A Neural Rewriting System to Solve Algorithmic Problems

Flavio Petruzzellis^{a,*}, Alberto Testolin^{a,b} and Alessandro Sperduti^a

^aDepartment of Mathematics, University of Padova, Padova, Italy ^bDepartment of General Psychology, University of Padova, Padova, Italy

Modern neural network architectures still struggle to Abstract. learn algorithmic procedures that require to systematically apply compositional rules to solve out-of-distribution problem instances. In this work, we focus on formula simplification problems, a class of synthetic benchmarks used to study the systematic generalization capabilities of neural architectures. We propose a modular architecture designed to learn a general procedure for solving nested mathematical formulas by only relying on a minimal set of training examples. Inspired by rewriting systems, a classic framework in symbolic artificial intelligence, we include in the architecture three specialized and interacting modules: the Selector, trained to identify solvable sub-expressions; the Solver, mapping sub-expressions to their values; and the Combiner, replacing sub-expressions in the original formula with the solution provided by the Solver. We benchmark our system against the Neural Data Router, a recent model specialized for systematic generalization, and a state-of-the-art large language model (GPT-4) probed with advanced prompting strategies. We demonstrate that our approach achieves a higher degree of out-of-distribution generalization compared to these alternative approaches on three different types of formula simplification problems, and we discuss its limitations by analyzing its failures.

1 Introduction

Whether neural networks can learn to perform systematic reasoning is a long-standing question in artificial intelligence and cognitive science [10], and the recent success of deep learning on challenging reasoning tasks has reinvigorated this debate [11]. However, despite some promising achievements [15], there is general consensus that even state-of-the-art transformer architectures and large language models still lack systematic and compositional generalization skills [6, 9, 28].

In this work, we investigate this issue by considering a type of systematic generalization benchmark that requires simplifying symbolic formulas to a minimal form. Specifically, we consider the problem of simplifying nested formulas of operations on lists of integers [22], as well as arithmetic and algebraic expressions. In principle these problems could be tackled by automatic symbolic solvers, but here we consider them as benchmarks to assess the reasoning capabilities of trained neural networks under controlled settings [7, 14].

We formalize formula simplification problems and propose a simple solution method inspired by rewriting systems rooted in the tradition of symbolic artificial intelligence [8]. We implement this solution method as a modular neural architecture called the Neural Rewriting System, which represents a case study of a neural architecture in which the roles and interaction mechanisms between neural modules are clearly defined, allowing greater interpretability of the resulting system. We exploit our framework to further evaluate two neural architectures that have been recently shown able to tackle reasoning tasks: the Neural Data Router [7], a transformer encoder trainable end-to-end to solve algorithmic problems, and GPT-4 [24] probed using self-consistency zero-shot Chain-of-Thought [32], which is an advanced prompting method that improves the performance of large language models on reasoning tasks.

Through extensive experiments and analyses, we show that the proposed modular architecture achieves greater accuracy compared to advanced alternative approaches, especially when solving deeply nested formulas. Furthermore, we show that our approach allows us to understand in detail what parts of a systematic generalization problem could represent the greatest obstacles for neural architectures. Indeed, when analyzing the causes of failure of the Neural Rewriting System, we find that accurately selecting the portion to be simplified in very long formulas is the most challenging task for the architecture, shedding light on the limitations of current approaches to deal with length generalization in algorithmic tasks with transformers.

2 Related works

Systematic Generalization A consistent stream of research investigates the extent to which neural networks can learn to reason in a systematic and compositional way [11, 30]. Several benchmarks have been proposed to measure this capacity: among the most popular, ListOps [22] and CTL [18] can be seen as instances of formula simplification problems, in which the learning system is tasked to produce the evaluation of a formula which typically requires to compute intermediate results, either implicitly or explicitly. Other formula simplification benchmarks include derivation and integration problems [16] or polynomial simplification [1]. Modern neural architectures achieve different degrees of systematic compositional generalization [7, 15, 23]. Interestingly, recent work has shown the emergence of specialized modules in Mixture of Experts architectures [21], suggesting that modular architectures can outperform monolithic ones in both in- and out-of-distribution testing conditions.

Rewriting systems and neural networks The idea of implementing neural architectures inspired by symbolic rewriting systems was initially explored in unsupervised learning settings, where individual network weights represented tokens to be rewritten [13]. Others im-

^{*} Corresponding Author. Email: flavio.petruzzellis@phd.unipd.it

plemented a rewriting system for algebraic problems using custom feature engineering methods and a feed-forward network [3], while in [4] the authors propose a reinforcement learning-based system that can learn a general rewriting mechanism by first choosing a region to simplify and then an appropriate rewriting rule.

LLMs Research on prompt engineering techniques has shown that appropriate prompting methods can boost the reasoning capabilities of LLMs [32, 33]. In particular, Chain-of-Though prompting improves LLMs' accuracy on reasoning tasks by leveraging their autoregressive nature: the models are elicited to give the final result using multi-step reasoning chains, which enable to process contextual information more effectively. Among the reasoning tasks on which Large Language Models often exhibit impressive problem-solving capabilities we find 'symbolic reasoning tasks'. Some examples are coin flip, last letter concatenation [33] and boolean variable assignment [2]. These tasks are remarkably similar in spirit to those we consider in this work, since problem instances are synthetically generated and can be solved by applying simple algorithms.

