

A REINFORCEMENT LEARNING ENVIRONMENT FOR MATHEMATICAL REASONING VIA PROGRAM SYNTHESIS

Anonymous authors

Paper under double-blind review

ABSTRACT

We convert the DeepMind Mathematics Dataset into a reinforcement learning environment by interpreting it as a program synthesis problem. Each action taken in the environment adds an operator or an input into a discrete compute graph. Graphs which compute correct answers yield positive reward, enabling optimization of a policy to construct compute graphs conditioned on problem statements. Baseline models are trained using Double DQN on various subsets of problem types, demonstrating the capability to learn to correctly construct graphs despite the challenges of combinatorial explosion and noisy rewards.

1 INTRODUCTION

The DeepMind Mathematics Dataset (Saxton et al., 2019) consists of synthetically generated math problems. They cover a range of problem types: numbers, comparison, measurement, arithmetic, algebra, polynomials, calculus, and probability. These problem types are arranged into a collection of 56 modules each containing different problem subtypes. The dataset provides the problems in the form of question-answer pairs represented as ASCII text.

The predominant algorithmic approach for learning to produce the answers conditioned on the question statements has been to train seq2seq models (Saxton et al., 2019)(Schlag et al., 2020)(Henighan et al., 2020). For some of the modules this approach yields very nearly 100% accuracy, however for other modules this does extremely poorly. For example, with this approach the best reported test accuracy for the module "numbers_list_prime_factors" is less than 25% (Henighan et al., 2020).

This poor performance is not surprising considering the nature of the problems on which this occurs. For instance, the module "numbers_is_prime" requires discriminating primes from non-primes. Apart from memorizing answers, there are only a few useful hacks that a learned model can easily pick up, for example, the fact that any even number (except for 2) is not prime. Algorithms to correctly test for primality require a sequence of divisions to be performed very precisely, and before a perfectly correct algorithm is obtained loss may not be much improved beyond random guessing. Thus, it's unclear if the gradient of loss obtained by comparing computed answers to correct answers could be sufficient to learn such algorithms.

A further difficulty with such systems is that they are unlikely to be sufficiently accurate or interpretable for real-world use. Models trained in this manner are effectively black boxes which only return answers. Probability estimates assigned to tokens sampled from the model can potentially provide clues about the likelihood of correctness (assuming they are well calibrated). However, the reasoning process by which the result was arrived at cannot be easily inspected.

We think it would be desirable to train neural networks to make use of existing programs for computing various mathematical operations (Andor et al., 2019)(Gupta et al., 2020). Just as human programmers use libraries when writing larger programs, learned algorithms should be able to build upon existing operators. With predefined operators available, the neural network wouldn't have to rediscover for itself algorithms and operators which are already well known, and could instead focus on learning how to compose them. Common compositional patterns can be further abstracted into operators of their own.

This is program synthesis: automatically constructing programs (Gulwani et al., 2017). Core to this is the idea *specification*: what the program is supposed to do. If the specification is sufficiently precise, then one may be able to prove that a given program meets it.

Unifying this with RL, we get a more formal problem statement: the learning problem is to learn a policy which maps a problem statement represented as text, to a program, represented as a graph composed of discrete operators. By conditioning on problem statements, we can learn a policy with reinforcement learning because the question-answer pairs can be used to provide a reward signal to adjust the parameters of the policy.

2 THE PROBLEM USING REINFORCEMENT LEARNING

To frame the problem in terms of reinforcement learning we need to define an environment. To explain how this is done we first need to explain how inputs are extracted from problem statements and to specify how the operators are defined.

2.1 INPUTS

In problems in the DeepMind Mathematics Dataset there are often parts of the text which are explicitly mathematical. We call these the inputs of the problems.

For example, in the following problem:

Let $h(t) = t^3 + t^2 + 1$. Let $v(d) = 6d^3 + 24d^2 + 4$. Let $w(j) = 4h(j) - v(j)$. What is the third derivative of $w(x)$ wrt x ?

The inputs are:

- 1) $h(t) = t^3 + t^2 + 1$
- 2) $v(d) = 6d^3 + 24d^2 + 4$
- 3) $w(j) = 4h(j) - v(j)$
- 4) $w(x)$
- 5) x

It turns out that for many problem types, a parser can be written which automatically extracts the inputs of the problem. By embedding such a parser within the environment implementation, given any problem statement the environment can determine which inputs are available. This will be used to help define the action space in section 2.3.

