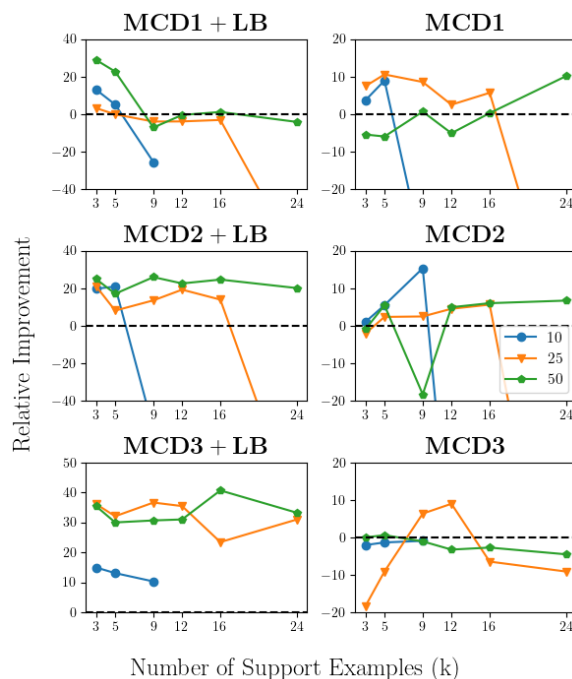


Figure 4: Exp. 4: Models evaluated with different numbers of support examples k sampled from the held-out portion of the *test* set for both with label shuffling (LB) (left column) and without (right). Relative improvement (RI) is calculated using $\frac{new-old}{old} \times 100$.

and MCD3, where the model improves for most (MCD2) and all (MCD3) values of k . This is interesting, because these are the more difficult splits of SCAN, where the *test* support examples can be much more informative. Contrarily, the models without label-shuffling (right column) do not show clearly an ability to learn from a novel distribution, especially for the more difficult splits of SCAN where the label-shuffled counterparts benefited the most. This demonstrates the existence of conflict between memorization and in-context learning in Transformers, where the possibility of memorization can negatively affect their ability to compositionally generalize using in-context learning upon its emergence.

For clarity, Figure 4 omits $k=1$ and $k=49$. $k=1$ leads to a significant positive RI for the LB models, and $k=49$ leads to significant deterioration for all models by failing to predict the end-of-sequence token properly (omitted for clarity). This is probably due to the model having to predict in positions far exceeding the maximum value of position seen during training, which has been known to be difficult for Transformers (Newman et al., 2020; Csordás et al., 2021).

5.5. Exp. 5: Meta-ICL for Pre-trained Models

Setup. Finally, we investigate how the effects of applying the meta-ICL framework carry over from models trained from scratch to the type of pre-trained causal language model typically used in NLP such as GPT-2 (Radford et al., 2019). We train two types of GPT-2 models without label shuffling for SCAN and COGS³: a supervised fine-tuned model ($M=1$) and meta-ICL models $M=\{10, 25\}$. Details on training, hyperparameters, and checkpointing are provided in Appendices A and B.

Result. Table 2 summarizes the result. Evidently, pre-trained models such as GPT-2 can also benefit from further training with meta-in context learning (compare rows 3 and 4). GPT-2 also outperforms the models trained from scratch with meta-ICL (row 2) in 3 out of 4 settings. However, the performance gain from meta-ICL training is diminished due to the pre-training (compare rows 3 and 4). In sum, these results confirm that pre-training on a large collection of natural language data lead to an acquisition of a prior more conducive to in-context learning, which agrees with the general consensus in the research community (Brown et al., 2020; Chowdhery et al., 2022; Min et al., 2022). Hence, our results are also relevant for the types of pretrained Transformer-based causal language models in widespread use in NLP today.

6. Conclusion

In this paper, we have studied the emergence of compositional generalization with a meta-learning regime which forces a sequence-to-sequence model to learn to *in-context learn*. We have investigated on three difficult datasets: the MCD splits of SCAN, COGS, and TMCD of GEO. Our main results showed that the meta-trained models show substantially better compositional generalization than the baselines in SCAN and GEO and closely matching in COGS.