3 Formula simplification problems

Many systematic generalization benchmarks require to evaluate a possibly very complex formula by replacing expressions with elementary operands that have an equivalent value. We call such problems 'formula simplification problems'. For example, in ListOps an operation on a list of integers can be replaced with a single integer that is the result of the operation. In order to use formula simplification problems as a benchmark to assess the systematic generalization abilities of neural architectures, we outline a formal framework that includes the class of problems we consider in this work.

In formula simplification problems, we can see formulas $f \in F$ as entities that are composed of two semantically distinct elements: operators $o \in O$ and arguments $a \in A$.¹ Arguments can be either atomic elements $e \in E$, such as integers, which are also the final values of any formula, or other formulas: therefore, in general, formulas have a recursive structure (see Figure 1). Furthermore, since in these problems each formula has always one and only one final value, any formula has the generic compact form $f = o(a_1, ..., a_n) = e$, where $o \in O$, $e \in E$ and $a_j \in F \cup E$, $\forall j \in [1, n]$. Finally, we can define leaf formulas $F^L \subset F$ as the subset of formulas whose arguments are all atomic elements: $f^L = o(a_1, ..., a_n)$ s.t. $a_k \in E$, $\forall k \in [1, n]$.

The simplest algorithm that can be applied to solve formula simplification problems is to solve leaf formulas one by one until the formula is completely reduced to an atomic element e. Therefore, for any problem, we can define a set of rewriting rules $r \in R$ which map leaf formulas to their values: $r : F^L \to E$, $\forall r \in R$. Taking into account the problems we consider in this work, we can further characterize the rewriting rules in R. First, the rules apply to contiguous sub-sequences of symbols. Second, each rewriting rule is contractive; that is, it maps sub-sequences to shorter ones. Third, the rewriting rules have mutually exclusive applicability scope, i.e. given a leaf formula, one and only one rule can be applied to reduce it. Fourth, several leaf formulas could be present in a formula at once, but the order in which rewriting rules are applied on them is irrelevant.

The notation introduced above allows us to outline a simple algorithm to solve nested formulas, consisting of the iterative application of the composition of three functions: $sel : F \to F^L$, that maps a formula to a leaf formula appearing in it; $sol : F^L \to E$ that solves



Figure 1: Examples of solution of ListOps, arithmetic and algebraic formulas. Formulas f are reduced to atomic values e by iteratively solving leaf formulas f^L , highlighted in yellow.

the leaf formula, i.e. applies an appropriate rewriting rule to map it to an equivalent atomic element; $com : F \times F^L \times E \to F$, that combines the initial formula f and the solution of the leaf formula f^L by replacing f^L with $sol(f^L)$ in the initial formula. Notice that the function sel can be applied again on the output of the function com which is a valid formula, thus enabling the implementation of an iterative solution procedure.

Throughout this work, we will use this simple algorithm as a reference to compare the way different models solve formula simplification problems and to critically evaluate their results. Furthermore, we propose a neural implementation of this algorithm in the form of a modular neural architecture called the Neural Rewriting System.

4 Neural Rewriting System

In this section, we describe the proposed Neural Rewriting System, a neural architecture composed of three integrated modules, called Selector, Solver, and Combiner, each being the neural implementation of one of the functions that can be used to solve formula simplification problems according to the algorithm described in Section 3. The modules are trained independently of each other and interact at test time to solve nested formulas iteratively. The model is designed to achieve strong out-of-distribution generalization, i.e. to solve very deep formulas while being trained on a set of much shallower formulas. We describe the composition of the datasets used in the development of the model in Section 5.2.1. A schematic representation of the system is shown in Figure 2.

4.1 Selector Module

The Selector module implements the function $sel : F \to F^L$, i.e. it is trained to map a formula to a leaf formula appearing therein. In analogy to what happens in humans when they deploy object-based attention to locate algebraic sub-expressions that can be simplified [19], the Selector is trained to identify the last leaf formula² occurring in the input formula on which a rewriting rule can be applied. When training the Selector, we assume that it will always receive syntactically correct formulas, i.e. we use teacher forcing.

We frame the problem as a sequence-to-sequence task and use a variant of the transformer encoder-decoder [31] to implement the Selector. We modify the vanilla transformer in two ways. First, given

 $[\]overline{1}$ We also sometimes refer to arguments as operands. We prefer to use the term arguments here to avoid confusion with operators O.

² To enable generalization on formulas with more arguments than seen during training, if the identified leaf formula has more than two arguments, we train the Selector to output the smallest fraction of the leaf formula which can be simplified. This corresponds to the operator and the two first arguments appearing in the leaf formula.



