

EXECUTING ARITHMETIC: FINE-TUNING LARGE LANGUAGE MODELS AS TURING MACHINES

Anonymous authors

Paper under double-blind review

ABSTRACT

Large Language Models (LLMs) have demonstrated remarkable capabilities across a wide range of natural language processing and reasoning tasks. However, their performance in the foundational domain of arithmetic remains unsatisfactory. When dealing with arithmetic tasks, LLMs often memorize specific examples rather than learning the underlying computational logic, limiting their ability to generalize to new problems. In this paper, we propose a Composable Arithmetic Execution Framework (CAEF) that enables LLMs to learn to execute step-by-step computations by emulating Turing Machines, thereby **achieving a true mastery** of computational logic. Moreover, the proposed framework is highly scalable, allowing composing learned operators to significantly reduce the difficulty of learning complex operators. In our evaluation, CAEF achieves nearly 100% accuracy across seven common mathematical operations on the LLaMA 3.1-8B model, effectively supporting computations involving operands with up to 100 digits, a level where GPT-4o falls short noticeably in some settings.

1 INTRODUCTION

Large Language Models (LLMs) have made significant strides in recent years, showcasing extraordinary capabilities across a range of natural language processing (NLP) tasks (Dubey et al., 2024; Jiang et al., 2024; Chowdhery et al., 2023), and in some cases, even surpassing human performance in specific benchmarks (Achiam et al., 2023). However, despite these advancements, LLMs still face significant challenges in performing arithmetic. Current research indicates that when presented with arithmetic problems, LLMs often rely on memorizing specific expressions and their corresponding outcomes rather than grasping the fundamental logic of arithmetic operations (Wu et al., 2023b). This inherent limitation poses a substantial barrier to their effective application in fields that demand essential computational skills.

To enhance the performance of LLMs in solving arithmetic problems, two primary approaches have been developed. The first approach positions the LLM as an agent that relies on an external calculator to perform computations (Hao et al., 2024; Ruan et al., 2023). In this setting, the LLM’s role is limited to providing the operands and invoking the appropriate operations. Although this method effectively simplifies the challenge of arithmetic for LLMs, it misses the opportunity for the models to learn computational logic, preventing LLMs from comprehending the underlying principles of arithmetic. Given that arithmetic serves as the foundation of mathematics, the lack of **arithmetic ability** may significantly impede the LLM’s capability to grasp more complex mathematical concepts. The second approach focuses on stimulating the LLM’s intrinsic capabilities, employing prompt engineering or fine-tuning techniques to enable the model to master arithmetic computations and solve problems through reasoning (Kojima et al., 2022; Huang et al., 2022; Yu et al., 2023). This approach typically involves the LLMs generating intermediate steps before reaching a final result.

Although the second approach is promising, it faces two significant challenges. The first challenge is that, under simple supervised fine-tuning, LLMs tend to memorize examples from the training set (Hu et al., 2024). As the length of the operands increases, the sample space expands exponentially, making it impractical for the LLM to memorize all possible examples. To fundamentally overcome this limitation, LLMs should primarily *learn and execute computational logic*, mirroring how humans systematically master arithmetic, rather than relying on memorization.

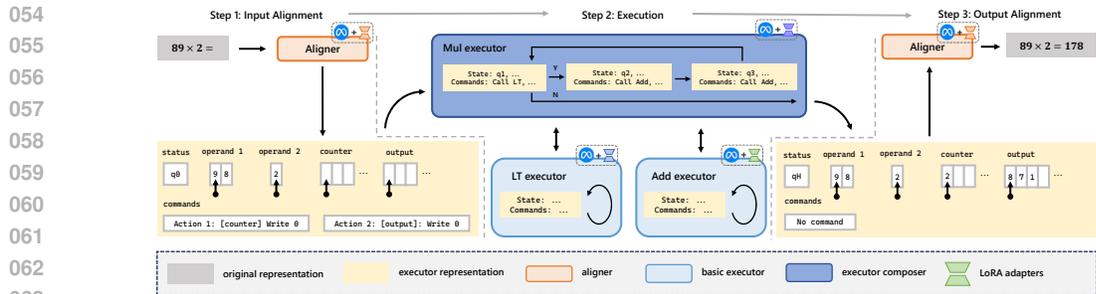


Figure 1: An illustrative CAEF flowchart demonstrates the execution of the *Multiplication* operation for 89×2 . The aligner converts the original arithmetic expression into a Turing Machine-like representation that the *Multiplication* executor can process. Acting as an executor composer, the *Multiplication* executor calls upon two basic executors, i.e., *Less_than* and *Addition*, to perform the actual computation. All the executors and the aligner are executed by the LLM.

The second challenge involves learning how to compose basic operators to build complex arithmetic operators. These complex operators are typically execution procedures that contain conditional statements (*if-then-else*) and iterative statements (*loop*), with the basic operators treated as function calls within these procedures. By doing this, LLMs could gradually learn more complex arithmetic operations by [focusing on their execution logic and calling the existing operators as necessary](#).

Mastering the execution of arithmetic is fundamentally equivalent to modeling computation. One famous mathematical model of computation is the Turing machine, which is formally introduced by Alan Turing (Turing et al., 1936). [If the LLM learns to execute computational logic by simulating executing a Turing machine based on its transition functions for each operator, it could solve arithmetic problems through a multi-query approach](#). This approach involves the LLM iteratively performing computations based on the current state and command, and then generating the next state and command.

In this paper, we propose a Composable Arithmetic Execution Framework (CAEF) for LLMs to solve arithmetic problems solely. Inspired by the Turing machine, CAEF aims to teach LLMs the computational logic, enabling them to *execute* the logic for specific arithmetic operators and *compose* arithmetic operators into more complex ones. CAEF has two key characteristics:

Executing arithmetic. As illustrated in Figure 1, CAEF employs a three-step procedure for each arithmetic operator, supported by two independent components within the LLM: the *executor* and the *aligner*. The executor, responsible for performing the actual computations, learns the underlying computational logic by modeling the transition function of the corresponding arithmetic Turing machine. This allows the LLM to iteratively generate intermediate results and ultimately produce the final output. The aligner serves as an interface, converting raw arithmetic expressions (e.g., $89 \times 2 =$) into a format that the executor can directly process. Once the executor completes its execution, the aligner transforms the executor’s output back into the final result. In our framework, both the executor and the aligner are implemented as separate LoRA adapters (Hu et al., 2021).

Composing operators. Complex operators can often be composed of basic or simpler ones, hierarchically or recursively. In CAEF, we design an *executor composer* that is responsible for the high-level execution procedures of complex operators and allows function calls to invoke other pre-learned arithmetic operators. Since each operator is implemented as a LoRA adapter, function calls in CAEF are executed by automatically switching LoRA adapters, following the LLM’s generated command. Thus, CAEF could facilitate the handling of more intricate computations.

Using the proposed framework, we have implemented seven operators: $+$, $-$, \times , \div , $>$, $<$, and $=$, along with two auxiliary operators (refer to Appendix A.4). Each of these operators is based on existing computational logic, such as the Turing machine or algorithms used in CPU design (e.g., the subtraction operator is modeled similarly to how modern CPUs handle the subtraction operation.). [Our experiments show that CAEF achieves high accuracy across all seven operators when using the LLaMA 3.1-8B model \(Dubey et al., 2024\)](#). Compared to GPT-4o, the LLM equipped with CAEF demonstrates minimal impact from changes in operand length, effectively supporting computations involving operands with up to 100 digits. [The main contributions of this paper are as follows:](#)

- We propose a framework CAEF enabling LLM learning to execute the computational logic of operators by imitating the execution of Turing machine. Also, CAEF can naturally support composing multiple learned operators for operators with complex logic.
- We implement executors and aligners for seven arithmetic operators based on the proposed framework. The executor is responsible for performing the step-by-step computations iteratively, while the aligner serves as an interface, facilitating the bidirectional conversion between the internal representation of the executor and the original representation.
- The extensive evaluation shows that CAEF outperforms the existing LLMs with seven classic arithmetic tasks. The proposed CAEF enables the LLM to achieve almost 100% accuracy when operands are up to 100 digits.

2 APPROACH: FRAMEWORK DESIGN

2.1 PROBLEM STATEMENT

Computational logic is fundamental to arithmetic. To truly master arithmetic, the LLM should learn and execute the underlying computational logic of arithmetic operations rather than merely memorizing examples of arithmetic expressions. For scalability, the LLM should be capable of constructing new operators by combining existing operators. For example, after learning *Addition* operation, the LLM could construct *Multiplication* by [learning](#) the computational logic of repeated addition could achieve multiplication.

Therefore, we need a framework that enables LLM to model arithmetic operators by learning to execute their underlying computational logic. In the field of automata, the Turing machine provides a suitable framework for describing this logic. [Following the examples \(e.g., Turing machines introduced in Sipser \(1996\)\), we could build a Turing machine for common arithmetic operations, which can be a reference](#) to create adequate datasets of execution steps for LLM training. Furthermore, the Turing machine inherently supports the combination of multiple Turing machines, making it ideal for constructing complex operations from existing ones. By emulating Turing machines, LLM can be designed to integrate multiple models, enabling it to execute more intricate arithmetic tasks.

2.2 LLM EXECUTES AS TURNING MACHINE

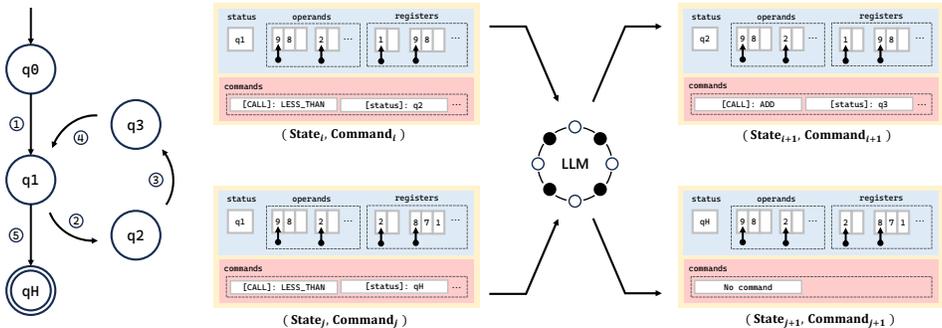
A Turing machine can be formally defined as a septuple $T = (Q, \Sigma, \Gamma, b, q_0, F, \delta)$, where Q is a finite set of states, $\Sigma \subseteq \Gamma$ is a finite alphabet for input, Γ is a finite tape alphabet, $b \in \Gamma$ is the blank symbol, $q_0 \in Q$ is the initial state, $F \subseteq Q$ is a set of final states, and δ is the transition function. When a Turing machine is in a non-halting state, the next action is determined by both the current symbol on the tape and the machine’s current state. In each action, the machine updates the symbol on the tape, transitions to a new state, and moves the tape head either left or right. This process repeats iteratively until the machine reaches a halting state, at which point the computation is complete, and the result is saved on the tape.

