

# KNOSSOS: COMPILING AI WITH AI

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

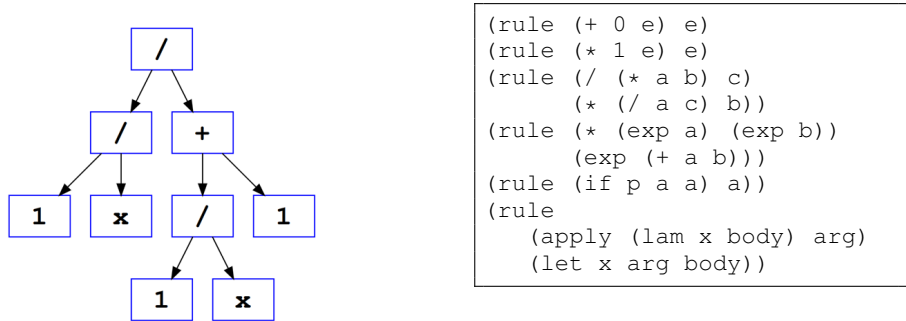
Machine learning workloads are often expensive to train, taking weeks to converge. The current generation of frameworks relies on custom back-ends in order to achieve efficiency, making it impractical to train models on less common hardware where no such back-ends exist. Knossos builds on recent work that avoids the need for hand-written libraries, instead compiles machine learning models in much the same way one would compile other kinds of software. In order to make the resulting code efficient, the Knossos compiler directly optimises the abstract syntax tree of the program. However in contrast to traditional compilers that employ hand-written optimisation passes, we take a rewriting approach driven by the  $A^*$  search algorithm and a learned value function that evaluates future potential cost reduction of taking various rewriting actions to the program. We show that Knossos can automatically learn optimisations that past compilers had to implement by hand. Furthermore, we demonstrate that Knossos can achieve wall time reduction compared to a hand-tuned compiler on a suite of machine learning programs, including basic linear algebra and convolutional networks. The Knossos compiler has minimal dependencies and can be used on any architecture that supports a C++ toolchain. Since cost model the proposed algorithm optimises can be tailored to a particular hardware architecture, the proposed approach can potentially be applied to a variety of hardware.

## 1 INTRODUCTION

While the development of any kind of software can benefit from compilers able to produce fast code, runtime efficiency is particularly important for modern machine learning. In particular, because modern models they can take weeks to train (OpenAI, 2018), compiler optimisations that lead to execution speed-ups are of huge value. In parallel, machine learning is being deployed on a variety of diverse devices ranging from wearables to huge clusters of powerful GPUs. Since each architecture has different performance profile and requires different code optimisations, it is difficult to provide tooling that works fast on all of them.

Traditionally, the tension between performance and interoperability is resolved by machine learning frameworks (Paszke et al., 2017; Abadi et al., 2016). In these frameworks, while code execution is outsourced to hardware-specific back-ends such as XLA (XLA authors, 2016). While this approach has seen huge initial success, the cost of providing customised back-ends for each target architecture is prohibitive. Moreover, the frameworks also custom front-ends that require the programmer to specify the model being trained as a compute graph. Since the compute graph has semantics separate from the host programming language, this process is often error-prone and time-consuming. In order to address these obstacles, a new generation of tools has recently appeared that transform machine learning code using the same techniques that have been used for compiling traditional software. The need for a separate front-end API for machine learning operations is eliminated by including automatic differentiation as a first-class feature of the compiled language (Innes et al., 2019; Frostig et al., 2018). Instead of custom back-ends, modern machine learning compilers use an intermediate representation and perform extensive code optimisations (Innes et al., 2019; Frostig et al., 2018; van Merriënboer et al., 2018; Wei et al., 2018; Sotoudeh et al., 2019; Rotem et al., 2018). In addition, program optimisation is being modelled as a machine learning task itself, with the compiler learning how to perform rewrites (Chen et al., 2018b;a).

Knossos expands on this line of work. The Knossos system includes a compiler which combines efficient program optimisation with an intermediate representation (IR) designed with machine learning



(a) Sample arithmetic expression tree (MDP state). The tree corresponds to  $\frac{1/x}{1/x+1}$ . (b) Example rewrite rules. Each rule defines a class of MDP transitions. See Appendix C for the full list of rules.

Figure 1: The Knossos MDP.

in mind. We formalize program optimisation as a finite-horizon Markov Decision Process (MDP), with the reward signal determined by the cost of executing a program. By solving this MDP, we are able to produce fast code tailor-made for any given task and architecture, without relying on backend-specific hand-written libraries. Knossos works by re-writing programs written in an intermediate representation (IR). Akin to JAX (Frostig et al., 2018) and Zygote (Innes et al., 2019), all Knossos functions are potentially differentiable, avoiding the syntactic awkwardness that arises from embedding a differentiable program in a host language. The IR can then be transpiled, allowing it to run on any platform that supports a C++ toolchain. This allows Knossos code to be seamlessly deployed on specialized or embedded hardware without the need of manual tuning, both for training and for deployment of models, enabling a much broader user base than competing approaches.

To our knowledge, Knossos is the first compiler that combines RL-based program optimisation, first-class support for deep learning primitives and the ability to target any architecture supporting the C++ toolchain. We defer detailed scope comparisons with prior work to Section 4. We empirically demonstrate the benefits of our program optimisation in Section 5, showing that Knossos was able to automatically learn *loop fusion*, a type of compiler optimisation that previously had to be applied manually.

## 2 CODE OPTIMISATION AS A REINFORCEMENT LEARNING PROBLEM

We model code optimisation as a finite-horizon Markov Decision Process (MDP). An MDP is defined (Puterman, 2014; Sutton & Barto, 2018) as a tuple  $(S, A, T, R, H, p_0)$ , where  $S$  denotes the state space,  $A$  denotes the action space,  $T$  denotes the transition dynamics,  $R$  denotes the rewards,  $H$  is the maximum time budget allowed to solve the problem (the horizon) and  $p_0$  is a fixed probability distribution over initial states. We provide a detailed description of the states, transitions and rewards later on this section.