Our results provide evidence that in-context learning can induce compositional generalization. We

confirm this relationship through various ablation studies, illustrating that compositional generalization can be improved by training on more in-context learning problems, in-context learning can emerge when learning from informative contexts despite the possibility of memorization, and the existence of conflict between memorization and compositional generalization. In this way, our study represents one step towards a deeper understanding of in-context learning. This can arguably improve our handling of out-of-distribution generalization, which is a fundamental challenge for effective machine learning (Ye et al., 2023), as well as making our learning models more plausible on the cognitive side, given the very limited memorization capabilities of humans (Fodor and Pylyshyn, 1988).

In future work, we plan to investigate the effect of using relative positional encoding (Dai et al., 2019) which are widely applied in recent LLMs as it has been shown to improve compositional generalization (Ontanon et al., 2022). Second, we believe that it would be worthwhile to investigate the effect of equipping the model with retrieval method to better choose the support examples in the future.

Limitations

Our results are relevant for generally elucidating the relationship between in-context learning and compositional generalization for a class of learning algorithms with memory. However, it must be observed that our work is predominantly empirical in nature and is not supported by theoretical guarantees. Also, it is conducted in an experimental lab-setting utilizing compositional generalization benchmarks that are well-established but are of synthetic nature. Hence, there might exist exceptions to the presented observations. The practical applications of the proposed methods for larger real world datasets are uncertain and need to be observed.

Furthermore, our approach in assessing a model’s compositional generalization follows the field’s predominant approach to this question: by generating a single data distribution and splitting them into a train and test in a systematic manner to test for compositionality. Although we have experimented with multiple datasets, there are still limitations in each chosen method of splitting used in our work. This implies that the degree of learned compositionality of our models might only partially encompass the entire spectrum of compositionality present in human languages.

Ethics Statement

Our work is concerned with foundational questions of learning generalizable models. It does not intro-

³We use Huggingface "gpt2" (Wolf et al., 2020)

duce new risks, nor does it involve sensitive applications. The datasets and pre-trained models are publicly available. Computational costs for training are relatively low. We do not believe that there are substantial ethical concerns in our work.

7. References

- Ekin Akyürek, Afra Feyza Akyürek, and Jacob Andreas. 2021. [Learning to recombine and resample data for compositional generalization](#). In *Proceedings of the International Conference on Learning Representations*.
- Shengnan An, Zeqi Lin, Qiang Fu, Bei Chen, Nanning Zheng, Jian-Guang Lou, and Dongmei Zhang. 2023. [How do in-context examples affect compositional generalization?](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11027–11052, Toronto, Canada. Association for Computational Linguistics.
- Jacob Andreas. 2020. [Good-enough compositional data augmentation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7556–7566, Online. Association for Computational Linguistics.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. Layer normalization. ArXiv:1607.06450.
- Marco Baroni. 2020. [Linguistic generalization and compositionality in modern artificial neural networks](#). *Philosophical Transactions of the Royal Society B: Biological Sciences*, 375(1791):20190307.
- Y. Bengio, S. Bengio, and J. Cloutier. 1991. [Learning a synaptic learning rule](#). In *International Joint Conference on Neural Networks*, volume ii, Seattle, WA.
- Leon Bergen, Timothy O' Donnell, and Dzmitry Bahdanau. 2021. [Systematic generalization with edge transformers](#). In *Advances in Neural Information Processing Systems*, volume 34, pages 1390–1402. Curran Associates, Inc.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Proceedings of NeurIPS*, volume 33, pages 1877–1901.
- Stephanie C. Y. Chan, Adam Santoro, Andrew K. Lampinen, Jane X. Wang, Aaditya Singh, Pierre H. Richemond, Jay McClelland, and Felix Hill. 2022. [Data Distributional Properties Drive Emergent In-Context Learning in Transformers](#). In *Proceedings of NeurIPS*.
- Yanda Chen, Ruiqi Zhong, Sheng Zha, George Karypis, and He He. 2022. [Meta-learning via language model in-context tuning](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 719–730, Dublin, Ireland. Association for Computational Linguistics.
- Noam Chomsky. 1965. *Aspects of the Theory of Syntax*. The MIT Press, Cambridge.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2022. Palm: Scaling language modeling with pathways. ArXiv:2204.02311.
- Henry Conklin, Bailin Wang, Kenny Smith, and Ivan Titov. 2021. [Meta-Learning to Compositionally Generalize](#). 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 3322–3335, Online. Association for Computational Linguistics.
- Róbert Csordás, Kazuki Irie, and Juergen Schmidhuber. 2021. [The devil is in the detail: Simple](#)