LLM is the generative model for text-based language, so how to transfer all information from a Turing machine to the LLM effectively is challenging. A tailored representation system is necessary for LLMs to accurately [learn](#) computational logic. To facilitate this transfer, the system must incorporate states analogous to those of the Turing machine, such as the machine state and tape state, to indicate the current status of the computation, in other words, the step in the execution process. Additionally, the system should include commands that specify the actions to be executed based on the current state to ensure correct transitions to the next state. Thus, CAEF provides a text-based representation $\langle s_i, c_i \rangle$ that effectively represents the state s_i and the command c_i for each step i in the computational logic. [Then, the state transition function \$f\$ \(e.g. LLM or LLM fine-tuned with LoRA adapters\) could use this representation at step \$i\$ as the input to generate the next representation at step \$i+1\$ as following:](#)

$$s_{i+1}, c_{i+1} = f(s_i, c_i) \tag{1}$$

By formulating the representation of both input and output for Equation 1, the LLM is enabled to perform computations in a manner similar to that of a Turing Machine by *executing* step-by-step transitions.

162
163
164
165
166
167
168
169
170
171
172



173 Figure 2: Diagram of the CAEF framework. The CAEF representation includes two required
174 components: state and command, corresponding to areas and in the figure. The state
175 part records the current status, operands, and registers that store intermediate variables and results,
176 etc. The command consists of a set of actions, such as write operations and call operations. Upon
177 receiving the state and command, the LLM generates the next state and the corresponding command,
178 with each step corresponding to a transition in the state diagram on the left.

179
180

181 **2.3 REPRESENTATION DESIGN**

182
183
184
185
186
187
188
189
190
191

In this paper, we design a structured representation for arithmetic problems to enable the LLM to accurately execute computational logic. As illustrated in Figure 2, representation of the arithmetic problem includes: 1) a status indicating the current step of the computation, and 2) a "tape" that records all operands and essential information, such as the number of digits processed, any carryover during addition, and other intermediate results. To facilitate the LLM’s learning of the execution process, the representation in CAEF explicitly includes the commands c required for execution. These commands involve the *call* to the next *state* s and other detailed actions, such as carrying over or moving the pointer. All the above elements are represented in text, which is friendly to LLM to deal with. Then, to make LLM execute based on the representation, CAEF structures the input as $\langle s_i, c_i \rangle$ for current step i , while the output is $\langle s_{i+1}, c_{i+1} \rangle$ for the next step $i+1$.

192
193
194
195
196
197
198
199

Besides modeling the representation, representation translation is another critical part of CAEF. In general, the original input of an arithmetic problem does not include the initial state or the first command to execute. Moreover, upon completing the computation, the raw output remains in the representation format. Thus, CAEF incorporates an aligner to manage the bidirectional translation between the original input/output and the representation. The aligner can also be implemented by fine-tuning a specific LoRA adapter. Notice that one key feature of the aligner should learn the ability to convert the left-to-right (L2R) representation of numbers into a right-to-left (R2L) format, as R2L is evaluated more effectively for LLM to operate the operands (Lee et al., 2023).

200
201

3 APPROACH: IMPLEMENTATION

202
203
204

Building on the conceptual design of CAEF, we present the detailed implementation of Equation 1, highlighting two key derived components: basic executors and executor composers with examples.

205
206

3.1 FINE-TUNING PROCESS AND IMPLEMENTATION DESIGN

207
208
209
210
211

CAEF offers a framework to enhance the arithmetic capabilities of LLMs. To implement Equation 1 for a specific arithmetic task, CAEF involves the following steps: 1) step 1: design a state machine and implement the representation $\langle s_i, c_i \rangle$ for the arithmetic task, and 2) step 2: sample pairs of input and output to create a dataset, which is then used to fine-tune the LLM for one-step execution.

212
213
214
215

Step 1. Designing a state machine can draw from existing Turing machines or other relevant state machines for the task. To implement state s_i and commands c_i in the representation, we transform the structured representation into plain text following two main guidelines: 1) numbers are formatted in R2L order, separated by |, and 2) each command is expressed in the format $\{[CMD] action\}$. For example, for the addition task $45 + 67 =$ where the current step involves adding the tens digits

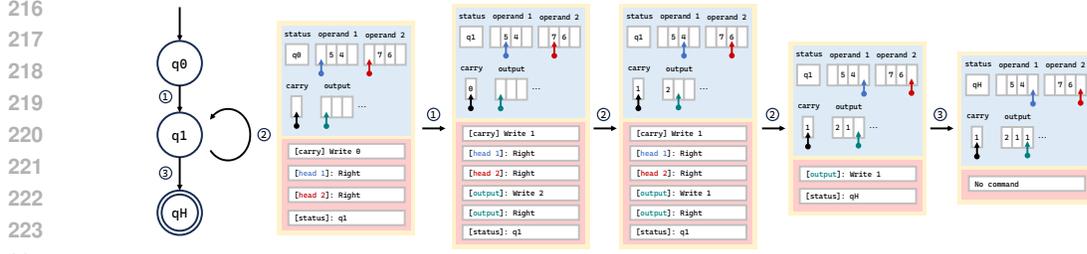


Figure 3: Execution process of $45 + 67$. The state diagram on the left abstracts the addition process. In step ②, a one-digit addition is performed, followed by updating the carry and output. The right side shows the actual sequence of state and command execution in the CAEF framework.

with a carry of 1 from the units place, the representation $\langle s_i, c_i \rangle$ may include several pointers: two HEADs pointing to the digits, a carry C for the carry value, and OUTPUT to record the results. All these pointers are moved using the RIGHT command. The representation is written as follows:

$$s_i : \text{ADD}, q1, | 5[\text{HEAD1}] | 4 | 7[\text{HEAD2}] | 6 [C]1 | 2[\text{OUTPUT}]$$

$$c_i : \text{CMD}: [C] 1, [\text{OUTPUT}] 1, [\text{OUTPUT}] \text{RIGHT}, [\text{HEAD1}] \text{RIGHT}, [\text{HEAD2}] \text{RIGHT}, q1$$

where $q1$ indicates the current status, and all pointers are presented in uppercase, enclosed in brackets with the pointed value on their right.

Step 2. To fine-tune the LLM, the dataset, including input and output representation pairs used for learning one-step execution. Continuing with the example, we create the representation $\langle s_{i+1}, c_{i+1} \rangle$ for the output of the one-step execution based on the above $\langle s_i, c_i \rangle$:

$$s_{i+1} : \text{ADD}, q1, | 5 | 4[\text{HEAD1}] | 7 | 6[\text{HEAD2}] [C]1 | 2 | 1[\text{OUTPUT}]$$

$$c_{i+1} : \text{CMD}: [\text{OUTPUT}] 1, [\text{OUTPUT}], [C], qH$$

where qH is the halting status. The details of the dataset refer to Section 4.1.

One efficient way for LLM to learn for one-step execution is LoRA fine-tuning. Since we target to learn $+$, $-$, \times , \div , $>$, $<$, and $==$ arithmetic operators, implementing multiple LLM instances leads to significant memory overhead. To mitigate this, we use a single base LLM model with multiple LoRA adapters that serve as learned executable arithmetic operators. Switching LoRA adapters based on function calls generated by the LLM can efficiently perform various operations, optimizing memory usage while maintaining flexibility in handling different arithmetic computations.

To implement a specific computational task, CAEF introduces two types of executors (i.e., *basic executor* and the *executor composer*) to learn to execute under the proposed representation. The basic executor is designed to handle tasks with well-defined computational logic, while the executor composer acts as a higher-level controller that orchestrates the process by calling other basic executors. In the following, we introduce the two types of executors through illustrative examples.

3.2 BASIC EXECUTORS

We use *addition* to illustrate the design of a basic executor. A natural way to implement addition is to imitate the accumulator, performing the addition of two corresponding digits from the operands once at a time, along with the value stored in the carry register. This process calculates the result for the current digit and simultaneously updates the carry register for the next higher digit's computation.

Thus, the state and the command for addition are constructed as follows. The state should include the following components: 1) the two operands, 2) two pointers indicating the current digits being processed, 3) the carry register, and 4) the output generated so far. The command part should at least include: 1) the actions to write the carry and output, 2) the actions to move the pointers, and 3) state transition actions to control the start, transitions, and halting of the addition. Based on this instruction, CAEF constructs the state machine based on the text-based represented $\langle s_i, c_i \rangle$. Figure 3 illustrates the computation process of CAEF for *addition*. The details of computations and dataset are listed in Appendix A.3 and Section 4.1, respectively. In this paper, we use similar procedures to design the operators for $>$, $<$ and $==$.

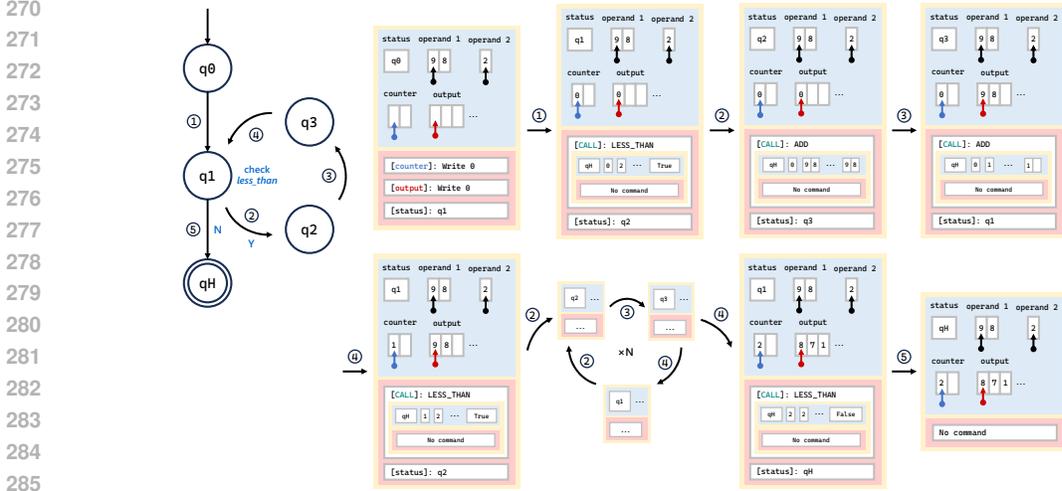


Figure 4: Execution process of 89×2 . The state diagram on the left abstracts the multiplication process, where in state q_1 , the less-than executor is performed. If true, the execution moves to state q_2 ; otherwise, it transitions to state q_5 and halts. Steps ③ and ④ execute the accumulation of the counter and output, respectively. The right side shows the actual execution in the CAEF framework.

3.3 EXECUTOR COMPOSERS

Executor composer designs to orchestrate the basic executors into intricate computational logic. Instead of performing computations directly, the executor composer "calls" other basic executors in a specific sequence to accomplish more complex tasks.