Figure 2: Schematic representation of the Neural Rewriting System.

the class of problems we consider in this work, we can assume that the problem the Selector needs to solve has a local nature: the leaf formula to be simplified is always located in a limited region of the input sequence. We therefore mask all entries in the self-attention matrix of the encoder but the ones around the main diagonal (i.e., we make them -inf). The diagonal window is 2k+1 tokens wide, where k is a hyperparameter (preliminary experiments showed that models with a vanilla self-attention achieved worse out-of-distribution generalization). Second, to enable generalization to out-of-distribution problem instances, we require the Selector to identify a leaf formula independently of the length of the input sequence, provided that it is syntactically correct. Recent research on length generalization in transformers provided evidence that their ability to generalize on longer samples can be influenced by the choice of positional encodings, especially when, at test time, these fall out of the range observed during training [5, 29]. We thus use Label-based Positional Encodings [17] to enable the Selector to identify leaf formulas in very long sequences. The positional information of an input sequence of L tokens is thus encoded in the following way: given a sequence of N sinusoidal positional encodings, where N is a large number that represents the maximum expected length of an input, L integers are sampled in the interval [0, N - 1] and then sorted. The encodings found in the positions corresponding to the sampled integers are then summed to the embeddings of the tokens in the input sequence.

Multi-output generation Label positional encodings introduce randomness in the representation of input sequences. We have observed experimentally that to mitigate the effect of randomness it can be useful to repeat the auto-regressive generation of transformer outputs. After sampling several output sequences from the probability distribution derived from the decoder's outputs, we choose the best one considering both a measure of confidence of the Selector and a measure of input-output agreement computed by the Combiner.

Given the specialized purpose of the Selector, we can see each output of the module as a candidate leaf formula \hat{f}^L . We generate

any token \hat{f}_i^L in an output sequence \hat{f}^L by sampling from the probability distribution obtained applying the softmax function to the logits produced by the final fully-connected layer of the decoder. We do not use any temperature parameter when sampling the output tokens. For any input formula f, we repeat the stochastic generation process M times, thus generating a sequence of candidate leaf formulas $\hat{F}^L = \langle \hat{f}^{L,1}, ..., \hat{f}^{L,M} \rangle$. We define the confidence of the Selector on any $\hat{f}^{L,j}, 1 \leq j \leq M$ as the joint probability of sampling its tokens: $c(\hat{f}^{L,j}) = \prod_{i=1}^N p_i^i$, where N is the number of tokens in $\hat{f}^{L,j}$, and p_i^j is the probability to sample token \hat{f}_i^L in $\hat{f}^{L,j}$. We also define an agreement score $a(\hat{f}^{L,j}, f) \in [0,1]$ which gives information about the fraction of $\hat{f}^{L,j}$ which is exactly present in the input formula f. This measure is computed by the Combiner and thus it is formally defined in Section 4.2. We then select the final output f^L of the Selector as the one with the highest Selector confidence which has an agreement score equal to 1 — that is, it matches the input sequence exactly. More formally, $f^L = \hat{f}^{L,j} \in \hat{F}^L$ s.t. $c(\hat{f}^{L,j}) \geq c(\hat{f}^{L,k}) \, \forall j, k \in [1,M] \wedge a(\hat{f}^{L,j}, f) = 1$.

Dynamic windowing We implement a dynamic windowing mechanism on longer input sequences that allows us to increase the model's generalization capacity on complex problem instances. The core idea behind this mechanism is to repeat the process of selecting a leaf formula several times, changing each time the window of the input formula that the Selector observes, and then relying on the confidence $c(\hat{f}^L)$ to pick the best output. We apply this mechanism on top of multi-output generation by modifying its behavior for sequences longer than a given threshold T. Given an input formula f, if |f| < T the computation is executed as described before. Otherwise, we generate M copies of the input $\langle f^{(1)}, ..., f^{(M)} \rangle$, whose lengths will be reduced by applying a window function w. Considering any input f as a sequence $f_1, ..., f_N$ of N tokens, we define the window function $w(f, k) = f_{k+1}, ..., f_N$ which reduces the length of the input by giving as output its last k tokens. Since the Selector is trained to output the last leaf formula appearing in the input, the window function reduces the input length starting from the first tokens. We divide the sequence of copies of the input $\langle f^{(1)},...,f^{(M)}\rangle$ into 20 groups $F^{(1)}, \dots, F^{(20)}$ of equal size. Intuitively, in each group the length of the input is reduced by a different percentage of tokens. More formally, the window function will be parameterized by $k = \text{floor}(|f^{(i)}| \cdot \frac{j}{20}) \forall f^{(i)} \in F^{(j)}, \forall j \in [1, 20]$. We then pick the final leaf formula using the confidence and agreement scores, as described in the previous paragraph. This ensures that the model can observe the whole input sequence and select a leaf expression in the part of the input where it can identify one with more confidence.

4.2 Solver Module

Unlike classical rewriting systems that use symbolic rule dictionaries, we learn rewriting rules directly from data by implementing the function $sol : F^L \to E$ with a trainable neural network. Given a leaf formula f^L by the Selector, the Solver is trained to produce the equivalent reduction e according to a valid rewriting rule. When training the Solver, we assume that the Selector always generates perfect outputs (i.e., we use teacher forcing). The Solver also learns to recognize the computation's termination state, signaling when such a state is reached. It maps atomic elements, representing a formula's final value, to the special symbol ω , indicating the end of computation. Similar to the Selector, we frame the task as a sequence-tosequence problem. We use a vanilla transformer encoder-decoder for the Solver since it learns simple input-output mappings.