2.2 PREDEFINED OPERATORS

To utilize program synthesis, we need a starting list of primitive operators. Each of them takes some fixed number of inputs and returns a single output. For example, one operator called 'differentiate_wrt' (short for "differentiate with respect to" takes as input an expression and a variable, and it outputs the derivative of the expression with respect to the variable.

A full list of operators is defined in Appendix A.

2.3 DEFINING THE REINFORCEMENT LEARNING ENVIRONMENT

Observations from the environment consist of a representation of the question concatenated with a sequence of indices corresponding to actions taken previously in the current episode.

The representation of the question can either be provided in encoded form (i.e. as an array of indices into a vocabulary of tokens) or unencoded form (i.e. raw text) depending on how the environment is configured. If the environment is configured to return questions in encoded form, then it uses a predefined byte pair encoding constructed from a corpus of questions. If the representation is provided in encoded form then it is also padded up to a fixed maximum length. Note that the raw text of the question is available at every step regardless of how configurations are set (through the info object returned by the "step" method of the environment).

The action space is discrete and corresponds to the set of operators and available inputs which can be introduced into the graph at any given time. The graph is built up in breadth first order. Fixing the order of graph construction simplifies the action space because otherwise the actions would need to both specify which operator is being applied and where in the graph it is being introduced.

The maximum number of available inputs is set as a parameter of the environment. Thus if the maximum number of available inputs is set to n_inputs , and there are n_ops operators, then the action space is represented by the set of integers from 0 to $n_ops + n_inputs - 1$. The inputs are represented within the action space in the order in which they appear in the problem. For example, given the following problem:

Let $h(t) = t^3 + t^2 + 1$. Let $v(d) = 6d^3 + 24d^2 + 4$. Let $w(j) = 4h(j) - v(j)$. What is the third derivative of $w(x)$ wrt x ?

The action space would be as follows:

```
0 : 1st operator
1 : 2nd operator
...
n_ops - 1 : nth operator
n_ops + 0 : h(t) = t**3 + t**2 + 1
n_ops + 1 : v(d) = 6*d**3 + 24*d**2 + 4
n_ops + 2 : w(j) = 4*h(j) - v(j)
n_ops + 3 : w(x)
n_ops + 4 : x
n_ops + 5 : None
...
n_ops + n_inputs - 1 : None
```

Note that if less than the maximum number of possible inputs is present in a given problem, then the actions after the last available input are simply defined as None and are masked (see section 3.1.1).

The reward has a value of 1 if the compute graph computes the right answer and a value of 0 otherwise. Note that if the compute graph is incomplete then it computes the value None which always yields a reward of 0.

Here is an example trajectory from the environment in which the environment state is provided in unencoded format for legibility. Note that 5 is the action index corresponding to the "differentiate" operator and 14 is the action index corresponding to the 1st input (which in this case corresponds to $6k^2 - 101k + 2548$):

```
state t=0: What is the first derivative of  $6k^2 - 101k + 2548$ ?;
action t=0: 5 (differentiate)
state t=1: What is the first derivative of  $6k^2 - 101k + 2548$ ?; 5 (differentiate)
reward t=1: 0
action t=1: 14 (input)
state t=2: What is the first derivative of  $6k^2 - 101k + 2548$ ?; 5 (differentiate), 14 (input)
reward t=2: 1
```

Note that the environment will automatically terminate episodes with reward 0 if the configured maximum length is reached.

2.4 MODULES

For the purposes of this research, we restricted consideration to a subset of modules from the DeepMind Mathematics Dataset that have composed counterparts. The reason for this is that the composed problems require much larger graphs to reliably compute correct answers. This means that in a reinforcement learning context, these problems require longer action sequences to deliver optimal reward, and so constitute a more challenging reinforcement learning environment.

Furthermore, it is the case that some of the modules without uncomposed counterparts contain problems for which some of the relevant inputs are represented simply as words and so are difficult to parse as inputs. For example the module 'measurement_conversion' contains questions such as: "What is seven halves of a day in minutes?"