**States and transitions** An MDP state  $s = (e_s, t_s)$  consists of a Knossos program (or expression)  $e \in E$  and the remaining time budget  $t_s \in [0, 1, \dots, H]$  (i.e., the number of remaining steps), where  $H$  is the maximum budget. Any state with  $t_s = 0$  is terminating. The initial state distribution  $p_0$  models the expressions that the RL agent is likely to be asked to optimize. A sample Knossos expression is shown in Fig. 1a. The action set  $A$  corresponds to different possible ways of rewriting the same expression (see Fig. 1b). The transition function  $T : S \times A \rightarrow S$  returns the next state after taking an action. For example, the first rule in Fig. 1b says that adding zero to any expression can be simplified to the expression itself. Once the action is chosen, the transition is deterministic. Because rewrite rules can be applied to different subexpressions, we specify  $A$  using generic rewrite rules, which are applied by pattern matching. There are over 50 rules like this – we provide the details in Appendix C. An essential feature of the rewrites is that they do not change the meaning of the program, i.e. by simplifying from one expression to another we also implicitly generate a proof that the expressions are equivalent.

**Policies and Value Functions** The RL agent maintains a policy  $\pi(a|s)$ , which defines the probability of taking an action in state  $s$  given there are  $t_s$  steps remaining till the total time budget is exhausted. A policy  $\pi$  generates rollouts  $\tau_\pi$ . A rollout  $\tau_\pi$  is defined as a sequence of states, actions and rewards obtained from the MDP  $\tau_\pi = (s_1, a_1, r_1, s_2, a_2, r_2, \dots, s_H, r_H)$ . Since the policy  $\pi$  can be stochastic, it is modelled as a random variable. The goal of RL agent is to find an optimal policy  $\pi^* = \arg \max_\pi J_\pi$ , which attains the best possible return. The return is defined as  $J_\pi = E_{\tau_\pi} \left[ \sum_{t=0}^{t=H-1} R(s_t, s_{t+1}) \right]$ . Given a policy and the number of timesteps  $t$  remaining till the end of episode, we define a value function  $V(s) = E_\tau \left[ \sum_{i=0}^{t_s-1} R(s_i, s_{i+1}) \mid s_0 = s \right]$ , where  $t_s$  denotes the remaining time budget at state  $s$ . The optimal value function  $V^*$  is defined as the value function of an optimal policy  $\pi^*$ .

**Rewards and the Cost Model** We assume access to a function  $c(s)$ , which provides the *cost model*, i.e. the computational cost of running  $e_s$ , the expression represented by state  $s = (e_s, t_s)$ , on representative inputs. While developing a perfect cost models is theoretically impossible due to the intractability of the halting problem (Turing, 1937), very good cost models exist for the particular subset of programs that compilers are asked to optimise. The ideal cost model  $c_B$  would correspond to the run-time of the program on typical inputs, but evaluating costs by benchmarking is very computationally intensive. In practice, one can often find a surrogate cost function such that for most initial programs  $s_0$ , the state that is reachable from  $s_0$  and minimizes the surrogate cost function  $c$  agrees with that for the ideal cost function  $c_B$ , that is,

$$\arg \min_{s \sim s_0} c(s) = \arg \min_{s \sim s_0} c_B(s), \quad (1)$$

which is much easier to acquire. In other words, the cost function  $c$  does not have to produce the same run-time but the same minimum over programs. We show experimentally in Section 5 that it is indeed possible to reduce the wall clock time of running a program by optimising such a proxy cost model. Knossos has a modular architecture, making it easy to change the cost function. This makes it possible to quickly re-tune Knossos programs for any target hardware. We stress that the formalism allows us to find optimisations even in case getting to the optimized version of the code requires using intermediate programs of higher cost (see Fig. 8).

Our reward function is based on this cost model. The rewards  $R(s_1, s_2) = c(s_2) - c(s_1)$  correspond to the attained reduction in cost when rewriting expression  $e_{s_1}$  into  $e_{s_2}$ . This formulation ensures that return  $J_\pi$  equals the total cost reduction attained along the length of the rollout  $\tau$ . Similarly, the value function corresponds to the expected cost reduction under the current policy. Since our MDP includes a ‘no-op’ rewrite rule that allows us to keep the current expression and hence the cost, the optimal value function is monotonic in  $t$  i.e.

$$V^*((e, t')) \geq V^*((e, t)) \quad \text{for any } e, t' \geq t. \quad (2)$$

### 3 TRAINING THE RL AGENT

**Hard and Easy Aspects of Rewriting** There are two main ways in which the task of rewriting expressions is *more challenging* than typical RL benchmarks. First, the allowed set of actions not only changes from state to state, but grows with the size of the expression. This makes exploration hard. Second, the states of the MDP, which correspond to the expressions being rewritten, are represented as graphs, whose size and topology varies as optimisation progresses. This is unlike traditional deep Reinforcement Learning (Mnih et al., 2013), which learns either from pixels or from data of fixed shape. While the rewriting task has many features that make it difficult, *it is also easier* than many traditional RL tasks for three reasons. First, MDP transitions are completely deterministic. Second, the task has a large degree of locality in the sense that the performance of a program can often be substantially improved by optimising its parts separately. Third, we can generate state transitions in any order convenient to us, as opposed to the traditional RL setting, where we are constrained by the order imposed by the environment. Overall, we have a problem similar to traditional planning, but which requires us to generalise well in order to obtain competitive solutions. To do this, Knossos uses a custom RL algorithm, based on  $A^*$  search supported by value function learned with a graph neural networks (Algorithm 1). We describe how to obtain the heuristic in Section 3.2, and the search algorithm in Section 3.1.