Multiplication is a typical example of the executor composer, which can be implemented by calling the $+$ and $<$ basic executors. CAEF uses two accumulators ($+$ involved) to implement $a \times b$. The first accumulator increments by 1 with each loop iteration, while the second adds a during each iteration. This process continues until the first accumulator reaches b ($<$ involved), and then the value in the second accumulator represents the final result. LLM is fine-tuned to execute control flow and loops, by calling the $<$ executor and, based on its result, either halts or continues the loop. Figure 4 illustrates the computation process for 89×2 using our implementation. Since the executor composer decouples the computational logic into several executors, the fine-tuning process could be done separately for each executor, showing the ability of executor composition.

Besides *multiplication*, we also design *subtraction* (considering only non-negative results) and *division* (floor division) executor composers using similar methodologies. Specifically, we draw inspiration from how subtraction is handled in CPUs to construct the *subtraction* executor composer and the detailed implementation can be found in Appendix A.4.

4 EVALUATION

4.1 SETTING

Models. We utilize the LLaMA 3.1-8B pretrained model (non-instruct version) as the base model. During LoRA fine-tuning, all linear modules in the decoder layer are involved in training, with the hyperparameters fixed at $r = 8$, $\alpha = 16$, and a learning rate of 5×10^{-5} . The fine-tuning process is conducted in two stages. In the first stage, we introduce an exhaustive explanation in the prompt, detailing the computation goal of an executor, the required input/output format, and providing an example. This explanation is followed by the actual sample, as illustrated in Appendix A.1. In the second stage, we remove the long explanation from the prompt and present only the sample, expecting the model to predict the next state and the subsequent commands directly. We use batch sizes of 8 and 16 for the first and second stages, respectively. All experiments are conducted on a server equipped with six H800 GPUs. The code and the models are available¹.

¹The implementation code is accessible at <https://github.com/HNDRXwjrmY/CAEF>, and the checkpoints are available at https://huggingface.co/HNDRXwjrmY/CAEF_llama3.1_8b

Table 1: Overall evaluation results across seven operators. "LLaMA 3.1 (L)" refers to the LLaMA fine-tuned with LoRA, while "LLaMA 3.1 (I)" refers to the LLaMA 3.1-8B-Instruct model.

Operator	Model	5-digits	10-digits	50-digits	100-digits	1~10-digits
Addition	CAEF	100.0	99.6	99.9	98.6	-
	LLaMA 3.1 (L)	92.1 ^{-7.9}	64.8 ^{-34.8}	0.0 ^{-99.9}	0.0 ^{-98.6}	-
	LLaMA 3.1 (I)	93.5 ^{-6.5}	35.0 ^{-64.6}	0.0 ^{-99.9}	0.0 ^{-98.6}	-
	GPT-4o	98.4 ^{-1.6}	94.0 ^{-5.6}	65.0 ^{-34.9}	43.0 ^{-55.6}	-
Subtraction	CAEF	98.7	99.5	98.8	98.0	-
	LLaMA 3.1 (L)	82.8 ^{-15.9}	61.0 ^{-38.5}	0.0 ^{-98.8}	0.0 ^{-98.0}	-
	LLaMA 3.1 (I)	92.6 ^{-6.1}	60.3 ^{-39.2}	0.0 ^{-98.8}	0.0 ^{-98.0}	-
	GPT-4o	98.6 ^{-0.1}	95.9 ^{-3.6}	84.0 ^{-14.8}	71.6 ^{-26.4}	-
Greater_than	CAEF	99.2	99.0	99.2	97.2	-
	LLaMA 3.1 (L)	93.0 ^{-6.2}	90.0 ^{-9.0}	46.3 ^{-52.9}	10.0 ^{-87.2}	-
	LLaMA 3.1 (I)	99.3 ^{+0.1}	97.7 ^{-1.3}	72.1 ^{-27.1}	70.0 ^{-27.2}	-
	GPT-4o	99.8 ^{+0.6}	99.6 ^{+0.6}	99.0 ^{-0.2}	93.2 ^{-4.0}	-
Less_than	CAEF	99.7	99.3	99.6	98.0	-
	LLaMA 3.1 (L)	96.2 ^{-3.5}	93.6 ^{-5.7}	84.0 ^{-15.6}	45.0 ^{-53.0}	-
	LLaMA 3.1 (I)	93.9 ^{-5.8}	86.3 ^{-13.0}	74.6 ^{-25.0}	67.4 ^{-30.6}	-
	GPT-4o	99.9 ^{+0.2}	100.0 ^{+0.7}	99.3 ^{-0.3}	89.2 ^{-8.8}	-
Equal	CAEF	99.4	99.6	99.1	98.4	-
	LLaMA 3.1 (L)	57.5 ^{-41.9}	66.2 ^{-33.4}	59.2 ^{-39.9}	54.0 ^{-44.4}	-
	LLaMA 3.1 (I)	100.0 ^{+0.6}	98.8 ^{-0.8}	99.6 ^{+0.5}	99.6 ^{+1.2}	-
	GPT-4o	100.0 ^{+0.6}	100.0 ^{+0.4}	100.0 ^{+0.9}	100.0 ^{+1.6}	-
Multiplication	CAEF	-	-	-	-	99.3
	LLaMA 3.1 (L)	-	-	-	-	61.8 ^{-37.5}
	LLaMA 3.1 (I)	-	-	-	-	61.4 ^{-37.9}
	GPT-4o	-	-	-	-	97.7 ^{-1.6}
Division	CAEF	-	-	-	-	99.3
	LLaMA 3.1 (L)	-	-	-	-	98.4 ^{-0.9}
	LLaMA 3.1 (I)	-	-	-	-	96.5 ^{-2.8}
	GPT-4o	-	-	-	-	99.1 ^{-0.2}

Baseline. We compare our approach against three baselines on +, -, ×, ÷, ==, >, and < operators. The first is a LLM fine-tuned with LoRA on LLaMA 3.1-8B (non-instruct version). Additionally, we include two unmodified LLMs, GPT-4o and LLaMA 3.1-8B Instruct, both of which directly generate the computational results based on the arithmetic expressions through a single model query. The prompts used for these models are in Appendix A.5.

Dataset. In CAEF, an operator requires an executor and an aligner, each supported by a specific LoRA adapter. To generate training datasets for these adapters, we implement a Turing machine prototype for each operator. For the executor, we generate random arithmetic expressions and run the Turing machine from its initial state until it halts, recording states and commands before and after each transition. This produces a sequence of states and commands, from which we sampled to train the executor. By generating multiple sequences through random initialization, an adequate training dataset for the executor can be created. It is notable that for arithmetic expressions with long operands, the sequences tend to be lengthy. Simple random sampling may lead to a dataset dominated by intermediate steps, potentially omitting samples from the first and final transitions. To address this, we ensure that the first and last steps are always included. Similarly, for the aligner, we generate two alignment processes: one aligning the original arithmetic expression with the executor’s initial state and first command, while another aligning the executor’s halt state with the final result of the original arithmetic expression.

For the test sets, we generate a dataset consisting of pure arithmetic expressions using predefined templates (refer to in Appendix A.2). Specifically, for +, -, ==, >, and < operations, we create test sets with two operands of equal length, consisting of 5, 10, 50, and 100 digits. For multiplication and division, to avoid excessively large values, we adjusted the data range based on the characteristics of these two operators. In multiplication of the form $a \times b = c$, we restricted a to be a random number

Table 2: Accuracy of the executor and aligner across seven operators. The executor’s accuracy refers to the probability of completing the entire computation correctly from the initial state to the final step, with each step being accurate. The aligner’s accuracy is divided into two parts: the conversion from the original input to the executor’s representation, denoted as aligner (I), and the conversion from the executor’s final representation to the output, denoted as aligner (O).

Operator	Component	5-digits	10-digits	50-digits	100-digits	1~10-digits
Addition	executor	100.0	100.0	99.9	99.6	-
	aligner (I)	100.0	99.7	100.0	99.6	-
	aligner (O)	100.0	99.9	100.0	99.4	-
Subtraction	executor	100.0	100.0	99.6	99.2	-
	aligner (I)	98.8	99.7	99.5	99.6	-
	aligner (O)	99.9	99.7	99.7	99.2	-
Greater_than	executor	100.0	100.0	99.8	99.6	-
	aligner (I)	99.2	99.1	99.4	98.6	-
	aligner (O)	100.0	99.9	100.0	99.2	-
Less_than	executor	100.0	100.0	100.0	100.0	-
	aligner (I)	99.8	99.3	99.7	98.4	-
	aligner (O)	99.9	100.0	99.8	99.6	-
Equal	executor	100.0	100.0	99.8	99.4	-
	aligner (I)	99.4	99.6	99.6	98.8	-
	aligner (O)	100.0	100.0	99.8	99.8	-
Multiplication	executor	-	-	-	-	99.5
	aligner (I)	-	-	-	-	99.8
	aligner (O)	-	-	-	-	100.0
Division	executor	-	-	-	-	99.4
	aligner (I)	-	-	-	-	100.0
	aligner (O)	-	-	-	-	99.9

with 1-10 digits, and b to fall within the value range $[1, 15]$. In division of the form $a \div b = c$, we constrained c to be within $[1, 15]$ and b to be a random number with 1-10 digits.

Metrics. We employ accuracy as the evaluation metric. Each arithmetic problem is computed once, and the result is compared with the ground truth using the Exact Match criterion.

4.2 MAIN RESULTS

Table 1 presents the evaluation results of our method and baseline models across the seven operators. Compared to the baselines, the proposed approach performs stably on all operators with high accuracy. Specifically for tasks with long numbers, such as 100-digit addition, LLM with CAEF effectively learns the computational logic to execute the addition process.

To further explore the actual performance of the executor and aligner during the computation process, we separately evaluate their accuracy on the same dataset. As the results shown in Table 2, we observe that even though the executor must generate numerous intermediate steps in an iterative manner, while the aligner only performs two conversion steps, the executor still outperforms the aligner overall. The executor achieves over 99% accuracy in all experimental settings, indicating that it has effectively learned the arithmetic logic. When provided with the correct initial state and command, it functions correctly in the vast majority of cases. On the other hand, the aligner shows lower accuracy when converting the original input compared to converting the executor’s output in most cases, suggesting that the bottleneck in the overall computation process lies in the reversal of operands, rather than in the computation itself. [Due to the page limit, more detailed analysis are presented in Appendix A.6, the analysis of computational complexity in CAEF is detailed in A.7, and an extended experiment exploring the merging of the aligner and executor for individual operators is presented in A.8.](#)

5 RELATED WORK

LLMs in Mathematical Contexts. Prior research has focused on enhancing LLM performance in mathematical tasks, often relying on external tools for calculations and primarily addressing math word problems rather than pure arithmetic. A common external tool is a calculator, as exemplified by Schick et al. (2024), which introduces a self-supervised method where the model learns when to call external tools via API access. [Similar strategies can be found in Gou et al. \(2023\) and He-Yueya et al. \(2023\), and it was employed in even earlier work \(Andor et al., 2019\).](#) Another tool is a programming language interpreter, such as Python, where the model generates code and an external interpreter executes it to obtain the result. A representative example is Lu et al. (2024), which treats the LLM as a planner that generates code and submits it to an external Python executor to handle math problems in tabular contexts. Wang et al. (2023) employs supervised fine-tuning to improve code-based problem-solving, while Zhou et al. (2023) proposes a zero-shot prompting method to enable code-driven self-verification, thereby improving mathematical performance.