Figure 1: Type Hierarchy

A full list of supported modules is defined in Appendix A.

3 CHALLENGES

3.1 COMBINATORIAL EXPLOSION

The experiments we present here use 15 operators and require up to 3 inputs, which gives an action space of size 18. We also set a limit on graph size of at most 7 nodes. Even with this restriction, there are $18^7 \approx 6.12 * 10^8$ possible graphs. Given such a large state-space, naive search would be unlikely to stumble upon correct graphs. Here we present strategies for mitigating the impact of combinatorial explosion.

3.1.1 MASKING INVALID ACTIONS

To reduce the effective size of the search space we mask actions which are guaranteed to produce invalid graphs. In particular we introduce a type hierarchy and assign types to the parameter(s) and output of each operator. Also, each input is automatically assigned a type by the environment. Taken together, typed operators and typed inputs allow the application of type constraints limit which actions are valid at any given time. To implement this we define a method of the environment (called `compute_mask`) which generates a boolean masking vector over the action space. This method can be called by an agent before selecting an action to determine which actions are invalid, although formally the environment does not force this constraint. So invalid actions can still be taken, but they will result in graph that fails to compute a meaningful result (i.e. output of the graph will be "None").

Recall that actions add nodes to the graph in breadth-first order. Thus, each action after the first will correspond to a specific parameter of an operator that has already been applied. The type of the parameter being filled-in by the action determines the type constraint applied to that action. The form of the constraint is that if the parameter requires type X, then only types at or below X in the type hierarchy are permitted.

The type hierarchy consists of the following custom types: Equation (e.g. $2*x + y = 3$), Expression (e.g. $2*x + 1$), Function (e.g. $f(x) = 2*x + 1$), Value (e.g. 2), Variable (e.g. x), Rational (e.g. $1/2$). The hierarchy for the custom types is provided in Figure 1. We also utilize the following types which are built into Python: object, list, dict, and bool. Note that list, dict, bool and all the custom types are subclasses of object, and thus fall below it in the type hierarchy (Van Rossum & Drake, 2011).

In addition to type constraints, there are two other additional constraints implemented by masking. The first action always corresponds to the root node of the graph and has the additional constraint that it must correspond to an operator (i.e. it cannot be an input). There is also the additional constraint that any action defined as None is always masked. This occurs for problems which have less than the maximum number of inputs in which case the action space is padded up to a fixed size with actions defined as None.

3.1.2 SUBGRAPH ABSTRACTION

Another strategy for addressing combinatorial explosion which can be applied but was not implemented in our experiments is the process of abstracting frequently rewarded subgraphs into new operators. The concept is to identify subgraphs which frequently occur in graphs that yield reward, then to identify the inputs and output of that subgraph such that it could be redefined as a new operator. This is known more broadly in the computer science literature as Frequent Subgraph Mining (Jiang et al., 2013).

For example:

```
differentiate_wrt(differentiate_wrt(Expression('-3*z**5 + 13*z**3 + 41*z
**2'),Variable('z')),Variable('z')) reward = 1
```

This subgraph can be abstracted into a new operator as follows:

```
diff_wrt_2(p0, p1) = differentiate_wrt(differentiate_wrt(p0, p1), p1)
```

New operators defined in this way could be introduced into the action space and would provide the possibility of finding shorter graphs to compute the same operators. This is similar to the operator abstraction process in Dreamcoder (?). This would increase the size of the action space and hence the number of possible graphs, however it would reduce the required length of frequently rewarded graphs and hence would have a multiplicative effect in reducing the effective size of the search space.

3.2 NOISY REWARD

An additional challenge with this environment is that the rewards are noisy with respect to graph correctness. In other words, a particular graph might compute the right answer for a given question but do so in the wrong way. For example:

```
Is 5340 a multiple of 10?; not(is_prime(Value('10'))) = True, reward: 1
```

If the inputs of the problem were changed the same graph may no longer compute the right answer because it is simply performing the wrong operations and so cannot reliably generalize to different inputs. This effectively means that a positive reward from the environment does not guarantee that the constructed graph is correct.

4 APPROACH

We use Double DQN (van Hasselt et al., 2015) to learn a value function that maps a state and an action to the expected discounted sum of rewards. An off-policy learning algorithm is used because it simplifies the maintenance of exploration which is critical in the presence of noisy rewards. Epsilon-greedy exploration is utilized with a step-wise linear annealing schedule on the value of epsilon.