**Algorithm 1** Knossos

---

```

function TRAIN( $E_{\text{train}}, H, M$ )
  Input:  $E_{\text{train}}$ : Set of expressions to train
           $H$ : Maximum depth of the search
           $M$ : Number of epochs
  Initialize a value function  $\hat{V}$ 
  for  $M$  iterations do
    for  $e \in E_{\text{train}}$  do
       $S, V^{\text{target}} \leftarrow A^*(e, \hat{V}, H)$  ▷ Optimize  $e$  and store the learned values  $V^{\text{target}}$ 
       $\hat{V} \leftarrow \text{FIT}(S, V^{\text{target}}, \hat{V})$  ▷ Fit the value function - see Section 3.2
    end for
  end for
end function

function OPTIMIZE( $e, \hat{V}, H$ )
  Input:  $e$ : Expression to optimize
           $\hat{V}$ : Learned value function
           $H$ : Maximum depth of the search
   $S, - \leftarrow A^*(e, \hat{V}, H)$  ▷ Optimize  $e$ 
  return  $\arg \min_{s \in S} c(s)$  ▷ Return the best expression found during the process
end function

```

---

**Algorithm 2** Deep  $A^*$  search

---

```

function  $A^*(s_0, \hat{V}, H)$ 
  Input:  $s_0$ : Expression to optimize
           $\hat{V}$ : Learned value function
           $H$ : Maximum depth of the search
   $s_0 \leftarrow (e_0, H)$ 
   $O \leftarrow \{s_0\}$  ▷  $O$  is an open list storing all unexplored states
   $C \leftarrow \{s_0\}$  ▷  $C$  is a closed list storing all states ▷  $t$  represents the step budget left for the state.
  while not TERM_CONDITION() and  $O$  is not empty do
     $s \leftarrow \arg \max_{s \in O} f(s)$ . ▷  $f$  is the heuristic function defined in equation 3.
     $O \leftarrow O \setminus \{s\}$ 
    for all  $a \in A(s)$  do
       $n \leftarrow T(s, a)$  ▷ Obtain the next state  $n = (e_n, t_n)$ 
      if  $(e_n, t) \notin C (\forall t)$  or  $(e_n, t) \in C$  and  $t_n > t$  then
        if  $t_n > 0$  then
           $O \leftarrow O \cup \{n\}$ 
        end if
        if  $n \notin C$  then
           $C \leftarrow C \cup \{n\}$ 
        end if
      end if
    end for
  end while
  for all  $s \in C$  do
     $V^{\text{target}}(s) \leftarrow \max_{s' \in C \text{DISTANCE}(s, s') \leq t_s} (c(s) - c(s'))$ 
    ▷ Empirical estimate of maximum cost reduction achievable from  $s$ 
  end for
  return  $C, t, V^{\text{target}}$ 
end function

```

---




Figure 2: Graph neural network

### 3.1 SEARCHING THE SPACE OF REWRITES WITH $A^*$

We use the  $A^*$  algorithm (Hart et al., 1968) both to train the compiler and to deploy it.  $A^*$  maintains two priority queues. One queue ( $O$ ) stores the frontier, i.e. states from which transitions have not been explored yet. The other one ( $C$ ) stores the states visited so far and is used to avoid exploring the same path twice. The states are explored in the order induced by the  $A^*$  heuristic, which in our case corresponds to the learned value function  $\hat{V}$ , obtained from previous iterations. In particular, node priority is set as follows:

$$f(s) = f(\underbrace{(e_s, t_s)}_s) = \hat{V}(s) + \underbrace{\sum_{i=0}^{H-t_s-1} R(s_i, s_{i+1})}_{c(s_0) - c(s)}. \quad (3)$$

Here,  $\hat{V}(s)$  is the estimated future cost reduction obtained from state  $s$  within  $t$  remaining time-steps. The quantity  $c(s_0) - c(s)$  corresponds to the cost reduction that has already been achieved by time  $t$ , measured against the cost of the initial expression. Thus,  $f(s)$  is an estimate of the maximum possible cost improvement from a trajectory passing through state  $s$  at time  $t$ .

After the search, we compute the empirical estimate of the maximum cost reduction achievable ( $V^{\text{target}}(s)$ ) for each visited state. The estimated value of  $s$  with  $t_s$  timesteps is the maximum cost reduction found from  $s$  within  $t_s$  steps.  $\text{DISTANCE}(s, s')$  in Algorithm 2 is the number of steps required to reach  $s'$  from  $s$ . The algorithm stops after the value function was evaluated a set number of times. In the code this is represented with the function `TERM-CONDITION`.

$A^*$  is well-suited for the rewriting task because it exploits its characteristic features. In particular, it exploits determinism by assuming that a cost reduction achievable once can always be achieved again. It exploits the availability of reset by considering nodes in the order defined by the heuristic function. It exploits locality by preferring re-writes that need a small number of rule applications. Before deciding on  $A^*$ , we also performed experiments with Monte Carlo Tree Search (MCTS). MCTS does not make use of reset and had worse empirical performance (see Appendix D for details).

### 3.2 LEARNING VALUE FUNCTIONS WITH GRAPH NEURAL NETWORKS

States in the Knossos MDP correspond to computation graphs. In order to apply deep RL to these graphs, we need to be able to construct differentiable embeddings of them. To do this, we employ Graph Neural Networks based on Gated Recurrent Units (Li et al., 2016). During the forward pass, the GNN begins with an initial embedding of the graph nodes. It then iteratively applies a diffusion process to the graph. At each step, the obtained representation is fed into a gated recurrent unit (GRU). The process implicitly encodes the edge structure of the graph in the obtained representation.

**Graph representation** We represent a Knossos expression as a graph. The graph nodes correspond to subexpressions (see Fig. 1a). The graph edges are of two kinds. The first kind of edges connects the nodes with their parents. In addition, we use another kind of edges, which is used to explicitly provide the information that two subexpressions are identical. See Table 2b for a list of all edge types. Edges can be directed or undirected, with the directed edges going in opposite ways considered different.

**Graph neural network** To compute the value function for an expression  $e$  and time budget  $t$ , we start from computing the initial node embedding  $h_v^0 \in \mathbb{R}^d$  for all node  $v \in N(e)$ , where  $N(e)$  is the set of vertices in expression  $e$ . The initial node embedding consists of a one-hot encoding of the node type (constant, variable, etc) followed by zero padding.

This embedding is then fed into the following recurrent computation (see Fig. 2a):

$$h_v^{t+1} = f \left( h_v^t, \bigoplus_{\substack{v' \in A_p(v) \\ p \in P}} m_p(h_{v'}^t) \right) \quad (t = 0, \dots, T - 1), \quad (4)$$

where  $p \in P$  indexes different edge types,  $A_p(v)$  is the set of neighbors of node  $v$  with respect to the  $p$ th edge type. We choose the message function  $m_p$  to be a single dense layer for each edge type  $p$  and the aggregation operator  $\bigoplus$  as the sum of all incoming messages. We use the GRU cell (Cho et al., 2014) as the recurrent unit  $f$ . The final node embedding  $h_v^T$  is computed by unrolling equation 4 for  $T$  time steps.