4.3 Combiner Module

The last module in the architecture is the Combiner, a neural implementation of the function $com : F \times F^L \times E \to F$. Its purpose is thus to produce a simplified version of the original formula, given the formula itself f, the leaf formula f^L identified by the Selector, and its reduction e computed by the Solver.

In order to carry out its task, the first operation that the Combiner must perform is finding the position in f where the leaf formula f^L appears. We notice that the convolution is a suitable operation to detect which portion of an input sequence has the highest match with another sequence used as a filter, so we implement this operation using a 2D Convolutional Neural Network (CNN) whose filters are set dynamically at execution time using the output of the Selector, rather than being learned with backpropagation.

More precisely, we represent both the input sequence f and the leaf formula f^L as sequences of 1-hot vectors over the same vocabulary. Since the leaf formulas found for different sequences in a batch can have different lengths, we pad each one with zeros to prevent the padding to match in the input. Then, we set the filter of the 2D CNN to the 1-hot representation of f^L . We refer to the CNN parameterized in this way as CNN_{f^L} . Doing so allows us to obtain from the output of the convolution both information on the location of the best match of f^L in f and on the number of tokens in f^L that match f exactly in some point. Indeed, we can compute the location of the best match as $pos(f^L, f) = argmax CNN_{f^L}(f)$. Furthermore, we can compute the agreement score $a(f^L, f) = \frac{\max \operatorname{CNN}_{f^L}(f)}{|f^L|}$, where $|f^L|$ is the number of tokens in f^L . Dividing by $|f^L|$ makes the score normalized, which allows us to compare the agreement scores of leaf formulas with different lengths. Indeed, as described in Section 4, the Selector uses this score for multi-output generation to discard the outputs that do not have an exact match in the input formula. Notice that in this case the CNN is parameterized using candidate leaf formulas \hat{f}_j^L whose accuracy scores with f are compared. If there is no Selector output such that $a(\hat{f}_i^L, f) = 1$, the computation on the input sequence f is stopped, and this is considered a failure of the model.

After finding the position of the leaf expression in f, the Combiner replaces f^L with e in f, to compute the simplified version of the formula f'. We implement this operation as a deterministic operator with input f, f^L , e, and $pos(f^L, f)$.

5 Experiments and results

We evaluated the capacity of our model to solve formula simplification problems on ListOps, arithmetic, and algebraic formulas. We compare the performance of the Neural Rewriting System to the Neural Data Router and GPT-4, one of the most advanced LLMs currently available. We make available the code and data used in our experiments [25] and a Technical Appendix containing methodological details and additional results [27].

5.1 Datasets

For all three tasks, the automatic generation of formulas is parameterized by specifying the nesting level of a formula. Any formula is nested at each level in two points: exactly two arguments in the formulas on that level will be other formulas. Formulas are tokenized at the character level for all tasks and trained models.

ListOps The ListOps dataset [22] was introduced to evaluate the capacity of neural networks to build parse trees of nested formulas. The original dataset consists of formulas composed of operations on

lists of integers, including minimum, maximum, median and sum modulo 10 of a list of integers. We modified the ListOps dataset so that formulas had exactly 2 nesting points at each nesting level, and we made it possible to specify the number of arguments that appear in formulas at any level. Since we are interested in the system's capacity to generalize on highly nested formulas rather than learning specific operations, we reduced the set of operations to minimum, maximum and sum modulo 10.

Arithmetic We generated formulas composed by sum, subtraction and multiplication operations between two integers sampled in the interval [-99, 99]. Since in this work we do not investigate the capacity to generalize to numbers with more digits than those seen during training, we used the modulo 100 of the intermediate results in the solution process.

Algebra We considered a subset of algebraic formulas that can always be reduced up to a minimal form in a deterministic way. Such formulas involve sum and subtractions between two monomials and their final value is always a monomial. The numerical coefficients of monomials were sampled in the interval [-99, 99], and each monomial can contain up to four literal variables sampled from the set $\{a, b, x, y\}$. All monomials appearing in a given formula have the same literal variables. As in the case of Arithmetic, all intermediate numerical values were taken modulo 100 when computing the final value of the formula.

5.2 Experiments

5.2.1 Neural Rewriting System

We describe here how we built the training and validation sets for the Selector and Solver modules. Further methodological details can be found in the Technical Appendix.

Selector In Arithmetic and Algebra problems we included in the training set formulas with nesting levels 1, 2 and 3. In the ListOps problem, in which we could specify the number of arguments that appear in the formulas, we included in the training set formulas with nesting levels 1 and 2 that have 2 or 3 arguments for each operator. In all three problems, the training set also contains atomic elements corresponding to the initial formula's final value. Solving formula simplification problems by iteratively simplifying leaf formulas produces several simplifications of the initial formula. To show to the Selector the complete solution process of formulas in the training set, we have also included in it formulas that appear as intermediate solution steps of the formulas described above.