We also apply prioritized experience replay (Schaul et al., 2016) to improve the efficiency of learning. We sample steps from replay memory with priority directly proportional to the most recently computed TD-error. After every training batch, we re-compute replay priorities for both the steps used to construct the batch and an additional random sample of steps from the replay buffer. The additional random sample is taken to encourage all replay priorities to remain consistent with the current model parameters, even if they haven't been used to construct a training batch in some time.

We initialize the replay buffer with trajectories (containing up to 50k steps) collected by a uniform random policy, where a one-to-one balance is kept between trajectories with positive reward and trajectories with 0 reward. This balancing is done to prevent a heavy skew towards trajectories with reward 0, due to reward sparsity in the environment. After an initial period of training on the replay buffer as initialized, new experience is continuously incorporated into the replay buffer. However, a one-to-one balance between trajectories with positive reward and 0 reward is maintained since otherwise trajectories with 0 reward would overwhelm the replay buffer.

The model used in experiments is a transformer encoder (Vaswani et al., 2017) with 6 encoder blocks followed by 2 dense layers. We use 4 attention heads and a hidden layer size of 256. The dense blocks use a hidden layer size of 256. We apply dropout (Srivastava et al., 2014) of 0.1 in the

Transformer blocks and between the dense layers. We use a learning rate of 5×10^{-5} and a batch size of 512. The epsilon value for epsilon-greedy exploration is initialized to 0.4 and is linearly annealed to 0.05 by an increment of 2.5×10^{-5} per step. The full set of hyperparameters used to produce the results reported here is provided in the code (https://github.com/joepalermo/dm_math_solvers/blob/master/hparams-for-paper.cfg).

5 RESULTS

We conduct experiments on different subsets of modules to evaluate how it affects model performance. For each experiment we run 5 trials with the same hyperparameters and different random seeds (as recommended in Henderson et al. 2017 (Henderson et al., 2019)). We sample 1.01 million examples from each module under consideration with 800k/200k/10k across train/validation/test. The trials are run for 50k steps each except for the trials in the "Interference" experiment (see below) which are run for 100k steps.

In Table 1 we report mean test reward for the median trial. We define the median trial as the trial for which the mean test reward across modules is the median across trials. In appendix B we report results from all trials and also provide validation curves corresponding to the results reported in Table 1.

In our experiments we consider only the uncomposed modules. Composed modules are distinguished by containing much longer problems (they are the modules suffixed by "_composed" or "_compose"). Based on preliminary experiments, the methods employed here do not perform well on the composed modules due to the increased challenge of combinatorial explosion. Note also that in the experiments we report here the "calculus_differentiate" module is filtered to remove multi-variate problems as this reduces the maximum required graph size (the option to remove this filter is provided as a hyperparameter of the environment).

In these experiments we also limit the action space to contain only the 15 operators required to successfully compute correct graphs on the modules selected. The list of operators used is provided in appendix B.2.

In our first experiment we train on all uncomposed modules simultaneously and observe significant interference between modules reflected by lower final test performance than in smaller subsets of modules. We hypothesize that by including multiple modules in which answers to problems are similarly expressed, interference between modules is magnified. For example, the modules "numbers_is_factor" and "numbers_is_prime" both have true/false answers. In the case in which both modules are trained on simultaneously the operators "divides" and "is_prime" respectively will be frequently misused during exploration (e.g. Is 5340 a multiple of 10? `not(is_prime(Value('10')))` = True, which results in a positive reward).

To investigate this hypothesis we select two additional subsets of modules. The first of these module subsets contains "numbers_is_factor" and "numbers_is_prime", and we refer to it as the "Interference" experiment. The second contains "numbers_is_prime", "numbers_list_prime_factors", "numbers_calculus_differentiate", "numbers_div_remainder", and "numbers_gcd", and we refer to it as the "No Interference" experiment. The modules in the "Interference" experiment are selected because the form of the answers in those modules are both true/false and so will result in the type of collisions described above (i.e. noisy rewards). The modules in the "No Interference" experiment are selected because the form of the answers in those modules is such that they are unlikely to result in the type of collisions we describe above. However, notably the final results show that "numbers_is_prime" had better performance in the "Interference" experiment which is evidence against the hypothesis.