Finally the value of expression  $e$  is computed by taking a weighted sum of the final node embedding  $h_v^T$  and passing through a dense layer as follows:

$$\mathbf{V}(e) = O \left( \sum_{v \in N(e)} \sigma(g(h_v^T, h_v^0)) \cdot r(h_v^T) \right) \quad (5)$$

where  $g : \mathbb{R}^{2d} \rightarrow \mathbb{R}^d$ ,  $r : \mathbb{R}^d \rightarrow \mathbb{R}^d$ ,  $O : \mathbb{R}^d \rightarrow \mathbb{R}^H$  are all one-layer dense networks,  $\sigma$  denotes the sigmoid function, and  $\mathbf{V}(e) = [V((e, 0)), V((e, 1)), \dots, V((e, H - 1))]^\top \in \mathbb{R}^H$ .

We train the above GNN to approximate the optimal value function  $V^*$ . Let  $\bar{V}(s) = \mathbf{V}(e_s)[t_s]$  the value function  $\mathbf{V}$  computed for expression  $e_s$  and time budget  $t_s$ . To track an approximate lower bound of the optimal value function  $V^*$ , we minimize the loss  $l(\bar{V}(s) - V^{\text{target}}(s)/t_s)$ . The target value  $V^{\text{target}}(s)$  is defined in Algorithm 2 and corresponds to the best cost improvement obtained with the current policy in  $t_s$  steps. Normalization by  $t_s$  is introduced to ease optimisation by ensuring that target values for all outputs of  $\mathbf{V}$  are in a similar magnitude. Thus the value function estimate  $\hat{V}(s)$  can be obtained from per-step value estimate  $\bar{V}(s)$  as  $\hat{V}(s) = t_s \cdot \bar{V}(s)$ . For the loss function  $l$  we use the Huber loss. Details about the optimiser used to minimize the loss  $l$  are given in Appendix B. In the pseudocode in Algorithm 1, this optimisation is represented with the function FIT.

## 4 RELATED WORK

Knossos builds on a long tradition of compiler technology. Similarly to traditional compilers (Santos & Peyton-Jones, 1992; Lattner & Adve, 2004) and the more recent deep learning compilers such as Myia (van Merriënboer et al., 2018), DLVM (Wei et al., 2018), ISAM (Sotoudeh et al., 2019) and GLOW (Rotem et al., 2018), Knossos uses an intermediate representation to optimize programs. However, while these approaches rely on layers of hand-coded optimisation heuristics, Knossos *learns* the algorithm used to optimize its programs.

In this respect, Knossos is a spiritual successor of benchmark-driven hardware-agnostic optimisation approaches in computational linear algebra (Padua, 2011) and signal processing (Frigo & Johnson, 1998). However, unlike these approaches, Knossos is a fully-fledged compiler, and can optimize arbitrary programs. Moreover, thanks to its Reinforcement Learning-driven optimizer, Knossos has an advantage over existing approaches that attempt to learn how to optimize arbitrary code. For example, Bunel et al. (2017) learns parameters of a code optimizer with a hard-coded hierarchy. REGAL (Paliwal et al., 2019) only learns the hyper-parameters for a fixed genetic algorithm that preforms the actual optimisation. The TVM compiler (Chen et al., 2018a) learns a cost model over programs, but uses simple simulated annealing to perform the optimisation. Similarly, Chen et al. (2018b) handles only index summation expressions and again relies on simulated annealing. LIFT (Steuer et al., 2017) defines an intermediate language suited for expressing numerical computation, but focuses on providing the right set of rewrite rules rather than on the program optimisation process itself. In Section 5, we demonstrate that the RL optimizer used by Knossos outperforms this approach by a large margin.

Knossos is also related to JAX (Frostig et al., 2018), which performs just-in-time compilation of Python code using the XLA backend (XLA authors, 2016). Knossos differs from JAX in two ways.




Figure 3: Example of tricky optimisation task. Two expressions are similar but the optimal rewrite strategies are different.

First, it uses efficient RL code optimisation, which is architecture-agnostic. In fact, since Knossos generates C++ code, it supports a much broader variety of target architectures. Also, unlike JAX, it makes use of the benefits of a statically typed languages. In terms of scope, Knossos is also similar to Zygote for Julia (Innes et al., 2019). However, unlike these compilers, Knossos makes use of an RL-driven code optimizer.

Since Knossos provides first class support for automatic differentiation, it is also related to established deep learning frameworks (Maclaurin et al., 2015; Abadi et al., 2016; Paszke et al., 2017). However, unlike Knossos, these frameworks do not learn how to optimize code, instead relying on manually-prepared back-ends. Moreover, using them either requires meta-programming, where the user has to use a high-level language to specify the desired computation graph using constructions external to the language (Abadi et al., 2016), or is constrained to a restricted subset of the language (Paszke et al., 2017). In contrast, the Knossos language can be used directly, without manually specifying computation graph constructs or restricting oneself to an allowed subset of the language.

In parallel, the idea of automated rewriting to achieve a given objective was explored in the context of automated theorem provers. This is conceptually related to our approach since finding an equivalence between formulae is the same as finding a proof that they are equal. However, recent work in this space has substantial differences in scope. In particular, state-of-the-art work that searches for refutational proofs in first-order logic (Zombori et al., 2019; Kaliszyk et al., 2018) uses hard-coded features and cannot learn any new ones. Also, the optimal objective is very different. While a mathematical proof is only correct when completely reduced to a tautology, we are satisfied with simplifying an expression by a certain margin, not necessarily in the most optimal way possible.

For the Reinforcement Learning part, our algorithm differs from standard techniques in that it has a much larger action space and a state space that consists of graphs, which makes the application of traditional RL algorithms like DQN (Mnih et al., 2013), A2C (Mnih et al., 2016) and PPO (Schulman et al., 2017) ineffective. AlphaGo (Silver et al., 2016; 2017), which also performs a search over a large state space, but differs from Knossos in that it learns for pixel observations and uses an action space of bounded size. Reinforcement Learning has also been applied to expression rewriting and scheduling problems (Chen & Tian, 2019). However, since this approach used actor-critic RL that does not exploit reset, it is less well-suited for compilation tasks as described in Section 3.