LLMs in Arithmetic Scenarios. Another series of work focuses solely on arithmetic, which we consider directly related to our research. The common characteristic of these studies is their effort to teach LLMs computational logic and improve calculation accuracy through step-by-step processes. Among these works, Nye et al. (2021) is an early and far-reaching study, predating the popular Chain-of-Thought (CoT) approach. It introduces a similar idea in the arithmetic domain, where the language model outputs intermediate steps to a buffer called a "scratchpad," significantly improving performance in integer addition. Hu et al. (2024) observes that transformers tend to approach arithmetic problems using "case-based reasoning" and proposes a Rule-Following Fine-Tuning technique that guides the model to execute calculations step by step. Zhou et al. (2024) combines four techniques (i.e., FIRE position encodings, Randomized position encodings, Reversed format (R2L format), and Index hints) to develop a new model that achieves a $2.5\times$ improvement in length generalization for two-integer addition.

6 LIMITATIONS

Prone to errors with repeated digit patterns. Both the executor and the aligner tend to generate incorrect steps when encountering patterns of repeated digits, such as sequences like "999..." where a single digit repeats, or "456456..." where multiple digits repeat. These errors typically manifest as extra or missing repetitions of the pattern. While this issue might be mitigated by intentionally generating more such expressions to increase the representation of similar samples in the training set, we believe the root cause lies in limitations inherent to generative LLMs.

Efficiency Issue. In our method, completing a single computation requires generating the full sequence of intermediate steps, which essentially means repeatedly calling the `model.generate()` function. For computations involving hundreds of steps, this process can be extremely time-consuming. One potential solution lies in optimizing the use of the KV cache. In our approach, the input to the LLM at two consecutive steps is highly similar. However, since different parts of the input shift position, the KV cache from the previous step cannot be effectively reused. The KV cache functions like a ROM. If we could transform it into a RAM-like structure that supports simple editing operations, such as swapping adjacent tokens while maintaining the correct tokens and positional embeddings, this could significantly improve computational efficiency.

Implementation of the Turing machine prototype. When generating the training set for the executor, CAEF wants to ensure the correctness of the samples and enable the executor to learn key computational steps, such as carrying over or exiting loops. A practical approach is to construct a Turing machine prototype corresponding to the target operator and record its execution process. While there are many existing Turing machine designs, the implementation process may take some human-involved effort. A future work could design a generation process to translate existing Turing machines into CAEF required Turing machine prototypes.

Inability to Solve Math Word Problems. Currently, our method requires manual selection of the active LoRA adapter, rather than enabling the model to autonomously determine the appropriate adapter. This limitation hinders the direct application of our approach to solving math word problems. However, our method can be regarded as a modular component. Several studies have explored integrating Mixture of Experts (MoE) and LoRA techniques to facilitate the automatic

486 selection and switching of active LoRA adapters based on the input (Wu et al., 2023a; Zadouri et al.,
487 2023; Huang et al., 2023; Xu et al., 2024). These studies are orthogonal to our approach, and we
488 posit that combining these techniques with our method could enable effective application to math
489 word problems. For example, leveraging the CAEF plug-in, a large language model (LLM) could
490 dynamically switch to CAEF to handle arithmetic computations as part of the reasoning process in
491 solving such problems.

492

493 7 CONCLUSION

494

495 This paper proposes a framework that enables LLMs to learn to execute step-by-step arithmetic
496 computational logic by imitating the behavior of a Turing machine. This approach significantly
497 enhances LLMs' computational capability without relying on any external tools. Moreover, the
498 framework is highly scalable, allowing the construction of complex executors by composing learned
499 basic executors, reducing the difficulty of [learning](#) the complex logic. We hope that our work
500 provides a new perspective for enabling LLMs to learn rule-based computation.

501

502

503

504

505

506

507

508

509

510

511

512

513

514

515

516

517

518

519

520

521

522

523

524

525

526

527

528

529

530

531

532

533

534

535

536

537

538

539

REFERENCES

- 540
541
542 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni
543 Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4
544 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- 545 Daniel Andor, Luheng He, Kenton Lee, and Emily Pitler. Giving bert a calculator: Finding
546 operations and arguments with reading comprehension. In *Proceedings of the 2019 Conference on*
547 *Empirical Methods in Natural Language Processing and the 9th International Joint Conference*
548 *on Natural Language Processing (EMNLP-IJCNLP)*, pp. 5947–5952, 2019.
- 549 Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam
550 Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm:
551 Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):
552 1–113, 2023.
- 553 Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha
554 Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models.
555 *arXiv preprint arXiv:2407.21783*, 2024.
- 556 Zhibin Gou, Zhihong Shao, Yeyun Gong, Yujiu Yang, Minlie Huang, Nan Duan, Weizhu Chen,
557 et al. Tora: A tool-integrated reasoning agent for mathematical problem solving. *arXiv preprint*
558 *arXiv:2309.17452*, 2023.
- 559 Shibo Hao, Tianyang Liu, Zhen Wang, and Zhiting Hu. Toolkengpt: Augmenting frozen language
560 models with massive tools via tool embeddings. *Advances in neural information processing*
561 *systems*, 36, 2024.
- 562 Joy He-Yueya, Gabriel Poesia, Rose E Wang, and Noah D Goodman. Solving math word problems
563 by combining language models with symbolic solvers. *arXiv preprint arXiv:2304.09102*, 2023.
- 564 Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang,
565 and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint*
566 *arXiv:2106.09685*, 2021.
- 567 Yi Hu, Xiaojuan Tang, Haotong Yang, and Muhan Zhang. Case-based or rule-based: How do
568 transformers do the math? *arXiv preprint arXiv:2402.17709*, 2024.
- 569 Chengsong Huang, Qian Liu, Bill Yuchen Lin, Tianyu Pang, Chao Du, and Min Lin.
570 LoraHub: Efficient cross-task generalization via dynamic lora composition. *arXiv preprint*
571 *arXiv:2307.13269*, 2023.
- 572 Jiaxin Huang, Shixiang Shane Gu, Le Hou, Yuexin Wu, Xuezhi Wang, Hongkun Yu, and Jiawei
573 Han. Large language models can self-improve. *arXiv preprint arXiv:2210.11610*, 2022.
- 574 Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris
575 Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand,
576 et al. Mixtral of experts. *arXiv preprint arXiv:2401.04088*, 2024.
- 577 Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large
578 language models are zero-shot reasoners. *Advances in neural information processing systems*,
579 35:22199–22213, 2022.
- 580 Nayoung Lee, Kartik Sreenivasan, Jason D Lee, Kangwook Lee, and Dimitris Papailiopoulos.
581 Teaching arithmetic to small transformers. *arXiv preprint arXiv:2307.03381*, 2023.
- 582 Pan Lu, Baolin Peng, Hao Cheng, Michel Galley, Kai-Wei Chang, Ying Nian Wu, Song-Chun
583 Zhu, and Jianfeng Gao. Chameleon: Plug-and-play compositional reasoning with large language
584 models. *Advances in Neural Information Processing Systems*, 36, 2024.
- 585 Maxwell Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin,
586 David Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, et al. Show
587 your work: Scratchpads for intermediate computation with language models. *arXiv preprint*
588 *arXiv:2112.00114*, 2021.
- 589
590
591
592
593

- 594 Jingqing Ruan, Yihong Chen, Bin Zhang, Zhiwei Xu, Tianpeng Bao, Hangyu Mao, Ziyue Li, Xingyu
595 Zeng, Rui Zhao, et al. Tptu: Task planning and tool usage of large language model-based ai
596 agents. In *NeurIPS 2023 Foundation Models for Decision Making Workshop*, 2023.
597
- 598 Timo Schick, Jane Dwivedi-Yu, Roberto Dessi, Roberta Raileanu, Maria Lomeli, Eric Hambro,
599 Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can
600 teach themselves to use tools. *Advances in Neural Information Processing Systems*, 36, 2024.
- 601 Michael Sipser. *Introduction to the Theory of Computation*. International Thomson Publishing, 1st
602 edition, 1996. ISBN 053494728X.
603
- 604 Alan Mathison Turing et al. On computable numbers, with an application to the
605 entscheidungsproblem. *J. of Math*, 58(345-363):5, 1936.
- 606 Ke Wang, Houxing Ren, Aojun Zhou, Zimu Lu, Sichun Luo, Weikang Shi, Renrui Zhang, Linqi
607 Song, Mingjie Zhan, and Hongsheng Li. Mathcoder: Seamless code integration in llms for
608 enhanced mathematical reasoning. *arXiv preprint arXiv:2310.03731*, 2023.
609
- 610 Xun Wu, Shaohan Huang, and Furu Wei. Mixture of lora experts. In *The Twelfth International
611 Conference on Learning Representations*, 2023a.
- 612 Zhaofeng Wu, Linlu Qiu, Alexis Ross, Ekin Akyürek, Boyuan Chen, Bailin Wang, Najoung Kim,
613 Jacob Andreas, and Yoon Kim. Reasoning or reciting? exploring the capabilities and limitations
614 of language models through counterfactual tasks. *arXiv preprint arXiv:2307.02477*, 2023b.
615
- 616 Jingwei Xu, Junyu Lai, and Yunpeng Huang. Meteora: Multiple-tasks embedded lora for large
617 language models. *arXiv preprint arXiv:2405.13053*, 2024.
- 618 Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T Kwok,
619 Zhenguo Li, Adrian Weller, and Weiyang Liu. Metamath: Bootstrap your own mathematical
620 questions for large language models. *arXiv preprint arXiv:2309.12284*, 2023.
621
- 622 Ted Zadouri, Ahmet Üstün, Arash Ahmadian, Beyza Ermis, Acyr Locatelli, and Sara Hooker.
623 Pushing mixture of experts to the limit: Extremely parameter efficient moe for instruction tuning.
624 In *The Twelfth International Conference on Learning Representations*, 2023.
- 625 Aojun Zhou, Ke Wang, Zimu Lu, Weikang Shi, Sichun Luo, Zipeng Qin, Shaoqing Lu, Anya Jia,
626 Linqi Song, Mingjie Zhan, et al. Solving challenging math word problems using gpt-4 code
627 interpreter with code-based self-verification. *arXiv preprint arXiv:2308.07921*, 2023.
- 628 Yongchao Zhou, Uri Alon, Xinyun Chen, Xuezhi Wang, Rishabh Agarwal, and Denny
629 Zhou. Transformers can achieve length generalization but not robustly. *arXiv preprint
630 arXiv:2402.09371*, 2024.
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647

648 A APPENDIX

649 A.1 EXAMPLE OF SAMPLES IN TRAINING SET OF EXECUTOR AND ALIGNER

650 A.1.1 ADDITION

651 Addition executor:

652 *Input:*

653 The following is a state paired with a command to be executed of a Turing Machine that
654 performs addition.

655 The state includes the current operator, the current state of the machine, the current tape
656 contents, and the current head positions.