Table 1: Test reward per module for the median trial of each experiment.

Test Results			
Module	Interference	No Interference	All Uncomposed Modules
numbers__is_factor	0.7800	-	0.3669
numbers__is_prime	1.000	0.7382	0.6567
numbers__list_prime_factors	-	1.000	1.000
calculus__differentiate	-	0.8511	0.3350
polynomials__evaluate	-	-	0.9517
numbers__div_remainder	-	1.000	0.8159
numbers__gcd	-	1.000	0.9990
numbers__lcm	-	-	1.000
algebra__linear_1d	-	-	0.1124
algebra__polynomial_roots	-	-	0.8253
algebra__linear_2d	-	-	0.2660
Mean Reward across Modules	0.8900	0.9179	0.6830

6 DISCUSSION

There are two main contributions of this paper. We have introduced a new reinforcement learning environment by interpreting the DeepMind Mathematics Dataset as a problem in program synthesis. We have also trained a baseline model on several subsets of the uncomposed modules despite the dual challenge of combinatorial explosion and noisy rewards.

From preliminary experiments it’s clear that performance on significantly longer graphs is poor, however we believe that by implementing subgraph abstraction (as described in section 3.1.2) performance on longer graphs could be significantly improved. Furthermore we suspect that content based attention (as in (Bahdanau et al., 2016)) could be a useful architectural component to integrate into the learned model due to the dynamic nature of the action space. We consider these to be interesting topics for future research.

CODE

We provide a light-weight repository containing an implementation of the reinforcement learning environment and setup instructions:

https://anonymous.4open.science/r/math_prog_synth_env-D46E/

We also provide the full code used to produce the results reported here:

https://anonymous.4open.science/r/dm_math_solvers-C310/

ACKNOWLEDGMENTS

The authors would like to thank Alex Krizhevsky, Rayhane Mama, Hashiam Kadhim, Marc Tyndel, and Ragavan Thurairatnam for helpful discussions.

REFERENCES

- Daniel Andor, Luheng He, Kenton Lee, and Emily Pitler. Giving bert a calculator: Finding operations and arguments with reading comprehension, 2019.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate, 2016.
- Sumit Gulwani, Alex Polozov, and Rishabh Singh. *Program Synthesis*, volume 4. NOW, August 2017. URL <https://www.microsoft.com/en-us/research/publication/program-synthesis/>.
- Nitish Gupta, Kevin Lin, Dan Roth, Sameer Singh, and Matt Gardner. Neural module networks for reasoning over text, 2020.
- Peter Henderson, Riashat Islam, Philip Bachman, Joelle Pineau, Doina Precup, and David Meger. Deep reinforcement learning that matters, 2019.
- Tom Henighan, Jared Kaplan, Mor Katz, Mark Chen, Christopher Hesse, Jacob Jackson, Heewoo Jun, Tom B. Brown, Prafulla Dhariwal, Scott Gray, Chris Hallacy, Benjamin Mann, Alec Radford, Aditya Ramesh, Nick Ryder, Daniel M. Ziegler, John Schulman, Dario Amodei, and Sam McCandlish. Scaling laws for autoregressive generative modeling, 2020.
- Chuntao Jiang, Frans Coenen, and Michele Zito. A survey of frequent subgraph mining algorithms. *The Knowledge Engineering Review*, 28(1):75–105, 2013. doi: 10.1017/S0269888912000331.
- David Saxton, Edward Grefenstette, Felix Hill, and Pushmeet Kohli. Analysing mathematical reasoning abilities of neural models, 2019.
- Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay, 2016.
- Imanol Schlag, Paul Smolensky, Roland Fernandez, Nebojsa Jojic, Jürgen Schmidhuber, and Jianfeng Gao. Enhancing the transformer with explicit relational encoding for math problem solving, 2020.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 15(1): 1929–1958, January 2014. ISSN 1532-4435.
- Hado van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning, 2015.
- G. Van Rossum and F.L. Drake. *The Python Language Reference Manual*. Python Manual. Network Theory Limited, 2011. ISBN 9781906966140. URL <https://books.google.ca/books?id=Ut4BuQAACAAJ>.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need, 2017.