### Generalisation to Unseen Data

## 5 BENCHMARKS

We evaluated Knossos in three settings. First, to understand how close and reliably we can achieve the best optimisation, we applied Knossos to a manually curated set of arithmetic expressions, where we know the best available sequence of rewrites. Second, we applied Knossos to a set of linear algebraic operations, which are representative of typical workloads in numerical computing. Third,

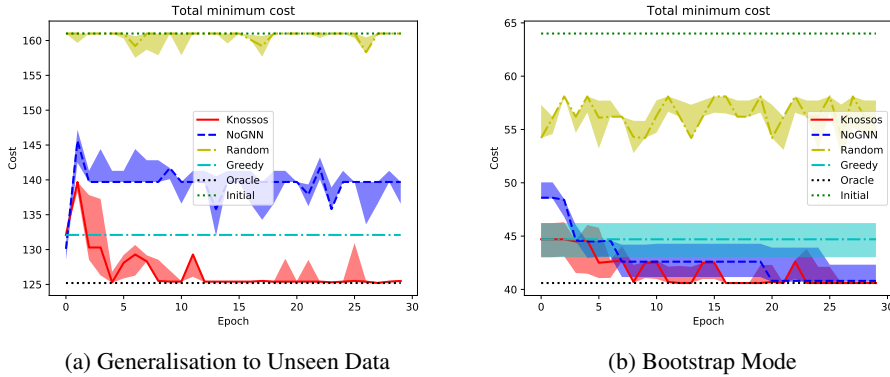


Figure 5: Performance of Knossos on a set of arithmetic expressions. Horizontal axis shows epochs, vertical axis shows cost. Shaded area spans the 20% and 80% percentile over 10 repetitions.

in order to reflect common modern use-cases, we evaluated Knossos on a computer vision task that uses a convolutional neural network. Since Knossos is as an alternative to traditional compilers, we compare it to a hand-written rule-based transpiler of the Knossos IL, which we call *k<sub>SC</sub>*. Both Knossos and *k<sub>SC</sub>* output C++, which is compiled to binary using *gcc with optimisation enabled*, ensuring a fair comparison. We describe the results below.

While arithmetic expressions are simple, optimising them is not always a simple task. Figure 3 shows an example of two similar arithmetic expressions. Although they look very similar, they require different optimisation strategy to reach to the optimal form. The left expression gets to optimal by an arithmetic simplification ( $\times x$  to a denominator and a numerator) but the right expression gets to optimal by a common subexpression elimination. It is difficult for a rule-based compiler to distinguish the two and optimise such similar expressions using different strategies. To test Knossos on arithmetic expressions, we used a training set of 36 arithmetic expressions and a test set of 12 different ones. The details of the experimental setup are given in Appendix B. In this setting, we pick 6 expressions randomly from a training set to train in each epoch. We ran training for 30 epochs and running 10 repetitions for each experiment with different random seeds. Search depth was limited to 10 and the termination condition in  $A^*$  was set to 5000 evaluations of the value function. See Appendix B for the full details including network parameters. We show the results in Figure 5a. It can be seen from the figure that Knossos achieved the oracle cost for all expressions. We also performed an ablation, comparing Knossos to  $A^*$  algorithm (shown as NoGNN) that does not perform the GNN recurrence in equation 4. As a baseline, we compared to greedy best-first search, which picks a next state to explore greedily without using the value function  $f(s) := c(s_0) - c(s)$ . We also show a comparison to random search and the initial cost of the expression, before any optimisation.

**Bootstrap Mode** Similarly to a traditional compiler, where we are given a concrete program to optimize, the expressions used to evaluate Knossos in this benchmark were the same ones that we used during training. Even in this setup,

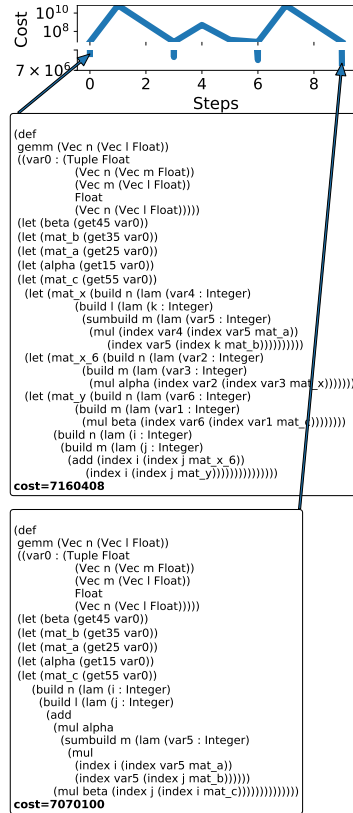


Figure 4: General Matrix Multiply (GEMM) program rewrite sequence obtained by Knossos. The initial expression was obtained from our rule-based *k<sub>SC</sub>* compiler and shown in the middle. The final expression was obtained after 10 rewriting steps and shown in the bottom.






Figure 6: Performance of Knossos on basic linear algebra and convolutional network. Shaded area indicates one standard deviation.




Figure 7: Comparison of wall time. Shaded area indicates one standard deviation.

Knossos still generalises, but it does it across sub-expressions of the expressions in the training set. We tested that on 8 expressions, training for 30 epochs. Other experimental setup is the same as Arithmetic Expressions. Figure 5a shows the comparison of the minimum cost achieved by each agent. It can be seen from the figure that Knossos achieved the best possible cost for all expressions.

**Linear Algebra Primitives** Numerical linear algebra is fundamental to most calculations in scientific computing and machine learning. Primitives such as vector multiplication, plane rotation, matrix multiplications and similar primitives often represent the most time-consuming part of the given computation. To evaluate the performance of Knossos on in this setting, we trained on a set of 11 such linear algebra primitives and evaluated on General Matrix Multiplication (GEMM). We trained for 5 epochs, each of which included optimisation of cost of 6 primitives. Search depth was limited to 30 and the termination condition in  $A^*$  was set to 5000 evaluations of the value function. Figure 6a shows the cost of GEMM. The plot shows results for 10 independent runs of the Knossos code optimizer on the same input source file. We used an augmented set of training rules, which included vector operations (see Table 4 in Appendix). Because of the complexity of the task, we split the search into two phases of 15 steps each. The training phases differ in the set of allowed rules. In the first phase, we only allow rules that result in large changes to the cost (Table 4). In the second phase, we allow all rules. The shaded area represents one standard deviation across the runs of Knossos. Results show that Knossos produced code of lower cost than the output of the traditional `ksc` compiler according to our cost model. We also performed a benchmark using wall clock time, shown in Fig. 7a, again showing an improvement. In addition, we performed a qualitative evaluation of the output in Fig. 4. In the program obtained by `ksc` (middle listing), three temporary variables `mat_x`, `mat_x_6`, and `mat_y` corresponding to the result of  $A \cdot B$ ,  $\alpha \cdot \text{mat}_x$ , and  $\beta \cdot C$ , respectively, are created. In the output of Knossos (bottom listing), all the temporary variables are gone. Hence, Knossos has discovered a form of loop fusion – the type of optimisation that previously had to be built into a compiler by a laborious manual process.