657 - There are three states in the machine: q0, q1, and qH. The machine starts in state q0 and
658 halts when it reaches state qH. q1 is the state where the machine does the addition and
659 calculates the carry out.

660 - The head positions are represented by [HEAD1] and [HEAD2], which indicate the
661 positions of the heads on the two operands.

662 - The carry out is represented by [C].

663 - The output position is represented by [OUTPUT].

664 The command includes a series of actions to be executed by the machine and they are
665 separated by commas.

666 - [OUTPUT] <number>: Write the number to the output position.

667 - [OUTPUT] <direction>: Move the output head to the direction.

668 - [C] <number>: Write the number to the carry out register.

669 - [HEAD1] <direction>: Move the head on the first operand to the direction.

670 - [HEAD2] <direction>: Move the head on the second operand to the direction.

671 - <state>: Move the machine to the state.

672 The machine performs addition by reading the digits from the two operands and writing the
673 sum to the output tape.

674 Based the current state and the command, predict the next state of the machine and next
675 command to be executed.

676 ADD, q0, [HEAD1] |5|4[HEAD2] |7|6 [C] [OUTPUT]

677 CMD: [C] 0, [HEAD1] RIGHT, [HEAD2] RIGHT, q1

678 *Output:*

679 ADD, q1, [HEAD1]|5|4 [HEAD2]|7|6 [C]0 [OUTPUT]

680 CMD: [C] 1, [OUTPUT] 2, [OUTPUT] RIGHT, [HEAD1] RIGHT, [HEAD2] RIGHT, q1

681 Addition aligner:

682 *Input:*

683 The following is an input to a Turing Machine or an output of a Turing Machine.

684 The task is doing an adaptation:

685 - If it is an input, adapt the original input to the format that the Turing Machine can
686 understand.

687 - If it is an output, adapt the original output to the format that represents the final result.

688 Input example:

689 ““

690 - input:

691 1504+2379=

702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755

- output:

ADD, q0, [HEAD1] |4|0|5|1|[HEAD2] |9|7|3|2 [C] [OUTPUT]
CMD: [C] 0, [HEAD1] RIGHT, [HEAD2] RIGHT, q1

““

Output example:

““

- input:

ADD, qH, |7|6|3|4|[HEAD1] |4|3|2|1|[HEAD2] [C]0 |1|0|6|5
No command to execute. Halt state.

- output:

4367+1234=5601

““

There are two lines that represent the Turing Machine:

- The first line is the current state of the machine.

- The second line is the command to be executed.

And this format is fit to both input and output as the examples shown above.

For the current state (the first line):

- There are at least 2 states in the machine: q0 and qH. The machine starts in state q0 and halts when it reaches state qH.

- The head positions are represented by [HEAD1] and [HEAD2], which followed by two operands.

- [C] represents the carry out register and [OUTPUT] represents the output position. And these two are empty at the beginning.

The command (the second line) includes a series of actions to be executed by the machine and they are separated by commas.

- [HEAD] <direction>: Move the head to the direction.

- [C] <number>: Write the number to the carry out register.

- <state>: Move the machine to the state.

Note that the number is represented in reverse order in machine, which is beneficial to the machine to perform the subtraction operation.

Based on the input, determine it is an input or an output, and adapt it to the format correspondingly.

45+67=

Output:

ADD, q0, [HEAD1] |5|4|[HEAD2] |7|6 [C] [OUTPUT]
CMD: [C] 0, [HEAD1] RIGHT, [HEAD2] RIGHT, q1

A.1.2 SUBTRACTION

Subtraction executor:

Input:

The following is a input to be executed of a Turing Machine that performs subtraction.

To solve a subtraction problem by the machine, the machine is required to provide the initial state and command for other basic machines, including addition, reflection and left mask.

For example, for $47819 - 12345 = 35474$, the machine will perform the following steps:

- step 1: call reflection, $99999 - 12345 = 87654$

- step 2: call addition, $47819 + 87654 = 135473$

- step 3: call addition, $135473 + 1 = 135474$

756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809

- step 4: call left mask, left_mask(135474) = 35474

The input may includes four lines or the original subtraction problem.

When it is original problem, generate the initial subtraction state, command and prepare the initial state and the first command of the first called machine.

When it includes four lines, it means the previous state, command and the result of the called machine. In detail:

- The first line is the current state of the machine.
- The second line is the command to be executed.
- The third line and the fourth line are halt state of another machine which is called by the subtraction machine at previous step.

For the current state (the first line):

- There are five states in the machine: q0, q1, q2, q3 and qH. The machine starts in state q0 and halts when it reaches state qH.
- The head positions are represented by [HEAD1] and [HEAD2], which followed by two operands.

The command (the second line) includes a series of actions to be executed by the machine and they are separated by commas.

- [CALL] <operation>: Call another machine to perform the operation.
- <state>: Move the machine to the state.

When the commands include [CALL], another extra two lines are needed to specify the initial state and the first command of the machine to be called.

As for initial state, it should include the operation, q0 state, operands and the head positions.

As for the first command:

- [OUTPUT] <number>: Write the number to the output position.
- [OUTPUT] <direction>: Move the output head to the direction.
- [HEAD1] <direction>: Move the head on the first operand to the direction.
- [HEAD2] <direction>: Move the head on the second operand to the direction.
- <state>: Move the machine to the state.

The machine performs subtraction by reading the digits from the two operands and calling other machines to complete the subtraction operation.

Based on the current input, predict the output which includes next state, next command and the initial state and the first command of the machine to be called.

SUB, q0, [HEAD1]|7|4 [HEAD2]|2|1
CMD q1

Output:

SUB, q1, [HEAD1]|7|4 [HEAD2]|2|1
CMD [CALL] REFLECTION, q2
REFLECTION, q0, [HEAD1] |9|9[HEAD2] |2|1 [OUTPUT]
CMD [HEAD1] RIGHT, [HEAD2] RIGHT, q1

Subtraction aligner:

Input:

The following is an input to a Turing Machine or an output of a Turing Machine.

The task is doing an adaptation:

- If it is an input, adapt the original input to the format that the Turing Machine can understand.

810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863

- If it is an output, adapt the original output to the format that represents the final result.

Input example:

““

- input:

4531-1504=

- output:

SUB, q0, [HEAD1]|1|3|5|4 [HEAD2]|4|0|5|1

CMD q1

““

Output example:

““

- input:

SUB, qH, [HEAD1]|1|3|5|4 [HEAD2]|4|0|5|1 |7|2|0|3

No command to execute. Halt state.

- output:

4531-1504=3027

““

There are two lines that represent the Turing Machine:

- The first line is the current state of the machine.

- The second line is the command to be executed.

And this format is fit to both input and output as the examples shown above.

For the current state (the first line):

- There are at least 2 states in the machine: q0 and qH. The machine starts in state q0 and halts when it reaches state qH.

- The head positions are represented by [HEAD1] and [HEAD2], which followed by two operands.

The command (the second line) includes a series of actions to be executed by the machine and they are separated by commas.

- [HEAD] <direction>: Move the head to the direction.

- [OUTPUT] <number>: Write the number to the output position.

- <state>: Move the machine to the state.

Note that the number is represented in reverse order in machine, which is beneficial to the machine to perform the subtraction operation.

Based on the input, determine it is an input or an output, and adapt it to the format correspondingly.

46-28=

Output:

SUB, q0, [HEAD1]|6|4 [HEAD2]|8|2

CMD q1

A.1.3 MULTIPLICATION

Multiplication executor:

Input:

The following is a input to be executed of a Turing Machine that performs multiplication.

864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917

To solve a multiplication problem by the machine, the machine is required to provide the initial state and command for other basic machines, including addition and less_than machines.

For example, for $4513 * 3 = 13539$, the machine will perform the following algorithm:

- step 1: cnt = 1, sum = 4513(operand1)
- step 2: call less_than, determine whether cnt < 3(operand2), if yes, go to step 3, otherwise, go to step 5
- step 3: call addition, sum = sum + 4513(operand1)
- step 4: call addition, cnt = cnt + 1, go to step 2
- step 5: current machine halts

The input includes at least two lines and may have two more lines.

- The first line is the current state of the machine.
- The second line is the command to be executed.

When there are two more lines:

- The third line and the fourth line are halt state of another machine which is called by the multiplication machine at previous step.

For the current state (the first line):

- There are five states in the machine: q0, q1, q2, q3 and qH. The machine starts in state q0 and halts when it reaches state qH. q1, q2 and q3 are used to perform the loop structure.
- The head positions are represented by [HEAD1] and [HEAD2], which followed by two operands.

The command (the second line) includes a series of actions to be executed by the machine and they are separated by commas.

- [OUTPUT] <number>: Write the number to the output position.
- [COUNT] <number>: Write the number to the count register.
- [CALL] <operation>: Call another machine to perform the operation.
- <state>: Move the machine to the state.

When the commands include [CALL], another extra two lines are needed to specify the initial state and the first command of the machine to be called.

As for initial state, it should include the operation, q0 state, operands and the head positions.

As for the first command:

- [OUTPUT] <number>: Write the number to the output position.
- [OUTPUT] <direction>: Move the output head to the direction.
- [HEAD1] <direction>: Move the head on the first operand to the direction.
- [HEAD2] <direction>: Move the head on the second operand to the direction.
- <state>: Move the machine to the state.

The machine performs multiplication by reading the digits from the two operands and calling other machines to complete the multiplication operation.

Based on the current input, predict the output which includes next state, next command and the initial state and the first command of the machine to be called.

```
MUL, q3, [HEAD1]|3|8|6 [HEAD2]|8|6 [COUNT]|5|4 [OUTPUT]|8|1|4|1|3
CMD [CALL] ADD, q1
ADD, qH, |5|4[HEAD1] |1[HEAD2] [C]0 |6|4
No command to execute. Halt state.
```

Output:

```
MUL, q1, [HEAD1]|3|8|6 [HEAD2]|8|6 [COUNT]|6|4 [OUTPUT]|8|1|4|1|3
CMD [CALL] LESS_THAN, q2
```

LESS_THAN, q0, [HEAD1]|6|4[HEAD2] |8|6 [OUTPUT]
 CMD [HEAD1] RIGHT, [HEAD2] RIGHT, [OUTPUT] False, q1

Multiplication aligner:

Input:

The following is an input to a Turing Machine or an output of a Turing Machine.