- tricks improve systematic generalization of transformers. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 619–634, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le, and Ruslan Salakhutdinov. 2019. [Transformer-XL: Attentive language models beyond a fixed-length context](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2978–2988, Florence, Italy. Association for Computational Linguistics.
- Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Lukasz Kaiser. 2019. [Universal transformers](#). In *Proceedings of the International Conference on Learning Representations*.
- Li Dong and Mirella Lapata. 2016. [Language to logical form with neural attention](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 33–43, Berlin, Germany. Association for Computational Linguistics.
- Yan Duan, John Schulman, Xi Chen, Peter L. Bartlett, Ilya Sutskever, and Pieter Abbeel. 2017. RL²: Fast reinforcement learning via slow reinforcement learning. ArXiv:1611.02779.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. [Model-agnostic meta-learning for fast adaptation of deep networks](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1126–1135. PMLR.
- Jerry A Fodor and Zenon W Pylyshyn. 1988. Connectionism and Cognitive Architecture: A Critical Analysis. *Cognition*, 28:3–71.
- Yao Fu, Hao Peng, Ashish Sabharwal, Peter Clark, and Tushar Khot. 2023. [Complexity-based prompting for multi-step reasoning](#). In *Proceedings of the International Conference on Learning Representations*.
- Daniel Furrer, Marc van Zee, Nathan Scales, and Nathanael Schärli. 2021. [Compositional Generalization in Semantic Parsing: Pre-training vs. Specialized Architectures](#). ArXiv:2007.08970 [cs].
- Shivam Garg, Dimitris Tsipras, Percy S Liang, and Gregory Valiant. 2022. [What can transformers learn in-context? a case study of simple function classes](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 30583–30598. Curran Associates, Inc.
- Jonathan Gordon, David Lopez-Paz, Marco Baroni, and Diane Bouchacourt. 2020. [Permutation equivariant models for compositional generalization in language](#). In *Proceedings of the International Conference on Learning Representations*.
- Jiatao Gu, Yong Wang, Yun Chen, Victor O. K. Li, and Kyunghyun Cho. 2018. [Meta-learning for low-resource neural machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3622–3631, Brussels, Belgium. Association for Computational Linguistics.
- Yinuo Guo, Zeqi Lin, Jian-Guang Lou, and Dongmei Zhang. 2020a. [Hierarchical poset decoding for compositional generalization in language](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 6913–6924. Curran Associates, Inc.
- Yinuo Guo, Hualei Zhu, Zeqi Lin, Bei Chen, Jian-Guang Lou, and Dongmei Zhang. 2020b. Revisiting iterative back-translation from the perspective of compositional generalization. In *Proceedings of AAAI*.
- Jonathan Herzig, Peter Shaw, Ming-Wei Chang, Kelvin Guu, Panupong Pasupat, and Yuan Zhang. 2021. [Unlocking Compositional Generalization in Pre-trained Models Using Intermediate Representations](#). ArXiv:2104.07478 [cs].
- Sepp Hochreiter, A. Steven Younger, and Peter R. Conwell. 2001. Learning to learn using gradient descent. In *Proceedings of the International Conference on Artificial Neural Networks*, page 87–94, Berlin, Heidelberg. Springer-Verlag.
- Arian Hosseini, Ankit Vani, Dzmitry Bahdanau, Alessandro Sordani, and Aaron Courville. 2022. [On the compositional generalization gap of in-context learning](#). In *Proceedings of the Fifth BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 272–280, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- 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.
- Yichen Jiang and Mohit Bansal. 2021. [Inducing transformer’s compositional generalization ability via auxiliary sequence prediction tasks](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6253–6265, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