**Convolutional Network** In order to evaluate Knossos on workloads characteristic of modern machine learning pipelines, we also evaluated Knossos on a computer vision task. We optimize a code for training a convolutional deep network on the MNIST dataset (LeCun, 1998). The source code




Figure 8: Reverse mode of convolutional network rewrite sequence obtained by Knossos. The initial expression was obtained from our rule-based `ksc` compiler and shown in the bottom left. The final expression was obtained after 32 rewriting steps and shown in the bottom right. The expression in the bottom middle corresponds to the highest point in the cost sequence.

represents a typical implementation of a deep learning algorithm and contains primitives such as dense layers, convolutional layers, pooling layers, and so on. While MNIST is a basic benchmark, we stress that the goal of Knossos was *code optimisation* as opposed to the computer vision task itself. We trained on 5 expressions and evaluated on a reverse mode of a convolutional layer. We fixed the search depth to 40. The termination condition in  $A^*$  was set to 30000 evaluations of the value function. We used an augmented set of training rules and split the search into two phases of 20 steps each, allowing rules that result in large changes to the cost in the first phase and all rules in the second phase. Results are shown in Figure 6b for the cost model and Figure 7b for the wall clock time. The shaded area represents the standard deviation across the runs of Knossos and the resulting binary. As above, the Knossos optimizer produced code that outperformed the baseline.

**Summary of Benchmarks** We have demonstrated that Knossos is capable of producing code that is faster than the output of a traditional compiler. Moreover, unlike traditional compilers, Knossos does not rely on hand-crafted optimisation passes that are very laborious to implement. Instead, traditional optimisation passes are replaced by atomic rewrite rules that can be combined in many ways. In fact, in our benchmark of linear algebra primitives, Knossos was able to automatically discover *loop fusion*, an optimisation strategy long known to compiler designers. Knossos code in our experiments can perform both training and inference and can be run on any hardware supporting the C++ toolchain, including inexpensive embedded devices.

## 6 CONCLUSIONS

We have introduced Knossos, a new compiler targeting machine learning and numerical computation. Thanks to its automatic code optimisation, Knossos produces binaries that achieve better run-times than a traditional, rule-based compiler. Knossos can deal with complex code generated by automatic differentiation and automatically discover optimisations that previously required careful compiler design. We believe that Knossos will pave the way towards a new generation of future compilers, which will crucially rely on automatically inferring the correct optimisations.

## REFERENCES

- Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek Gordon Murray, Benoit Steiner, Paul A. Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: A System for Large-Scale Machine Learning. In Kimberly Keeton and Timothy Roscoe (eds.), *12th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2016, Savannah, GA, USA, November 2-4, 2016*, pp. 265–283. USENIX Association, 2016.
- Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.
- Cameron B Browne, Edward Powley, Daniel Whitehouse, Simon M Lucas, Peter I Cowling, Philipp Rohlfshagen, Stephen Tavener, Diego Perez, Spyridon Samothrakis, and Simon Colton. A survey of monte carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in games*, 4(1):1–43, 2012.
- Rudy Bunel, Alban Desmaison, M. Pawan Kumar, Philip H. S. Torr, and Pushmeet Kohli. Learning to superoptimize programs. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017.
- Tianqi Chen, Thierry Moreau, Ziheng Jiang, Lianmin Zheng, Eddie Q. Yan, Haichen Shen, Meghan Cowan, Leyuan Wang, Yuwei Hu, Luis Ceze, Carlos Guestrin, and Arvind Krishnamurthy. TVM: An Automated End-to-End Optimizing Compiler for Deep Learning. In Andrea C. Arpaci-Dusseau and Geoff Voelker (eds.), *13th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2018, Carlsbad, CA, USA, October 8-10, 2018*, pp. 578–594. USENIX Association, 2018a.
- Tianqi Chen, Lianmin Zheng, Eddie Q. Yan, Ziheng Jiang, Thierry Moreau, Luis Ceze, Carlos Guestrin, and Arvind Krishnamurthy. Learning to Optimize Tensor Programs. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada*, pp. 3393–3404, 2018b.
- Xinyun Chen and Yuandong Tian. Learning to perform local rewriting for combinatorial optimization. *NeurIPS*, abs/1810.00337, 2019. URL <http://arxiv.org/abs/1810.00337>.
- Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- Conal Elliott. The simple essence of automatic differentiation. *PACMPL*, 2(ICFP):70:1–70:29, 2018. doi: 10.1145/3236765.
- Matteo Frigo and Steven G. Johnson. FFTW: An adaptive software architecture for the FFT. In *Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP '98, Seattle, Washington, USA, May 12-15, 1998*, pp. 1381–1384. IEEE, 1998. ISBN 978-0-7803-4428-0. doi: 10.1109/ICASSP.1998.681704.
- Roy Frostig, Matthew Johnson, and Chris Leary. Compiling machine learning programs via high-level tracing. 2018. URL <http://www.sysml.cc/doc/2018/146.pdf>.
- Peter E Hart, Nils J Nilsson, and Bertram Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968.
- Mike Innes, Alan Edelman, Keno Fischer, Chris Rackauckus, Elliot Saba, Viral B Shah, and Will Tebbutt. Zygote: A differentiable programming system to bridge machine learning and scientific computing. *arXiv preprint arXiv:1907.07587*, 2019.
- Cezary Kaliszyk, Josef Urban, Henryk Michalewski, and Miroslav Olsák. Reinforcement Learning of Theorem Proving. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada*, pp. 8836–8847, 2018.
- Levente Kocsis and Csaba Szepesvári. Bandit based monte-carlo planning. In *European conference on machine learning*, pp. 282–293. Springer, 2006.