We constructed a separate in-distribution validation set with samples with the same structural characteristics as those in the training set. Unlike common machine learning tasks where models are tested on the same data distribution they are trained on, we want the Selector to possess out-of-distribution (OOD) generalization capabilities to identify the leaf formulas even in longer inputs than those seen during training. For this reason, we also built a separate out-ofdistribution validation set containing formulas with greater structural complexity, and use this set for model selection. For Arithmetic and Algebra, we included samples with nesting levels 4, 5 and 6. For ListOps, we included formulas with nesting levels 3 and 4 with 2, 3 or 4 arguments for each operator. To select the most capable model across the iterative resolution process, we also included formulas representing examples of intermediate resolution steps. To control the structural complexity of the formulas, the OOD validation sets are balanced across the nesting level of the leaf expressions.³

³ In the case of ListOps, we also balance across the number of arguments in



Figure 3: Performance of the Neural Data Router, GPT-4 and Neural Rewriting System on ListOps, arithmetic and algebraic formulas.

Solver The training and validation set of the Solver included two kinds of samples: leaf formulas, mapped to equivalent atomic elements, and atomic elements themselves, mapped to the end-of-computation special symbol ω . During training, we generated batches that included both kinds of samples with equal probability to avoid biasing the model toward solving leaf formulas.

5.2.2 Neural Data Router

The Neural Data Router (NDR) is a variant of a transformer encoder designed to learn to solve algorithmic problems with strong out-ofdistribution generalization capacity. In terms of the formalization of formula simplification problems we have defined in Section 3, the model takes as input a formula f and gives as output the equivalent atomic element e. The NDR thus implements the composition of the *sel*, *sol*, and *com* functions in one or more transformer layers, assuming that the effective integration between the three happens sub-symbolically in the embedding space during training.

This model has been previously evaluated on relatively simple algorithmic benchmarks that are nevertheless significantly similar to the problems we consider, including the solution of formulas in the original ListOps dataset and simple arithmetic formulas on singledigit integers. The main difference with the problems we consider here is that the operands in our Arithmetic and Algebra problems, we aim to generalize both on the number of operands and on the nesting depth of the formulas, whereas, in the original work, the number of operands observed at training time was the same as that observed at test time. Furthermore, we use significantly fewer samples during training and of lower complexity, as the NDR was originally trained on arithmetic and ListOps formulas with up to 5 nested operations.

We applied a minor modification to the architecture to adapt the model to the problems we have considered. Indeed, the problems considered in the original work always had a single-digit integer as a result, which the model was trained to output as the first token in the sequence given as output by the encoder. Since, in general, this no longer applies to our problems, we read the final answer from the first k positions of the sequence produced by the encoder, where k is the maximum length of a problem's targets.

We build all development sets for the NDR using the same toplevel formulas we included in the analogous sets for the Selector. Since the NDR was designed to compute intermediate results subsymbolically, we use as training target the final simplified form of the input formula and we do not include intermediate formulas as we do with the Neural Rewriting System. ⁴ Following the original experimental protocol, we made the training set balanced across nesting levels and number of operands. As done when training the Selector module, we create in- and out-of-distribution validation sets and use the latter to optimize hyperparameters, using the Weights and Biases Bayesian search on the same hyperparameters and intervals described in the original work. We report the final hyperparameter values for each task in the Technical Appendix.

5.2.3 GPT-4

Chain-of-Though (CoT) prompting improves the performance of LLMs on reasoning tasks by enabling step-by-step solution procedures. In terms of the formalization outlined in Section 3, we may observe how such reasoning chains resemble the iterative solution method of nested formulas based on rewriting. Therefore, we can say that LLMs prompted using CoT-like methods loosely implement the composition of *sel*, *sol*, and *com* functions in the whole model, where the iterative application of the function is implemented via the auto-regressive generation of outputs.⁵

We choose to prompt GPT-4 using self-consistency prompting [32] combined with zero-shot Chain-of-Thought (CoT) [12]. Zero-shot CoT has been introduced as a simpler alternative to CoT prompting, which allows users to achieve similar performance on reasoning benchmarks without the need to engineer exemplars for few-shot

⁴ For comparison, in the Technical Appendix we show results for models trained on larger datasets that also include intermediate formulas.

⁵ Notice that in general, at any given point of the solution process, the context given as input of the model will contain several different simplifications of the input function f rather than just one as it happens with the other models we consider.

reasoning. This is achieved by simply making the model's answer start with the sentence: "Let's think step-by-step". After the model has generated a response, it is prompted again to retrieve a wellformatted output. Self-consistency prompting operates on the concept that reasoning problems can have multiple valid paths leading to the same conclusion. To harness this, we generate 10 outputs for each input, allowing us to select the most consistently produced one. This method boosts confidence in the model's output and significantly enhances accuracy, resulting in a notable improvement in performance.