A ENVIRONMENT

A.1 PREDEFINED OPERATORS

We provide here a list of the predefined operators. The numbers associated to them indicates their index in the action space. They are also provided along with a type signature (e.g. $f(x: \text{type_of_x}) \rightarrow \text{return_type}$).

0. `lookup_value(mapping: Dict[Variable: Value], key: Variable) -> object`

Given a dictionary and a key to query it, it returns the corresponding value.

1. `solve_system(system: List[Equation]) -> Dict[Variable: Value]`

Given a system of equations returns a dictionary mapping variables to values.

2. `append(system: List[Equation], equation: Equation) -> List[Equation]`

Given a list of equations, appends a new equation to the end of that list.

3. `append_to_empty_list(equation: Equation) -> List[Equation]`

Returns a new list containing only the given equation.

4. `factor(inpt: Expression) -> Expression`

Converts a polynomial into irreducible factors over rational numbers.

5. `differentiate(expression: Expression) -> Expression`

Returns the first derivative of a polynomial. (This operator assumes univariate)

6. `mod(numerator: Value, denominator: Value) -> Value`

Returns the remainder of the numerator divided by the denominator

7. `gcd(x: Value, y: Value) -> Value`

Returns the greatest common divisor of x and y

8. `divides(numerator: Value, denominator: Value) -> bool`

Returns True if denominator is divisible by numerator

9. `is_prime(x: Value) -> bool`

Returns True if x is prime and False otherwise

10. `lcm(x: Value, y: Value) -> Value`

Returns the least common multiple of x and y

11. `lcd(x: Rational, y: Rational) -> Value`

Given two rationals return the least common denominator

12. `prime_factors(n: Value) -> Set[Value]`

Returns a set of all prime factors of n

13. `evaluate_function(function_definition: Function, function_argument : Expression) -> Value`

Evaluates a function by substituting the variable in function_definition by the function_argument. function_argument can either look like '2' or 'f(2)'

14. `not_op(x: bool) -> bool`
Returns the inverse of a boolean.
15. `differentiate_wrt(expression: Expression, variable: Variable) -> Expression`
Returns the first derivative of an expression with respect to a given variable.
16. `make_equation(expression1: Expression, expression2: Expression) -> Equation`
Returns an equation where expression1 is set equal to expression2.
17. `simplify(inpt: object) -> object`
Returns a simplification of inpt based on sympy heuristics.
18. `make_function(expression1: Expression, expression2: Expression) -> Function`
Returns a function where expression1 is set to be equal to expression2.
19. `replace_arg(function: Function, var: Variable) -> Function`
Replaces the argument in function with the given variable.
20. `lookup_value_equation(mapping: Dict[Variable: Value], key: Variable) -> Equation`
Given a dictionary and a key to query it, it returns an equation of "key = value"
21. `extract_isolated_variable(equation: Equation) -> Variable`
Given an equation it returns the isolated variable.
22. `substitution_left_to_right(arb: object, eq: Equation) -> object`
Returns the arb with all found instances of the equation's left hand side substituted by the equation's right hand side.

A.2 SUPPORTED MODULES

Here is the full list of supported modules:

```
numbers__is_factor
numbers__is_prime
numbers__list_prime_factors
calculus__differentiate
polynomials__evaluate
numbers__div_remainder
numbers__gcd
numbers__lcm
algebra__linear_1d
algebra__polynomial_roots
algebra__linear_2d
algebra__linear_1d_composed
algebra__linear_2d_composed
algebra__polynomial_roots_composed
calculus__differentiate_composed
numbers__div_remainder_composed
numbers__gcd_composed
numbers__is_factor_composed
numbers__is_prime_composed
numbers__lcm_composed
numbers__list_prime_factors_composed
polynomials__evaluate_composed
polynomials__compose
```