- Chris Lattner and Vikram S. Adve. LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation. In *2nd IEEE / ACM International Symposium on Code Generation and Optimization (CGO 2004), 20-24 March 2004, San Jose, CA, USA*, pp. 75–88. IEEE Computer Society, 2004. ISBN 978-0-7695-2102-2. doi: 10.1109/CGO.2004.1281665.
- Yann LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard S. Zemel. Gated graph sequence neural networks. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016. URL <http://arxiv.org/abs/1511.05493>.
- Dougal Maclaurin, David Duvenaud, and Ryan P. Adams. Autograd: Effortless gradients in numpy. In *ICML 2015 AutoML Workshop*, 2015.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pp. 1928–1937, 2016.
- OpenAI. AI and Compute. <https://openai.com/blog/ai-and-compute/>, 2018. Blog post.
- David A. Padua. Automatically Tuned Linear Algebra Software (ATLAS). In *Encyclopedia of Parallel Computing*, pp. 101. Springer, 2011. ISBN 978-0-387-09765-7. doi: 10.1007/978-0-387-09766-4\_2061.
- Aditya Paliwal, Felix Gimeno, Vinod Nair, Yujia Li, Miles Lubin, Pushmeet Kohli, and Oriol Vinyals. REGAL: Transfer Learning For Fast Optimization of Computation Graphs. *CoRR*, abs/1905.02494, 2019.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. 2017.
- Martin L. Puterman. *Markov Decision Processes.: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 2014.
- Nadav Rotem, Jordan Fix, Saleem Abdulrasool, Summer Deng, Roman Dzhabarov, James Hegeman, Roman Levenstein, Bert Maher, Nadathur Satish, Jakob Olesen, Jongsoo Park, Artem Rakhov, and Misha Smelyanskiy. Glow: Graph Lowering Compiler Techniques for Neural Networks. *CoRR*, abs/1805.00907, 2018.
- André L. M. Santos and Simon L. Peyton-Jones. On Program Transformation in the Glasgow Haskell Compiler. In John Launchbury and Patrick M. Sansom (eds.), *Functional Programming, Glasgow 1992, Proceedings of the 1992 Glasgow Workshop on Functional Programming, Ayr, Scotland, UK, 6-8 July 1992*, Workshops in Computing, pp. 240–251. Springer, 1992. ISBN 978-3-540-19820-8.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Vedavyas Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy P. Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016. doi: 10.1038/nature16961.
- David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354, 2017.
- Matthew Sotoudeh, Anand Venkat, Michael J. Anderson, Evangelos Georganas, Alexander Heinecke, and Jason Knight. ISA mapper: A compute and hardware agnostic deep learning compiler. In Francesca Palumbo, Michela Becchi, Martin Schulz, and Kento Sato (eds.), *Proceedings of the 16th ACM International Conference on Computing Frontiers, CF 2019, Alghero, Italy, April 30 - May 2, 2019*, pp. 164–173. ACM, 2019. ISBN 978-1-4503-6685-4. doi: 10.1145/3310273.3321559.
- Michel Steuwer, Toomas Rimmelg, and Christophe Dubach. Lift: a functional data-parallel IR for high-performance GPU code generation. In *Proceedings of the 2017 International Symposium on Code Generation and Optimization, CGO 2017, Austin, TX, USA, February 4-8, 2017*, pp. 74–85, 2017. URL <http://dl.acm.org/citation.cfm?id=3049841>.

Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, second edition, 2018.

Alan Mathison Turing. On computable numbers, with an application to the entscheidungsproblem. *Proceedings of the London mathematical society*, 2(1):230–265, 1937.

Bart van Merriënboer, Olivier Breuleux, Arnaud Bergeron, and Pascal Lamblin. Automatic differentiation in ML: Where we are and where we should be going. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada*, pp. 8771–8781, 2018.

Richard Wei, Lane Schwartz, and Vikram S. Adve. DLVM: A modern compiler infrastructure for deep learning systems. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Workshop Track Proceedings*. OpenReview.net, 2018.

XLA authors. TensorFlow XLA (Accelerated Linear Algebra). <https://www.tensorflow.org/xla/overview>, 2016.

Zsolt Zombori, Adrián Csiszárík, Henryk Michalewski, Cezary Kaliszyk, and Josef Urban. Towards Finding Longer Proofs. *CoRR*, abs/1905.13100, 2019.

Table 1: List of Experimental Settings

Arithmetic Expressions	
Parameter	Value
Number of training expressions	36
Number of test expressions	12
Rewrite rules	Simple rule set
Number of evaluation budget per epoch	5000
Minimum batch size for evaluation	16
Maximum Search Depth	10
Basic Linear Algebra	
Number of training expressions	11
Number of test expressions	1
Rewrite rules	Extended rule set
Number of evaluation budget per epoch	30000
Minimum batch size for evaluation	16
Maximum Search Depth	30
Convolutional Neural Network	
Number of training expressions	5
Number of test expressions	1
Rewrite rules	Extended rule set
Number of evaluation budget per epoch	30000
Minimum batch size for evaluation	16
Maximum Search Depth	40

## APPENDICES

### APPENDIX A KNOSSOS INTERMEDIATE REPRESENTATION

For background, we give a brief overview of the intermediate representation (IR) used by the Knossos compiler. The Knossos IR provides a convenient symbolic form for the reinforcement learning optimizer. It also has a LISP-like surface syntax, which we used to implement our programs. In the future, we plan to provide transpilers, allowing for the compilation of code written in other languages into Knossos. We provide a sample Knossos program in Figure 4. In order to facilitate Machine Learning workloads, the Knossos IL has native support for automatic differentiation. We use a new unified view of automatic differentiation as generalised transposition (Elliott, 2018). Rather than having an explicit distinction between forward mode and reverse mode AD, Knossos uses a type system together with a set of consistent rewrite rules. Whenever the gradient operator is used as part of a Knossos algorithm, the compiler first generates a syntax tree corresponding to the differentiated program and then applies rewrites to optimize the cost of its execution. This means that the resulting AD algorithm is tailor-made and optimized with that exact use case in mind. This is in contrast to systems such as PyTorch, which have hard-coded routines for backward-mode AD. From the perspective of the user, this process is completely transparent in the sense that taking gradients can be applied to any piece of Knossos code.