- Daniel Keyzers, 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 *Proceedings of the 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.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Louis Kirsch, James Harrison, Jascha Sohl-Dickstein, and Luke Metz. 2022. [General-purpose in-context learning by meta-learning transformers](#). In *Sixth Workshop on Meta-Learning at the Conference on Neural Information Processing Systems*.
- Brenden M Lake. 2019. [Compositional generalization through meta sequence-to-sequence learning](#). In *Proceedings of NeurIPS*, volume 32. Curran Associates, Inc.
- Brenden M. Lake and Marco Baroni. 2017. Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks. In *Proceedings of the International Conference on Machine Learning*.
- David B. Leake. 1996. *Case-Based Reasoning: Experiences, Lessons and Future Directions*, 1st edition. MIT Press, Cambridge, MA, USA.
- Yuanpeng Li, Liang Zhao, Jianyu Wang, and Joel Hestness. 2019. [Compositional generalization for primitive substitutions](#). 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 4293–4302, Hong Kong, China. Association for Computational Linguistics.
- Zhaoyi Li, Ying Wei, and Defu Lian. 2023. [Learning to substitute spans towards improving compositional generalization](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2791–2811, Toronto, Canada. Association for Computational Linguistics.
- Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han. 2020a. [On the variance of the adaptive learning rate and beyond](#). In *Proceedings of the International Conference on Learning Representations*.
- Qian Liu, Shengnan An, Jian-Guang Lou, Bei Chen, Zeqi Lin, Yan Gao, Bin Zhou, Nanning Zheng, and Dongmei Zhang. 2020b. [Compositional generalization by learning analytical expressions](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 11416–11427. Curran Associates, Inc.
- Sewon Min, Mike Lewis, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2022. [MetalCL: Learning to learn in context](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2791–2809, Seattle, United States. Association for Computational Linguistics.
- Benjamin Newman, John Hewitt, Percy Liang, and Christopher D. Manning. 2020. [The EOS decision and length extrapolation](#). In *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 276–291, Online. Association for Computational Linguistics.
- Farhad Nooralahzadeh, Giannis Bekoulis, Johannes Bjerva, and Isabelle Augenstein. 2020. [Zero-shot cross-lingual transfer with meta learning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4547–4562, Online. Association for Computational Linguistics.
- Maxwell I. Nye, Armando Solar-Lezama, Joshua B. Tenenbaum, and Brenden M. Lake. 2020. [Learning Compositional Rules via Neural Program Synthesis](#). ArXiv:2003.05562 [cs].
- Santiago Ontanon, Joshua Ainslie, Zachary Fisher, and Vaclav Cvicek. 2022. [Making transformers solve compositional tasks](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3591–3607, Dublin, Ireland. Association for Computational Linguistics.
- Pedro A. Ortega, Jane X. Wang, Mark Rowland, Tim Genewein, Zeb Kurth-Nelson, Razvan Pascanu, Nicolas Heess, Joel Veness, Alex Pritzel, Pablo Sprechmann, Siddhant M. Jayakumar, Tom McGrath, Kevin Miller, Mohammad Azar, Ian Osband, Neil Rabinowitz, András György, Silvia Chiappa, Simon Osindero, Yee Whye Teh, Hado van Hasselt, Nando de Freitas, Matthew