B EXPERIMENTS

B.1 FULL TEST RESULTS

Table 2: Test reward per module for runs on all uncomposed modules

Full Test Results for All Uncomposed Modules					
Module	Run 1	Run 2	Run 3	Run 4	Run 5
numbers__is_factor	0.2907	0.5239	0.2798	0.4807	0.3669
numbers__is_prime	0.7294	1.000	0.4860	1.000	0.6567
numbers__list_prime_factors	1.000	1.000	1.000	1.000	1.000
calculus__differentiate	0.3238	0.3611	0.6095	0.3377	0.3350
polynomials__evaluate	0.9990	0.9950	0.0000	0.9942	0.9517
numbers__div_remainder	1.000	1.000	0.7944	1.000	1.000
numbers__gcd	0.9981	0.8132	0.9317	1.000	0.9990
numbers__lcm	1.000	1.000	1.000	1.000	1.000
algebra__linear_1d	0.7470	0.7470	1.000	1.000	0.1124
algebra__polynomial_roots	0.4443	0.7159	0.9457	0.9143	0.8253
algebra__linear_2d	0.8110	0.1400	0.9600	1.000	0.2660
Mean reward across modules	0.6973	0.6816	0.6629	0.8843	0.6830

Table 3: Test reward per module for "Interference" experiment

Full Test Results for "Interference" Experiment					
Module	Run 1	Run 2	Run 3	Run 4	Run 5
numbers__is_factor	0.5155	0.7893	0.9383	0.9922	0.7800
numbers__is_prime	0.5229	0.9980	0.5258	1.000	1.000
Mean reward across modules	0.5192	0.8937	0.7321	0.9961	0.8900

Table 4: Test reward per module for "No Interference" experiment

Full Test Results for "No Interference" Experiment					
Module	Run 1	Run 2	Run 3	Run 4	Run 5
numbers__is_prime	0.5563	0.5400	0.7694	1.000	0.7382
numbers__list_prime_factors	1.000	1.000	1.000	1.000	1.000
calculus__differentiate	0.8112	0.9572	1.000	0.9653	0.8511
numbers__div_remainder	1.000	1.000	1.000	1.000	1.000
numbers__gcd	0.9990	1.000	1.000	1.000	1.000
Mean reward across modules	0.8733	0.8994	0.9539	0.9931	0.9179

B.2 OPERATORS SELECTED FOR EXPERIMENTS

1. lookup_value
2. solve_system
3. append
4. append_to_empty_list
5. factor
6. differentiate
7. mod
8. gcd
9. divides
10. is_prime
11. lcm

12. lcd
13. prime_factors
14. evaluate_function
15. not_op

B.3 VALIDATION CURVES

In the below graphs the dark centre line shows the median of the five trials and the shaded area bounds the 10th and 90th percentiles based on linear interpolation.