The actual prompt used for zero-shot CoT was built giving the model a minimal description of the problem at hand and then asking to solve it. For example, the zero-shot CoT prompt corresponding to the ListOps input sample [MIN[SM54] [MIN39]] is: "MIN, MAX and SM are operators on lists of single-digit integers which have the semantics of minimum, maximum and sum modulo 10, respectively. Solve the following expression involving these operators: [MIN [SM 5 4] [MIN 3 9]].". Following the zero-shot CoT prompting technique, the model was then prompted a second time to extract the well-formatted final answer.

5.3 Results

We report in Figure 3 the performances of all considered models. All models are tested on 100 samples for each level of complexity. We measure the sequence accuracy of the models, that is the exact match between the model's output and the target sequence. When we evaluate GPT-4 on the Algebra task, we take into account equivalent forms of the symbolic outputs using the SymPy library [20].

We first observe that all the models perform better on simple data splits and tend to worsen on more complex ones. This indicates that increasing the number of arguments and the nesting of formulas makes the problems more difficult for all models, as expected.

Analyzing the performance of the Neural Data Router, we find that the models with the best hyperparameters configurations have lower accuracy than other models on the simple splits and are not able to generalize the learned solution process to complex problem instances. Learning to solve simple formulas with a single nesting level - i.e., learning the sol function - appears to be the easiest task for the model, although it is learned with almost perfect accuracy only on arithmetic formulas. The lack of generalization capacity on complex formulas may, therefore, be determined by two different situations: making one or more errors while solving many leaf formulas, or learning a composition of the sel, sol, and com functions that cannot generalize effectively on nested formulas. Since the model is unable to generalize both on arithmetic and algebraic formulas, whose leaf formulas can be solved with substantially different levels of accuracy, we argue that both factors could play a role in causing the limited generalization capacity of the model.

We observe that GPT-4 can achieve perfect or almost perfect results on leaf formulas and simpler problem instances, surpassing the performance of the Neural Data Router and, in some cases, our architecture. However, the performance of GPT-4 is significantly lower on deeply nested formulas. Indeed, while self-consistency prompting improved the model's performance compared to vanilla zero-shot CoT prompting⁶, it has a limited impact on the capacity of the LLM to solve complex formulas, consistently with previous research [26].

Considering the model as an approximation of the composition of the *sel*, *sol*, and *com* functions, its limited generalization capability



Figure 4: Input length against the average confidence score of 1,000 outputs of the Selector. The vertical lines represent the maximum input length in the training set.

seems to be determined by the poor quality of this approximation obtained with the prompting method we used. However, by comparing the performance of the model on ListOps formulas with the other two tasks, we observe that the complexity of the individual leaf formulas (much simpler in the case of ListOps) seems to influence the capability of the model to apply a solution method that can generalize systematically on complex formulas.

Coherently with the model selection process, in Figure 3, we report the performance of our Neural Rewriting System with 1,000-way multi-output generation, and value of the Dynamic Windowing threshold T of 60, 150 and 200 for the ListOps, Arithmetic and Algebra problems, respectively. The Neural Rewriting System consistently outperforms the baselines on the three formula simplification tasks we consider. Importantly, the system shows a significantly high degree of generalization on formulas that are much more complex than those observed during training. Given the modular nature of our architecture, in the following paragraphs, we will analyze the role of each architecture component in achieving this result.

We first analyze the variation of the confidence score corresponding to variations in the length of the input sequences in the test set. We do so because the confidence score plays an important role in selecting the best Selector output in multi-output generation. Figure 4 represents input lengths versus the average confidence score of 1,000 Selector outputs. The vertical lines in the plot represent the maximum input length seen during training. While the average confidence score remains close to the training range for outputs that are moderately longer than the longest training samples, it drops significantly for sequences much longer than that. This can directly determine a decrease in model performance on very long formulas since Selector outputs with a low confidence score are more likely to be incorrect. Indeed, we set the value of the Dynamic Windowing threshold T, which determines the sequences on which the mechanism will be applied, by looking at the points where Selector confidence drops in each task. We also notice that confidence scores are lower for very

⁶ We report the performance of GPT-4 with vanilla zero-shot CoT prompting in the Technical Appendix for comparison.



Figure 5: Errors committed by the Neural Rewriting System classified by cause of error. Error bars corresponding to a complexity split refer to models where the Selector generates 10, 100 or 1,000 outputs per input and (optionally) employs the Dynamic Windowing mechanism. Data splits are defined by their nesting level N and number of arguments A.

short arithmetic and algebraic formulas. We suppose this is due to the fact that the number of formulas of a given nesting level grows exponentially with the value of the nesting level. Therefore, the Selector observed fewer examples of formulas with nesting level 1 during training. However, this does not hinder performance when the Selector generates many outputs per input.

We finally assessed the impact of the multi-output generation strategy on the model's performance by comparing the errors committed by the model when the Selector generated 10, 100 and 1,000 outputs per input. We measure the errors separately on subsets of formulas with different nesting levels and number of operands, to understand the impact of increasing the number of outputs generated by the model on the solution of both simple and complex formulas. Furthermore, for each model configuration and each subset of samples, we analyzed the root cause of the errors committed by the NRS. We classify the cause of the errors into three categories: errors caused by the Selector when it outputs a leaf formula that is not present in the input formula (missing-leaf); errors caused by the Selector when it outputs a leaf formula that is present in the input formula but does not correspond to a well-formed expression and whose substitution thus causes the corruption of the input formula (corrupted-leaf); and errors caused by the Solver when it outputs a wrong solution for the leaf formula identified by the Selector (wrong-solution). Figure 5 represents this multi-faceted measurement of the errors committed by the models.