- Botvinick, and Shane Legg. 2019. [Meta-learning of Sequential Strategies](#). ArXiv:1905.03030 [cs, stat].
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In *Proceedings of NeurIPS*, pages 8024–8035. Curran Associates, Inc.
- Arkil Patel, Satwik Bhattamishra, Phil Blunsom, and Navin Goyal. 2022. [Revisiting the compositional generalization abilities of neural sequence models](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 424–434, Dublin, Ireland. Association for Computational Linguistics.
- Kun Qian and Zhou Yu. 2019. [Domain adaptive dialog generation via meta learning](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2639–2649, Florence, Italy. Association for Computational Linguistics.
- Linlu Qiu, Peter Shaw, Panupong Pasupat, Pawel Nowak, Tal Linzen, Fei Sha, and Kristina Toutanova. 2022a. [Improving compositional generalization with latent structure and data augmentation](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4341–4362, Seattle, United States. Association for Computational Linguistics.
- Linlu Qiu, Peter Shaw, Panupong Pasupat, Tianze Shi, Jonathan Herzig, Emily Pitler, Fei Sha, and Kristina Toutanova. 2022b. [Evaluating the impact of model scale for compositional generalization in semantic parsing](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9157–9179, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. Technical report, OpenAI.
- Jake Russin, Jason Jo, Randall C. O’Reilly, and Yoshua Bengio. 2019. [Compositional generalization in a deep seq2seq model by separating syntax and semantics](#). ArXiv: 1904.09708.
- Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. 2016. Meta-learning with memory-augmented neural networks. In *Proceedings of ICML*, page 1842–1850, New York, NY.
- Jürgen Schmidhuber, Jieyu Zhao, and Marco Wiering. 1996. Simple principles of metalearning. Technical report, Istituto Dalle Molle di Studi sull’Intelligenza Artificiale.
- 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.
- 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 *Proceedings of NeurIPS*, page 6000–6010, Red Hook, NY, USA. Curran Associates Inc.
- Jane X. Wang, Zeb Kurth-Nelson, Dhruva Tirumala, Hubert Soyer, Joel Z. Leibo, Remi Munos, Charles Blundell, Dharshan Kumaran, and Matt Botvinick. 2017. [Learning to reinforcement learn](#). ArXiv:1611.05763 [cs, stat].
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V Le, and Denny Zhou. 2022. [Chain of thought prompting elicits reasoning in large language models](#). In *Proceedings of NeurIPS*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. 2022. [An Explanation of In-context Learning as Implicit Bayesian Inference](#). In *Proceedings of ICLR*.
- Nanyang Ye, Lin Zhu, Jia Wang, Zhaoyu Zeng, Jiayao Shao, Chensheng Peng, Bikang Pan,

Kaicun Li, and Jun Zhu. 2023. [Certifiable out-of-distribution generalization](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(9):10927–10935.

John M. Zelle and Raymond J. Mooney. 1996. Learning to parse database queries using inductive logic programming. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence - Volume 2, AAAI'96*, page 1050–1055. AAAI Press.

Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc V Le, and Ed H. Chi. 2023. [Least-to-most prompting enables complex reasoning in large language models](#). In *Proceedings of the International Conference on Learning Representations*.

A. Hyperparameters and Computing Resource

Causal Transformer We use the PyTorch (Paszke et al., 2019) implementation of RAdam (Liu et al., 2020a) as our choice of optimizer with the learning rate of 1×10^{-4} and $\beta = (0.9, 0.99)$ for all of our experiments. For stable training, we apply a linear warm-up for 500 steps for SCAN and GeoQuery and 5000 for COGS. We clip the gradient whenever the norm exceeds 5. We apply dropout rate of 0.1, ReLU activation, and batch-size of 5. Small batch size was used due to the limited availability of computing resources. Parameters are initialized according to the default initialization method of PyTorch. This means that the word embeddings are initialized by drawing from a standard normal. The embeddings have the same dimension as the model. We use the variant where the LayerNorm (Ba et al., 2016) is applied before each sub-block for SCAN and a normal configuration for COGS and GeoQuery. The resulting model size has 25.2 million parameters for all datasets. For the causal baseline (i.e., $M = 1$), use of a low batch-size leads to unstable training, hence we increase the batch-size to 256. Finally, we note that we did not perform any systematic hyperparameter tuning and most of the used hyperparameters were initial guesses.