The task is doing an adaptation:

- If it is an input, adapt the original input to the format that the Turing Machine can understand.
- If it is an output, adapt the original output to the format that represents the final result.

Input example:

““

- input:

44814*5=

- output:

MUL, q0, [HEAD1]|4|1|8|4|4 [HEAD2]|5 [COUNT] [OUTPUT]

CMD [COUNT] 1, [OUTPUT]|4|1|8|4|4, q1

““

Output example:

““

- input:

MUL, qH, [HEAD1]|4|1|8|4|4 [HEAD2]|5 [COUNT]|5 |0|7|0|4|2|2

No command to execute. Halt state.

- output:

44814*5=224070

““

There are two lines that represent the Turing Machine:

- The first line is the current state of the machine.
- The second line is the command to be executed.

And this format is fit to both input and output as the examples shown above.

For the current state (the first line):

- There are at least 2 states in the machine: q0 and qH. The machine starts in state q0 and halts when it reaches state qH.
- The head positions are represented by [HEAD1] and [HEAD2], which followed by two operands.

The command (the second line) includes a series of actions to be executed by the machine and they are separated by commas.

- [HEAD] <direction>: Move the head to the direction.
- [OUTPUT] <number>: Write the number to the output position.
- [COUNT] <number>: Write the number to the count register.
- <state>: Move the machine to the state.

Based on the input, determine it is an input or an output, and adapt it to the format correspondingly.

652202674*9560505=

Output:

MUL, q0, [HEAD1]|4|7|6|2|0|2|2|5|6 [HEAD2]|5|0|5|0|6|5|9 [COUNT] [OUTPUT]

972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

CMD [COUNT] 0, [OUTPUT] 0, q1

A.1.4 DIVISION

Division executor:

Input:

The following is a input to be executed of a Turing Machine that performs division.

To solve a division problem by the machine, the machine is required to provide the initial state and command for other basic machines, including addition and greater_than machines.

For example, for $4513 // 1504 = 3$, the machine will perform the following algorithm:

- step 1: output = 0, cnt = 1504(operand2)
- step 2: call greater_than, determine whether cnt > 4513(operand1), if yes, go to step 5, otherwise, go to step 3
- step 3: call addition, output = output + 1
- step 4: call addition, cnt = cnt + 1504, go to step 2
- step 5: current machine halts, output is the result

The input includes at least two lines and may have two more lines.

- The first line is the current state of the machine.
- The second line is the command to be executed.

When there are two more lines:

- The third line and the fourth line are halt state of another machine which is called by the division machine at previous step.

For the current state (the first line):

- There are five states in the machine: q0, q1, q2, q3 and qH. The machine starts in state q0 and halts when it reaches state qH. q1, q2 and q3 are used to perform the loop structure.
- The head positions are represented by [HEAD1] and [HEAD2], which followed by two operands.

The command (the second line) includes a series of actions to be executed by the machine and they are separated by commas.

- [OUTPUT] <number>: Write the number to the output position.
- [COUNT] <number>: Write the number to the count register.
- [CALL] <operation>: Call another machine to perform the operation.
- <state>: Move the machine to the state.

When the commands include [CALL], another extra two lines are needed to specify the initial state and the first command of the machine to be called.

As for initial state, it should include the operation, q0 state, operands and the head positions.

As for the first command:

- [OUTPUT] <number>: Write the number to the output position.
- [OUTPUT] <direction>: Move the output head to the direction.
- [HEAD1] <direction>: Move the head on the first operand to the direction.
- [HEAD2] <direction>: Move the head on the second operand to the direction.
- <state>: Move the machine to the state.

The machine performs division by reading the digits from the two operands and calling other machines to complete the division operation.

Based on the current input, predict the output which includes next state, next command and

1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079

the initial state and the first command of the machine to be called.

DIV, q1, [HEAD1]|0|5|6 [HEAD2]|8|3|2 [COUNT]|6|7|4 [OUTPUT]|1
 CMD [CALL] GREATER_THAN, q2
 GREATER_THAN, qH, |6|7|4[HEAD1] |0|5|6[HEAD2] False
 No command to execute. Halt state.

Output:

DIV, q2, [HEAD1]|0|5|6 [HEAD2]|8|3|2 [COUNT]|6|7|4 [OUTPUT]|1
 CMD [CALL] ADD, q3
 ADD, q0, [HEAD1] |1[HEAD2] |1 [C] [OUTPUT]
 CMD: [C] 0, [HEAD1] RIGHT, [HEAD2] RIGHT, q1

Division aligner:

Input:

The following is an input to a Turing Machine or an output of a Turing Machine.

The task is doing an adaptation:

- If it is an input, adapt the original input to the format that the Turing Machine can understand.
- If it is an output, adapt the original output to the format that represents the final result.

Input example:

““

- input:

4531//1504=

- output:

DIV, q0, [HEAD1]|3|1|5|4 [HEAD2]|4|0|5|1 [COUNT] [OUTPUT]
 CMD [COUNT]|4|0|5|1, [OUTPUT] 0, q1

““

Output example:

““

- input:

DIV, qH, [HEAD1]|3|1|5|4 [HEAD2]|4|0|5|1 [COUNT]|6|1|0|6 |3
 No command to execute. Halt state.

- output:

4531//1504=3

““

There are two lines that represent the Turing Machine:

- The first line is the current state of the machine.
- The second line is the command to be executed.

And this format is fit to both input and output as the examples shown above.

For the current state (the first line):

- There are at least 2 states in the machine: q0 and qH. The machine starts in state q0 and halts when it reaches state qH.
- The head positions are represented by [HEAD1] and [HEAD2], which followed by two operands.

The command (the second line) includes a series of actions to be executed by the machine and they are separated by commas.

- [HEAD] <direction>: Move the head to the direction.
- [OUTPUT] <number>: Write the number to the output position.

1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133

- [COUNT] <number>: Write the number to the count register.
- <state>: Move the machine to the state.

Based on the input, determine it is an input or an output, and adapt it to the format correspondingly.

8634010848//613431149=

Output:

DIV, q0, [HEAD1]|8|4|8|0|1|0|4|3|6|8 [HEAD2]|9|4|1|1|3|4|3|1|6 [COUNT] [OUTPUT]
CMD [COUNT]|9|4|1|1|3|4|3|1|6, [OUTPUT] 0, q1

A.1.5 GREATER_THAN

Greater_than executor:

Input:

The following is a state paired with a command to be executed of a Turing Machine that determines whether the first operand is greater than the second operand.

The state includes the current operator, the current state of the machine, the current tape contents, and the current head positions.

- There are three states in the machine: q0, q1, and qH. The machine starts in state q0 and halts when it reaches state qH. q1 is the state where the machine does the comparison.
- The head positions are represented by [HEAD1] and [HEAD2], which indicate the positions of the heads on the two operands.
- The output position is represented by [OUTPUT].

The command includes a series of actions to be executed by the machine and they are separated by commas.

- [OUTPUT] <number>: Write the number to the output position.
- [OUTPUT] <direction>: Move the output head to the direction.
- [HEAD1] <direction>: Move the head on the first operand to the direction.
- [HEAD2] <direction>: Move the head on the second operand to the direction.
- <state>: Move the machine to the state.

The machine performs comparison by reading the digits from the two operands and writing the result to the output tape.

Based on the current state and the command, predict the next state of the machine and next command to be executed.

GREATER_THAN, q1, |1|7|6|7|0|[HEAD1]|5|1|3|1 |5|6|4|1|7|[HEAD2]|8|1|4|7|4|8|8|3|2|7
[OUTPUT]False
CMD [HEAD1] RIGHT, [HEAD2] RIGHT, [OUTPUT] False, q1

Output:

GREATER_THAN, q1, |1|7|6|7|0|5|[HEAD1]|1|3|1 |5|6|4|1|7|8|[HEAD2]|1|4|7|4|8|8|3|2|7
[OUTPUT]False
CMD [HEAD1] RIGHT, [HEAD2] RIGHT, q1

Greater_than aligner:

1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187

Input:

The following is an input to a Turing Machine or an output of a Turing Machine.

The task is doing an adaptation:

- If it is an input, adapt the original input to the format that the Turing Machine can understand.
- If it is an output, adapt the original output to the format that represents the final result.

Input example:

““

- input:

45131>15040=

- output:

GREATER_THAN, q0, [HEAD1] |1|3|1|5|4[HEAD2] |0|4|0|5|1 [OUTPUT]

CMD [HEAD1] RIGHT, [HEAD2] RIGHT, [OUTPUT] False, q1

““

Output example:

““

- input:

GREATER_THAN, qH, |1|3|1|5|4[HEAD1] |0|4|0|5|1[HEAD2] True

No command to execute. Halt state.

- output:

45131>15040=True

““

There are two lines that represent the Turing Machine:

- The first line is the current state of the machine.
- The second line is the command to be executed.

And this format is fit to both input and output as the examples shown above.

For the current state (the first line):

- There are at least 2 states in the machine: q0 and qH. The machine starts in state q0 and halts when it reaches state qH.
- The head positions are represented by [HEAD1] and [HEAD2], which followed by two operands.

The command (the second line) includes a series of actions to be executed by the machine and they are separated by commas.

- [HEAD] <direction>: Move the head to the direction.
- [OUTPUT] <direction>: Move the output head to the direction.
- [OUTPUT] <result>: Write the result to the output position.
- <state>: Move the machine to the state.

Based on the input, determine it is an input or an output, and adapt it to the format correspondingly.

46989>82541=

Output:

GREATER_THAN, q0, [HEAD1] |9|8|9|6|4[HEAD2] |1|4|5|2|8 [OUTPUT]

CMD [HEAD1] RIGHT, [HEAD2] RIGHT, [OUTPUT] False, q1

A.1.6 LESS_THAN

Less_than executor:

1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241

Input:

The following is a state paired with a command to be executed of a Turing Machine that determines whether the first operand is less than the second operand.

The state includes the current operator, the current state of the machine, the current tape contents, and the current head positions.

- There are three states in the machine: q0, q1, and qH. The machine starts in state q0 and halts when it reaches state qH. q1 is the state where the machine does the comparison.
- The head positions are represented by [HEAD1] and [HEAD2], which indicate the positions of the heads on the two operands.
- The output position is represented by [OUTPUT].

The command includes a series of actions to be executed by the machine and they are separated by commas.

- [OUTPUT] <number>: Write the number to the output position.
- [OUTPUT] <direction>: Move the output head to the direction.
- [HEAD1] <direction>: Move the head on the first operand to the direction.
- [HEAD2] <direction>: Move the head on the second operand to the direction.
- <state>: Move the machine to the state.

The machine performs comparison by reading the digits from the two operands and writing the result to the output tape.

Based on the current state and the command, predict the next state of the machine and next command to be executed.