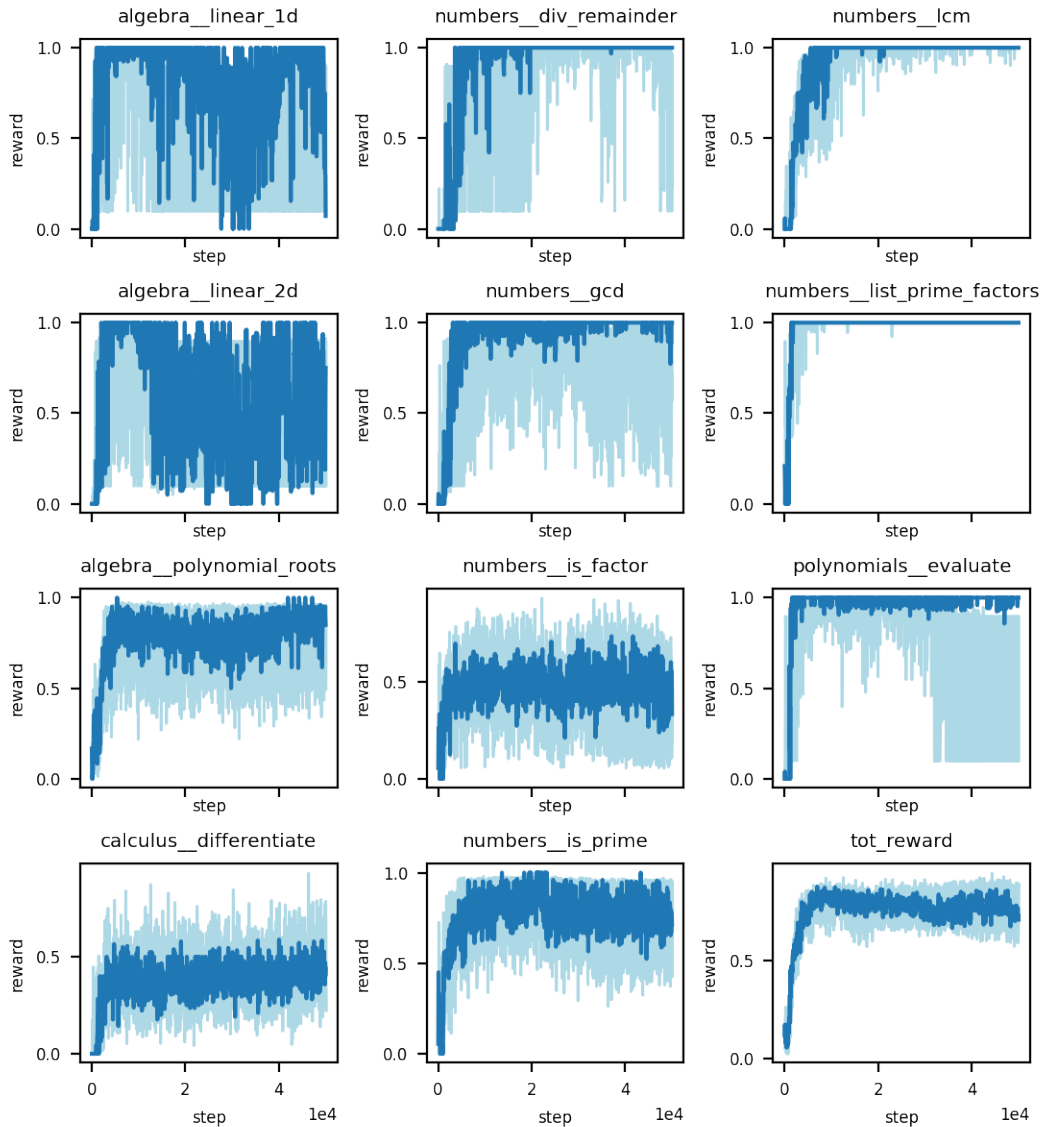


Figure 2: Reward for all uncomposed modules over 50000 steps.




Figure 3: Reward for "Interference" experiment over 100000 steps

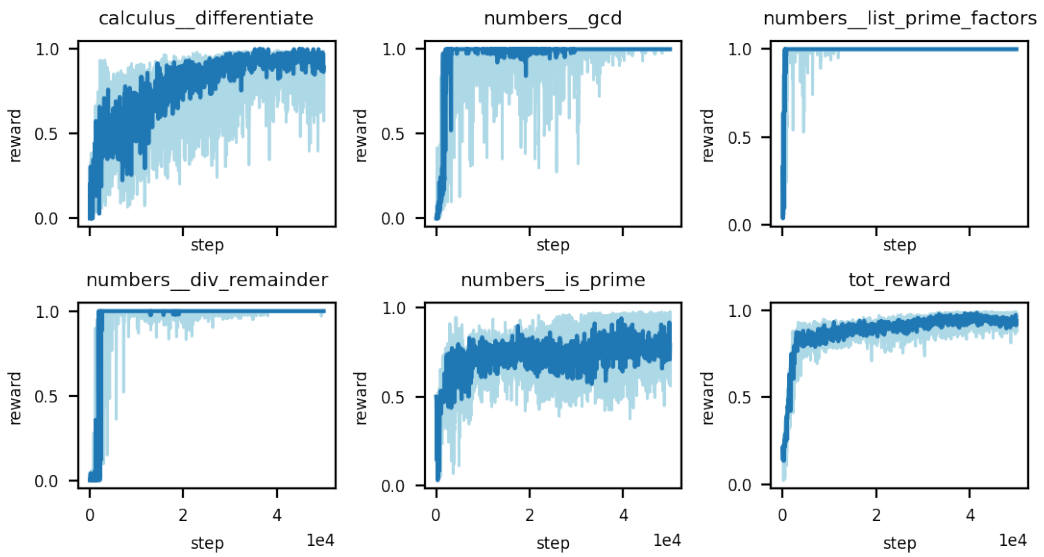


Figure 4: Reward for "No Interference" experiment over 50000 steps.