Transformer and Universal Transformer Both Transformer and Universal Transformer baselines are a 3-layer encoder-decoder architecture, and it is adapted from the code release of Csordás et al. (2021). We use the same learning rate 1×10^{-4} and $\beta = (0.9, 0.99)$ using Adam (Kingma and Ba, 2014) as our optimizer of choice. We do not use the default learning rate value of PyTorch as we saw the alternative configuration to be more stable. The dimension of 128 is used for both model state and word embeddings with 8 heads and 256 for feed-forward dimension. We use the dropout rate of 0.1 and batch size of 256.

GPT-2 We use the same set of hyperparameters as the causal Transformer models trained from scratch described above, except for the learning rate which we set to 5×10^{-5} .

Computing Resources We used a single GeForce RTX 2080 Ti 11G for our SCAN and GeoQuery experiments and a single GeForce GTX TITAN X 12G for our COGS experiments.

B. Checkpoint Selection for Evaluation

For SCAN, we follow the checkpoint selection method of [Conklin et al. \(2021\)](#) and use the available development set to pick the checkpoint for testing by training for 20k steps evaluating every 1000 steps. Usually, each model with meta-training takes around 10k steps to converge. The causal Transformer baselines, Transformers, and Universal Transformers all take much less time to converge, hence we only train for 10k steps.

For COGS, we simply train the models for 150k steps and take the last checkpoint for evaluation. This was similarly done in [Conklin et al. \(2021\)](#), but they use 10% of the test set to tune their hyperparameters. There is a validation set associated with the training set in COGS, but it is a widely known that tuning on this set does not work well ([Csordás et al., 2021](#)) as the model continues to improve on the test even when the model scores perfectly on the train and validation set.

For GeoQuery, we train the models for 50k steps and select the last checkpoint for evaluation.

C. Datasets and Preprocessing

SCAN The preprocessing decreases the average output length of the dataset from 14.3 to 12.2 and the maximum sequence length from 48 to 17. In the new format, the overall vocabulary size of SCAN is 30 with 11 output words, 4 special symbols and 15 input words. Table 3 shows a few examples.

COGS The resulting preprocessing is illustrated in Table 4. Before preprocessing, the average length is 51.07 and maximum of 175 which becomes 28.01 and 96 respectively after preprocessing. The resulting vocabulary size is 871.

GeoQuery We do not apply any pre-processing for intermediate representation. As mentioned in the main text, we only replace the entities with placeholders.

We noted that our intermediate representation method is different from the format found to be useful, which converts the sequence prediction task to a sequence tagging task ([Ontanon et al., 2022](#)). This is possible because in COGS, the target output is a concatenation of five semantically parsed "tags": a parent, the role of the parental relation, the category, the noun determiner, and the verb name. Hence, instead of generating the output sequence, a model can be made to tag each input token in parallel and combine the results for the final prediction.

Command (Input)	Before (Output)	After (Output)
run twice	RUN RUN	RUN * 2
jump after run	RUN JUMP	RUN + JUMP
jump around right	RTURN JUMP RTURN JUMP RTURN JUMP RTURN JUMP	(RTURN JUMP) * 4
jump around right and walk twice	RTURN JUMP RTURN JUMP RTURN JUMP RTURN JUMP WALK WALK	(RTURN JUMP) * 4 + WALK * 2

Table 3: Example SCAN action sequences (outputs) before and after preprocessing.

Sentence (Input)	Before (Output)	After (Output)
A rose was helped by a dog .	rose (x_1) AND help . theme (x_3 , x_1) AND help . agent (x_3 , x_6) AND dog (x_6)	rose 1 AND help . theme 3 1 AND help . agent 3 6 AND dog 6
Charlie loaned the cake in a house to the girl .	* cake (x_3) ; * girl (x_9) ; loan . agent (x_1 , Charlie) AND loan . theme (x_1 , x_3) AND loan . recipient (x_1 , x_9) AND cake . nmod . in (x_3 , x_6) AND house (x_6)	* cake 3 ; * girl 9 ; loan . agent 1 Charlie AND loan . theme 1 3 AND loan . recipient 1 9 AND cake . nmod . in 3 6 AND house 6

Table 4: Example COGS semantic parsing results before and after preprocessing.