LESS_THAN, q1, |4|1|0|[HEAD1]|2|0|6|1|[HEAD2]|2|7|6|[OUTPUT]True
CMD [HEAD1] RIGHT, [HEAD2] RIGHT, q1

Output:

LESS_THAN, q1, |4|1|0|2|[HEAD1]|0|6|1|2|[HEAD2]|7|6|[OUTPUT]True
CMD [OUTPUT] True, [OUTPUT], qH

Less_than aligner:

Input:

The following is an input to a Turing Machine or an output of a Turing Machine.

The task is doing an adaptation:

- If it is an input, adapt the original input to the format that the Turing Machine can understand.
- If it is an output, adapt the original output to the format that represents the final result.

Input example:

““

- input:

47182<83911=

- output:

LESS_THAN, q0, [HEAD1] |2|8|1|7|4|[HEAD2] |1|1|9|3|8|[OUTPUT]

CMD [HEAD1] RIGHT, [HEAD2] RIGHT, [OUTPUT] False, q1

““

Output example:

““

- input:

LESS_THAN, qH, |2|8|1|7|4|[HEAD1] |1|1|9|3|8|[HEAD2] True

1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295

No command to execute. Halt state.

- output:
47182<83911=True
““

There are two lines that represent the Turing Machine:

- The first line is the current state of the machine.
- The second line is the command to be executed.

And this format is fit to both input and output as the examples shown above.

For the current state (the first line):

- There are at least 2 states in the machine: q0 and qH. The machine starts in state q0 and halts when it reaches state qH.
- The head positions are represented by [HEAD1] and [HEAD2], which followed by two operands.

The command (the second line) includes a series of actions to be executed by the machine and they are separated by commas.

- [HEAD] <direction>: Move the head to the direction.
- [OUTPUT] <direction>: Move the output head to the direction.
- [OUTPUT] <result>: Write the result to the output position.
- <state>: Move the machine to the state.

Based on the input, determine it is an input or an output, and adapt it to the format correspondingly.

LESS_THAN, qH, |1|5|9|4|4|6|[HEAD1]|6|2|1|3|5|8|0|9|8 |3|7|2|6|4|2|[HEAD2] False

No command to execute. Halt state.

Output:

890853126644951<246273=False

A.1.7 EQUAL

Equal executor:

Input:

The following is a state paired with a command to be executed of a Turing Machine that performs equality comparison.

The state includes the current operator, the current state of the machine, the current tape contents, and the current head positions.

- There are three states in the machine: q0, q1, and qH. The machine starts in state q0 and halts when it reaches state qH. q1 is the state where the machine does the equality comparison.
- The head positions are represented by [HEAD1] and [HEAD2], which indicate the positions of the heads on the two operands.
- The output position is represented by [OUTPUT].

The command includes a series of actions to be executed by the machine and they are separated by commas.

- [OUTPUT] <number>: Write the number to the output position.
- [OUTPUT] <direction>: Move the output head to the direction.
- [HEAD1] <direction>: Move the head on the first operand to the direction.
- [HEAD2] <direction>: Move the head on the second operand to the direction.
- <state>: Move the machine to the state.

1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349

The machine performs equality comparison by reading the digits from the two operands and writing the result to the output tape.

Based on the current state and the command, predict the next state of the machine and next command to be executed.

EQUAL, q1, |0|5[HEAD1]|9 |0|5[HEAD2]|9 [OUTPUT]True
CMD [HEAD1] RIGHT, [HEAD2] RIGHT, q1

Output:

EQUAL, q1, |0|5|9[HEAD1] |0|5|9[HEAD2] [OUTPUT]True
CMD [OUTPUT], qH

Equal aligner:

Input:

The following is an input to a Turing Machine or an output of a Turing Machine.

The task is doing an adaptation:

- If it is an input, adapt the original input to the format that the Turing Machine can understand.
- If it is an output, adapt the original output to the format that represents the final result.

Input example:

““

- input:

45263==45263=

- output:

EQUAL, q0, [HEAD1] |3|6|2|5|4[HEAD2] |3|6|2|5|4 [OUTPUT]

CMD [HEAD1] RIGHT, [HEAD2] RIGHT, [OUTPUT] True, q1

““

Output example:

““

- input:

EQUAL, qH, |3|6|2|5|4[HEAD1] |3|6|2|5|4[HEAD2] True

No command to execute. Halt state.

- output:

45263==45263=True

““

There are two lines that represent the Turing Machine:

- The first line is the current state of the machine.
- The second line is the command to be executed.

And this format is fit to both input and output as the examples shown above.

For the current state (the first line):

- There are at least 2 states in the machine: q0 and qH. The machine starts in state q0 and halts when it reaches state qH.
- The head positions are represented by [HEAD1] and [HEAD2], which followed by two operands.

The command (the second line) includes a series of actions to be executed by the machine and they are separated by commas.

- [HEAD] <direction>: Move the head to the direction.
- [OUTPUT] <direction>: Move the output head to the direction.

1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403

- [OUTPUT] <result>: Write the result to the output position.
- <state>: Move the machine to the state.

Note that the number is represented in reverse order in machine, which is beneficial to the machine to perform the subtraction operation.

Based on the input, determine it is an input or an output, and adapt it to the format correspondingly.

EQUAL, qH, |6|5|6|8|8|9|7|1|6|7|7|1|2[HEAD1] |6|5|6|8|8|9|7|1|6|7|7|1|2[HEAD2] True
No command to execute. Halt state.

Output:
2177617988656==2177617988656=True

A.2 ARITHMETIC EXPRESSION TEMPLATE

Templates in Table 3 are used to generate arithmetic expressions in our experiment.

Table 3: Templates used for generating arithmetic expressions in training set and test set.

Operator	Template
Addition	{op1}+{op2}=
Subtraction	{op1}-{op2}=
Multiplication	{op1}*{op2}=
Division	{op1}://{op2}=
Greater	{op1}>{op2}=
Less	{op1}<{op2}=
Equal	{op1}=={op2}=

A.3 FULL COMPUTATION PROCESS OF THE EXAMPLES

The followings are the full computation process of the examples in 3.2 and 3.3.

A.3.1 ADDITION

Step 1 (aligner):
45+67=

Step 2 (executor):
*state*₀: ADD, q0, [HEAD1] |5|4[HEAD2] |7|6 [C] [OUTPUT]
*command*₀: CMD: [C] 0, [HEAD1] RIGHT, [HEAD2] RIGHT, q1

Step 3 (executor):
*state*₁: ADD, q1, [HEAD1]|5|4 [HEAD2]|7|6 [C]0 [OUTPUT]
*command*₁: CMD: [C] 1, [OUTPUT] 2, [OUTPUT] RIGHT, [HEAD1] RIGHT, [HEAD2] RIGHT, q1

Step 4 (executor):
*state*₂: ADD, q1, |5[HEAD1]|4 |7[HEAD2]|6 [C]1 |2[OUTPUT]
*command*₂: CMD: [C] 1, [OUTPUT] 1, [OUTPUT] RIGHT, [HEAD1] RIGHT, [HEAD2] RIGHT, q1

1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457

Step 5 (executor):
*state*₃: ADD, q1, |5|4[HEAD1] |7|6[HEAD2] [C]1 |2|1[OUTPUT]
*command*₃: CMD: [OUTPUT] 1, [OUTPUT], [C], qH

Step 6 (executor):
*state*₄: ADD, qH, |5|4[HEAD1] |7|6[HEAD2] [C]1 |2|1|1
*command*₄: No command to execute. Halt state.

Step 7 (aligner):
 45+67=112

A.3.2 MULTIPLICATION

Step 1 (aligner):
 89*2=

Step 2 (executor):
*state*₀: MUL, q0, [HEAD1]|9|8 [HEAD2]|2 [COUNT] [OUTPUT]
*command*₀: CMD [COUNT] 0, [OUTPUT] 0, q1

Step 3-1, before call (executor):
*state*₁: MUL, q1, [HEAD1]|9|8 [HEAD2]|2 [COUNT]|0 [OUTPUT]|0
*command*₁: CMD [CALL] LESS_THAN, q2
*callee_state*₀: LESS_THAN, q0, [HEAD1] |0|[HEAD2] |2 [OUTPUT]
*callee_command*₀: CMD [HEAD1] RIGHT, [HEAD2] RIGHT, [OUTPUT] False, q1

Step 3-1, after call (executor):
*state*₁: MUL, q1, [HEAD1]|9|8 [HEAD2]|2 [COUNT]|0 [OUTPUT]|0
*command*₁: CMD [CALL] LESS_THAN, q2
*callee_state*_H: LESS_THAN, qH, |0|[HEAD1] |2|[HEAD2] True
*callee_command*_H: No command to execute. Halt state.

Step 4-1, before call (executor):
*state*₂: MUL, q2, [HEAD1]|9|8 [HEAD2]|2 [COUNT]|0 [OUTPUT]|0
*command*₂: CMD [CALL] ADD, q3
*callee_state*₀: ADD, q0, [HEAD1] |9|8[HEAD2] |0 [C] [OUTPUT]
*callee_command*₀: CMD: [C] 0, [HEAD1] RIGHT, [HEAD2] RIGHT, q1

Step 4-1, after call (executor):
*state*₂: MUL, q2, [HEAD1]|9|8 [HEAD2]|2 [COUNT]|0 [OUTPUT]|0
*command*₂: CMD [CALL] ADD, q3
*callee_state*_H: ADD, qH, |9|8[HEAD1] |0|[HEAD2] [C]0 |9|8
*callee_command*_H: No command to execute. Halt state.

Step 5-1, before call (executor):
*state*₃: MUL, q3, [HEAD1]|9|8 [HEAD2]|2 [COUNT]|0 [OUTPUT]|9|8
*command*₃: CMD [CALL] ADD, q1
*callee_state*₀: ADD, q0, [HEAD1] |0|[HEAD2] |1 [C] [OUTPUT]
*callee_command*₀: CMD: [C] 0, [HEAD1] RIGHT, [HEAD2] RIGHT, q1

Step 5-1, after call (executor):
*state*₃: MUL, q3, [HEAD1]|9|8 [HEAD2]|2 [COUNT]|0 [OUTPUT]|9|8
*command*₃: CMD [CALL] ADD, q1
*callee_state*_H: ADD, qH, |0|[HEAD1] |1|[HEAD2] [C]0 |1

1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511

callee_command_H: No command to execute. Halt state.

Step 6-1, before call (executor):

state₄: MUL, q1, [HEAD1]|9|8 [HEAD2]|2 [COUNT]|1 [OUTPUT]|9|8

command₄: CMD [CALL] LESS_THAN, q2

callee_state₀: LESS_THAN, q0, [HEAD1] |1[HEAD2] |2 [OUTPUT]

callee_command₀: CMD [HEAD1] RIGHT, [HEAD2] RIGHT, [OUTPUT] False, q1

Step 6-2, after call (executor):

state₄: MUL, q1, [HEAD1]|9|8 [HEAD2]|2 [COUNT]|1 [OUTPUT]|9|8

command₄: CMD [CALL] LESS_THAN, q2

callee_state_H: LESS_THAN, qH, |1[HEAD1] |2[HEAD2] True

callee_command_H: No command to execute. Halt state.