Looking at the frequency of each of these classes of errors, we clearly see that generating only 10 outputs per input with the Selector is particularly ineffective on algebraic formulas, and on deeply nested arithmetic and ListOps formulas. Increasing the number of generated outputs to 100 or 1,000 considerably improves the System's performance on most formulas of all three tasks, apart from the most deeply nested arithmetic and algebraic formulas. On such samples, applying Dynamic Windowing drastically reduces the number of errors the model commits while also bringing further minor

improvements on simpler formulas.

As we described in Section 4, the modules in the Neural Rewriting System have been designed as neural implementations of the *sel*, *sol* and *com* functions that appear in the simple solution algorithm for formula simplification problems described in Section 3. For this reason, we can analyze in a greater level of detail to which modules are most error due, and thus which of the functions listed above have been more difficult to reproduce in a neural network.

Given our design choice to implement the *com* function using a CNN dynamically parameterized by the Selector, the errors due to this module can be directly traced back to errors made by the Selector. Looking at the classification of errors represented in Figure 5, we can see how the errors due to a wrong-solution of a leaf formula by the Solver are the minority. Therefore, the *sol* function appears to be relatively easy to implement in a transformer, consistent with our prior observations on NDR and GPT-4 models. Errors due to the Selector are instead the vast majority and are only partially mitigated by the generation of multiple outputs and the Dynamic Windowing, especially on algebraic formulas. The *sel* function, which formalizes the ability to identify leaf expressions, appears to be the hardest one to model with a transformer-based architecture, particularly due to the necessity to generalize on very long formulas.

6 Discussion

In this work, we have considered formula simplification problems, a kind of benchmark used to study the systematic generalization capabilities of neural networks. We proposed a formal description of this class of problems and used it to define a simple general solution algorithm involving three elementary functions. We then implemented a general neural architecture called the Neural Rewriting System, which models such algorithm by adopting a modular design.

We have compared the performance of the Neural Rewriting System to that of a highly specialized neural architecture, the Neural Data Router, and of GPT-4, a powerful general-purpose large language model. We have tested the models on nested arithmetic and algebraic formulas, and formulas from the ListOps dataset. We showed that while both the Neural Data Router and GPT-4 are, to different extents, able to solve simple formulas, both models struggle or fail when dealing with deeply nested formulas. The Neural Rewriting System achieves a considerably greater degree of generalization on much more complex formulas than those it was trained on. We analyzed models' failures in light of our formalization of formula simplification problems, highlighting the features of these problems that could represent the main obstacles to systematic generalization.

Limitations While we consider formula simplification problems with different number and kinds of operands, we also pose several constraints on these formulas. For instance, we limit our study to formulas with only one final form which can always be reached independently of the substitution applied. Both of these conditions are not satisfied, in general, by algebraic expressions, whose form could be made arbitrarily complex and not further reducible. Future work should thus study the systematic generalization capabilities of neural models in the more general case of formulas with several possible final forms, reachable from different solution paths.

7 Acknowledgements

We thank OpenAI for granting research access to the GPT APIs.