Step 7-1, before call (executor):

state₅: MUL, q2, [HEAD1]|9|8 [HEAD2]|2 [COUNT]|1 [OUTPUT]|9|8

command₅: CMD [CALL] ADD, q3

callee_state₀: ADD, q0, [HEAD1] |9|8[HEAD2] |9|8 [C] [OUTPUT]

callee_command₀: CMD: [C] 0, [HEAD1] RIGHT, [HEAD2] RIGHT, q1

Step 7-2, after call (executor):

state₅: MUL, q2, [HEAD1]|9|8 [HEAD2]|2 [COUNT]|1 [OUTPUT]|9|8

command₅: CMD [CALL] ADD, q3

callee_state_H: ADD, qH, |9|8[HEAD1] |9|8[HEAD2] [C]1 |8|7|1

callee_command_H: No command to execute. Halt state.

Step 8-1, before call (executor):

state₆: MUL, q3, [HEAD1]|9|8 [HEAD2]|2 [COUNT]|1 [OUTPUT]|8|7|1

command₆: CMD [CALL] ADD, q1

callee_state₀: ADD, q0, [HEAD1] |1[HEAD2] |1 [C] [OUTPUT]

callee_command₀: CMD: [C] 0, [HEAD1] RIGHT, [HEAD2] RIGHT, q1

Step 8-2, after call (executor):

state₆: MUL, q3, [HEAD1]|9|8 [HEAD2]|2 [COUNT]|1 [OUTPUT]|8|7|1

command₆: CMD [CALL] ADD, q1

callee_state_H: ADD, qH, |1[HEAD1] |1[HEAD2] [C]0 |2

callee_command_H: No command to execute. Halt state.

Step 9-1, before call (executor):

state₇: MUL, q1, [HEAD1]|9|8 [HEAD2]|2 [COUNT]|2 [OUTPUT]|8|7|1

command₇: CMD [CALL] LESS_THAN, q2

callee_state₀: LESS_THAN, q0, [HEAD1] |2[HEAD2] |2 [OUTPUT]

callee_command₀: CMD [HEAD1] RIGHT, [HEAD2] RIGHT, [OUTPUT] False, q1

Step 9-2, after call (executor):

state₇: MUL, q1, [HEAD1]|9|8 [HEAD2]|2 [COUNT]|2 [OUTPUT]|8|7|1

command₇: CMD [CALL] LESS_THAN, q2

callee_state_H: LESS_THAN, qH, |2[HEAD1] |2[HEAD2] False

callee_command_H: No command to execute. Halt state.

Step 10 (executor):

state₈: MUL, qH, [HEAD1]|9|8 [HEAD2]|2 [COUNT]|2 |8|7|1

command₈: No command to execute. Halt state.

Step 11 (aligner):

89*2=178

1512 A.4 IMPLEMENTATION OF SUBTRACTION OPERATOR
1513

1514 We implement subtraction in the CAEF framework by drawing inspiration from how subtraction is
1515 handled in CPUs. For subtraction in the form $a - b = c$, the process can be broken down into four
1516 steps:

- 1517
- 1518 1. Compute $\text{Reflection}(a, b)$: Generate a number a_9 , where all digits are 9 and it is the same
1519 length as a . Perform a reflection operation, which is essentially subtraction, between a_9
1520 and b . Since all digits of a_9 are 9, no borrowing occurs during this subtraction. Let the
1521 result of this step be p .
 - 1522 2. Compute $a + p$, and let the result be q .
 - 1523 3. Compute $q + 1$, and let the result be r .
 - 1524 4. Compute $\text{Left_mask}(r)$: Remove the leading 1 from the most significant digit of r . After
1525 this step, the final result, c , is obtained.
1526

1527 For example, in the case of $4531 - 1504 = 3027$, the process is as follows:
1528

1529

1530 *Step 1 (Reflection):*
1531 $\text{Reflection}(4531, 1504) = 9999 - 1504 = 8495$

1532 *Step 2 (Addition):*
1533 $4531 + 8495 = 13026$

1534

1535 *Step 3 (Addition):*
1536 $13026 + 1 = 13027$

1537

1538 *Step 4 (Left_mask):*
1539 $\text{Left_mask}(13027) = 3027$
1540
1541

1542 In CAEF, steps 2 and 3 can be handled using the addition executor, which has already learned the
1543 logic for addition, while the auxiliary operators needed for steps 1 and 4 are relatively simple to
1544 implement. The subtraction executor composer only needs to learn how to sequentially invoke these
1545 basic executors to perform subtraction.
1546

1547 A.5 PROMPTS USED IN BASELINE
1548

1549 Prompt used for LLaMA 3.1-8B pretrained model fine-tuned with LoRA:
1550

1551 *For addition, subtraction, multiplication, division:*
1552 Please calculate the expression.
1553 The expression is: {expr}.
1554 The final answer should be presented in integer form!
1555 Your output should be an integer.
1556 The answer is: {response}

1557 *For greater_than, less_than, equal:*
1558 Please judge the expression is true or false.
1559 The expression is: {expr}.
1560 The final answer should be True or False!
1561 Your output should be a word.
1562 The answer is: {response}
1563
1564

1565 Prompt used for LLaMA 3.1-8B-Instruct model:

1566
1567
1568
1569
1570
1571
1572
1573
1574
1575

For addition, subtraction, multiplication, division:

Please calculate the expression. The expression is: {expr}.

The final answer should be presented in integer form.

In your output, the final answer should be on its own line at the end, starting with 'Answer: '.

For greater_than, less_than, equal:

Please judge the expression is true or false. The expression is {expr}.

The final answer that you give should be true or false.

1576
1577

Prompt used for GPT-4o:

1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588

For addition, subtraction, multiplication, division:

Please calculate the expression. The expression is: {expr}.

The final answer should be presented in integer form.

In your output, the final answer should be on its own line at the end, starting with 'Answer: '.

For greater_than, less_than, equal:

Please judge the expression is true or false. The expression is {expr}.

The final answer that you give should be true or false.

1589
1590

A.6 FURTHER EXPERIMENT RESULTS ANALYSIS

1591
1592
1593
1594
1595
1596
1597

Using addition in the form of $a + b = c$ as an example, we generate the executor’s training dataset by dividing the expressions into equivalence classes based on the pair $(\text{len}(a), \text{len}(b))$, where 20 random arithmetic expressions are generated for each equivalence class. When the operand lengths are sufficiently large, 20 samples are sparse across the entire equivalence class space. However, the model still achieves high accuracy in tasks such as 100-digit addition, indicating that the LLM effectively learns the logic of the Turing machine’s transition function during training, thereby indirectly grasping the underlying logic of arithmetic computation.

1598
1599
1600
1601
1602
1603
1604
1605
1606

However, this sampling strategy alone can lead to poorer performance when operand lengths are relatively short, typically less than 10 digits, compared to longer operands. We believe this issue arises because the longest samples in the training set generally exceed 100 digits, and from the perspective of equivalence classes, the dataset becomes dominated by samples with operands of several dozen digits. Intuitively, although both cases involve a difference of 10 digits, the difference between 5 and 15 digits has a much larger impact than the difference between 90 and 100 digits, especially in the way the LLM perceives these distinctions. Therefore, in practice, we slightly increase the number of samples from equivalence classes with shorter operands. Additionally, for operators such as $==$, purely random sampling makes it difficult to obtain samples where the result is True, so some additional intervention is necessary.

1607
1608
1609

A.7 COMPUTATIONAL COMPLEXITY ANALYSIS

1610
1611
1612

For the seven operators implemented using the CAEF framework, we assume the longer operand has a length of d . Based on the computation mechanism of self-attention, the computational complexity of a single model inference is $O(d^2)$.

1613
1614
1615
1616
1617

For the *addition*, *greater_than*, *less_than*, and *equal* operators, the computation is performed digit by digit, requiring at most d model queries for a complete computation. Additionally, the aligner performs two representation conversions, resulting in a total of $d + 2$ model queries. Therefore, the overall computational complexity is $O(d^3)$.

1618
1619

For *subtraction*, the computational complexity of the auxiliary operators *Reflection* and *Left_mask* is also $O(d^3)$. Since subtraction involves one call each to *Reflection* and *Left_mask*, along with two calls to *addition*, the overall complexity remains $O(d^3)$.

For *multiplication*, the situation is slightly more complex. For a calculation of the form $a \times b$, we assume $\text{len}(a) = d_1$, $\text{len}(b) = d_2$, and $\text{len}(a \times b) = d_3$. The number of iterations in the loop is $b + 1$. During each iteration, *less_than* and two *addition* operations are performed, with the complexities as follows:

- For *less_than*, the longer operand has a length of d_2 , resulting in a total complexity of $O((b + 1)d_2^2)$ across all iterations.
- For the first *addition*, the longer operand has a length of d_3 , giving a total complexity of $O((b + 1)d_3^2)$ across all iterations.
- For the second *addition*, the longer operand again has a length of d_2 , resulting in a total complexity of $O((b + 1)d_2^2)$ across all iterations.

Thus, the overall complexity is the sum of these three parts, plus the two aligner conversions. The final computational complexity for multiplication is $O(bd_2^2 + bd_3^2)$.

For *division*, the situation is similar to multiplication. Assuming $\text{len}(a) = d_1$, $\text{len}(b) = d_2$, $a \div b = c$, and $\text{len}(c) = d_3$, the number of iterations in the loop is $c + 1$. The final computational complexity for division is $O(cd_1^2 + cd_3^2)$.

A.8 ATTEMPTS TO MERGE ALIGNER AND EXECUTOR

We attempt to combine the functionalities of the aligner and executor into one LoRA adapter. Table 4 shows the experimental results we obtained in the addition operator:

Table 4: Comparison of the results for merging the aligner and executor on the addition operator with the original CAEF method. The left side of the slash shows the results after merging the executor and aligner, while the right side presents the original results of CAEF.

Setting	5-digits	10-digits	50-digits	100-digits
executor & aligner	90.3/100.0	97.3/99.6	94.9/99.9	90.0/98.6
executor	100.0/100.0	99.9/100.0	99.6/99.9	97.4/99.6
aligner (I)	90.3/100.0	97.5/99.7	95.7/100.0	94.0/99.6
aligner (O)	100.0/100.0	99.9/99.9	99.5/100.0	98.0/99.4

From the results, we observe the following:

- The performance of the executor and aligner (O) shows a slight degradation compared to the original modular approach in most of the experimental settings.
- The aligner (I), however, experiences a significant performance drop. As a result, merging the aligner and executor leads to a substantial decline in the overall accuracy of addition.

Additionally, merging these two components introduces the challenge of determining the appropriate ratio for training samples from both parts. Therefore, based on the experimental results, we believe separating the executor and aligner remains the preferable approach.