References

- V. Agarwal, S. Aditya, and N. Goyal. Analyzing the nuances of transformers' polynomial simplification abilities. In *MATH-AI Workshop*. International Conference on Learning Representations, May 2021.
- [2] C. Anil et al. Exploring length generalization in large language models. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022, 2022.
- [3] C.-H. Cai et al. Learning of human-like algebraic reasoning using deep feedforward neural networks. *Biol. Inspired Cogn. Arch.*, 25:43–50, Aug. 2018.
- [4] X. Chen and Y. Tian. Learning to perform local rewriting for combinatorial optimization. In H. M. Wallach et al., editors, Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada, pages 6278–6289, 2019.
- [5] R. Csordás, K. Irie, and J. Schmidhuber. The devil is in the detail: Simple tricks improve systematic generalization of transformers. In M. Moens, X. Huang, L. Specia, and S. W. Yih, editors, *Proceedings* of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021, pages 619–634. Association for Computational Linguistics, 2021. doi: 10.18653/V1/2021.EMNLP-MAIN.49.
- [6] R. Csordás, K. Irie, and J. Schmidhuber. CTL++: evaluating generalization on never-seen compositional patterns of known functions, and compatibility of neural representations. In Y. Goldberg, Z. Kozareva, and Y. Zhang, editors, Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022, pages 9758–9767. Association for Computational Linguistics, 2022. doi: 10.18653/V1/2022. EMNLP-MAIN.662.
- [7] R. Csordás, K. Irie, and J. Schmidhuber. The neural data router: Adaptive control flow in transformers improves systematic generalization. In *The Tenth International Conference on Learning Representations, ICLR* 2022, Virtual Event, April 25-29, 2022, 2022.
- [8] N. Dershowitz and J.-P. Jouannaud. Rewrite systems. In Formal models and semantics, pages 243–320. Elsevier, 1990.
- [9] N. Dziri et al. Faith and fate: Limits of transformers on compositionality. In *Thirty-seventh Conference on Neural Information Processing* Systems, 2023.
- [10] J. A. Fodor and Z. W. Pylyshyn. Connectionism and cognitive architecture: a critical analysis. *Cognition*, 28(1-2):3–71, Mar. 1988.
- [11] D. Hupkes, V. Dankers, M. Mul, and E. Bruni. Compositionality decomposed: How do neural networks generalise? *Journal of Artificial Intelligence Research*, 67:757–795, 2020.
- [12] T. Kojima et al. Large language models are zero-shot reasoners. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022, 2022.
- [13] E. Komendantskaya. Parallel rewriting in neural networks. In Proceedings of the International Joint Conference on Computational Intelligence (IJCCI 2009) - ICNC, pages 452–458. INSTICC, SciTePress, 2009. ISBN 978-989-674-014-6. doi: 10.5220/0002319704520458.
- [14] B. M. Lake and M. Baroni. Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks. In J. G. Dy and A. Krause, editors, Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018, volume 80 of Proceedings of Machine Learning Research, pages 2879–2888. PMLR, 2018.
- [15] B. M. Lake and M. Baroni. Human-like systematic generalization through a meta-learning neural network. *Nature*, 623(7985):115–121, Nov. 2023.
- [16] G. Lample and F. Charton. Deep learning for symbolic mathematics. ArXiv, abs/1912.01412, 2019.
- [17] Y. Li and J. L. McClelland. Systematic generalization and emergent structures in transformers trained on structured tasks. In *All Things Attention: Bridging Different Perspectives on Attention*. Annual Conference on Neural Information Processing Systems, 2022.
- [18] A. Liska, G. Kruszewski, and M. Baroni. Memorize or generalize? searching for a compositional RNN in a haystack. *CoRR*, abs/1802.06467, 2018.
- [19] T. Marghetis, D. Landy, and R. L. Goldstone. Mastering algebra retrains the visual system to perceive hierarchical structure in equations. *Cognitive research: principles and implications*, 1:1–10, 2016.

- [20] A. Meurer et al. Sympy: symbolic computing in python. *PeerJ Computer Science*, 3:e103, Jan. 2017. ISSN 2376-5992. doi: 10.7717/peerj-cs.103.
- [21] S. Mittal, Y. Bengio, and G. Lajoie. Is a modular architecture enough? In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022, 2022.
- [22] N. Nangia and S. R. Bowman. Listops: A diagnostic dataset for latent tree learning. In S. R. Cordeiro, S. Oraby, U. Pavalanathan, and K. Rim, editors, Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 2-4, 2018, Student Research Workshop, pages 92–99. Association for Computational Linguistics, 2018. doi: 10.18653/V1/N18-4013.
- [23] S. Ontañón, J. Ainslie, Z. Fisher, and V. Cvicek. Making transformers solve compositional tasks. In S. Muresan, P. Nakov, and A. Villavicencio, editors, *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022,* pages 3591–3607. Association for Computational Linguistics, 2022. doi: 10.18653/V1/2022.ACL-LONG. 251.
- [24] OpenAI. Gpt-4 technical report, 2023.
- [25] F. Petruzzellis. flavio2018/itersolv-public: First release, 2024. URL https://zenodo.org/doi/10.5281/zenodo.12734121.
- [26] F. Petruzzellis, A. Testolin, and A. Sperduti. Benchmarking GPT-4 on algorithmic problems: A systematic evaluation of prompting strategies. Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation, LREC-COLING 2024, Turin (Italy), May, 20-25, 2024, 2024.
- [27] F. Petruzzellis, A. Testolin, and A. Sperduti. A neural rewriting system to solve algorithmic problems, 2024. URL https://arxiv.org/abs/2402. 17407.
- [28] O. Press et al. Measuring and narrowing the compositionality gap in language models. In H. Bouamor, J. Pino, and K. Bali, editors, *Findings* of the Association for Computational Linguistics: EMNLP 2023, pages 5687–5711, Singapore, Dec. 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.378.
- [29] A. Ruoss et al. Randomized positional encodings boost length generalization of transformers. In A. Rogers, J. L. Boyd-Graber, and N. Okazaki, editors, Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), ACL 2023, Toronto, Canada, July 9-14, 2023, pages 1889–1903. Association for Computational Linguistics, 2023. doi: 10.18653/V1/2023. ACL-SHORT.161.
- [30] A. Testolin. Can neural networks do arithmetic? a survey on the elementary numerical skills of state-of-the-art deep learning models. *Applied Sciences*, 2024.
- [31] A. Vaswani et al. Attention is all you need. In I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, and R. Garnett, editors, Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA, pages 5998– 6008, 2017.
- [32] X. Wang et al. Self-consistency improves chain of thought reasoning in language models. In *The Eleventh International Conference on Learn*ing Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023, 2023.
- [33] J. Wei et al. Chain-of-thought prompting elicits reasoning in large language models. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 -December 9, 2022, 2022.