While the details of this process are beyond the scope of this paper, from the perspective of this work, the important feature of AD is that it corresponds to a transformation of the abstract syntax tree. The resulting AST can then be optimised in the same way as any other code.

### APPENDIX B REPRODUCIBILITY AND DETAILS OF EXPERIMENTAL SETUP

We now describe the parameters used to perform the experiments reported on in the paper. The parameters used by  $A^*$  in the four tasks described in Sec. 5 are listed in Tab. 1. The hyper-parameters for the value network training are given in Tab. 2.

Table 2: List of hyperparameters

Model Parameter	Value Function
Number of features per node	200
Number of propagation steps	10
Activation function for GNN	tanh
Dropout rate for GNN	0
Number of hidden layers in MLP	2
Number of dimensions of hidden layers in MLP	200
Activation function for MLP	tanh
Dropout rate for MLP	0.2
<b>Training Parameter</b>	
Loss function	Huber loss
Optimizer	Adam optimizer
Learning rate	0.0001
Batch size	16

In the Graph Neural Network, initial node features are one-hot vectors that represent the node types. The used node types are: constant, variable, let, if, tuple, select, +, -, \*, /, exp, log, ==, >, >=, or, build, apply, lam, sum, sumbuild, constVec, and deltaVec. Edge types are listed in Tab. 2b. The auxiliary edge type “is-identical” is inserted to identify identical subexpressions. It was added so that it is easier to learn re-writes that rely on matching expressions. The GNN was implemented using a sparse adjacency matrix instead of dense matrix in order to conserve GPU memory in settings where some expressions grow beyond > 10000 nodes during training. We ran the GNN recursion 10 times. For optimization we used the Adam optimizer with learning rate 0.0001 and set the dropout rate zero for GNN and 0.2 for the MLP.

#### APPENDIX C REWRITE RULES

We list the basic rule set used in the arithmetic expressions benchmark in Tab. 3. The additional rewrite rules used in basic linear algebra and convolutional neural network are given in Tab. 4.

#### APPENDIX D ADDITIONAL ABLATIONS.

In addition to  $A^*$  search, we compare the performance of Monte Carlo Tree Search (Browne et al., 2012) using the UCT formula (Auer et al., 2002; Kocsis & Szepesvári, 2006). In order to disambiguate across subtly different versions of the algorithm, we describe it below.

Each iteration of MCTS consists of four steps (Algorithm 3).

1. Selection: Starting from the root, a tree policy is recursively descends through the tree until it reaches a leaf node.
2. Expansion: A child node is added to expand the tree.
3. Simulation: A rollout policy is applied from the new node until the end of the episode.
4. Back-up: The simulation result is backed up through the selected nodes to update their statistic.

The tree policy  $\pi_t$  and rollout policy  $\pi_r$  are defined as follows.

$$\pi_t(a|s) = \arg \max_{a \in A} \left( \frac{X(s')}{n(s')} + \beta \sqrt{\frac{\ln n(s)}{n(s')+1}} \right) \quad (6)$$

$$\pi_r(a|s) = \text{softmax}_{a \in A} (R(s, a) + V(s'), \alpha) \quad (7)$$

Here,  $n(s)$  is a visitation count of state  $s$ , and  $\beta$  is a constant to control the exploration bonus.  $X(s')/n(s')$  is the average cost reduction achieved by a set of trajectories which passed through

Table 3: List of rewrite rules used for arithmetic expressions.

LHS	RHS	Side conditions
<b>Simple Rule Set</b>		
<b>Arithmetic Rules</b>		
(add 0.0 a)	a	
(mul 1.0 a)	a	
(sub a 0.0)	a	
(div a 1.0)	a	
(add a (mul b -1.0))	(sub a b)	
(mul a b)	(mul b a)	
(add a b)	(add b a)	
(sub a a)	0.0	
(add (add a b) c)	(add (add a c) b)	
(sub (add a b) c)	(add (sub a c) b)	
(add (sub a b) c)	(sub (add a c) b)	
(sub (sub a b) c)	(sub (sub a c) b)	
(sub a (sub b c))	(sub (add a c) b)	
(mul (mul a b) c)	(mul (mul a c) b)	
(div (mul a b) c)	(mul (div a c) b)	
(mul (div a b) c)	(div (mul a c) b)	
(div (div a b) c)	(div (div a c) b)	
(div a (div b c))	(div (mul a c) b)	
(div (div a b) c)	(div a (mul b c))	
(mul a (add b c))	(add (mul a b) (mul a c))	
(add (mul a b) (mul a c))	(mul a (add b c))	
(mul a (sub b c))	(sub (mul a b) (mul a c))	
(add a (mul -1.0 b))	(sub a b)	
(sub (mul a b) (mul a c))	(mul a (sub b c))	
(add (div y x) (div z x))	(div (add y z) x)	
(sub (div y x) (div z x))	(div (sub y z) x)	
(div a (mul b c))	(div (div a b) c)	
(mul 0.0 a)	0.0	
<b>Binding Rules</b>		
e	(let (x e) x)	
(op (let (x e) x))	(let (x e) (op x))	$x \notin \text{op}$
(let (x e1) (let (x e1) e2))	(let (x e1) e2)	
(let (x e1) e2)	(let (x e1) e2[x/e1])	replace all free x in e2 with e1
(let (x e1) e2)	e2	$x \notin e2$

Table 4: List of rewrite rules used in addition to simple rule set (Table 3) for Basic Linear Algebra and Convolutional Neural Network benchmarks. In the two-phase strategy described in Sec. 5, only the rules with checkmarks in the last column are used for the first 15 epochs.

LHS	RHS	Side conditions	Include in phase 1
(size (build n (lam x b)))	n		
(index arg (build sz (lam x body)))	(let (x arg) body)		✓
(select i (tuple e1 e2 e3...))	ei		
(sum (build n (lam x b)))	(sumbuild n (lam x b))		✓
(sumbuild n (lam x (deltaVec m i v)))	(deltaVec m i (sumbuild n (lam x v)))	$x \notin m \cap x \notin i$	✓
(sumbuild n (lam x (tuple f g)))	(tuple (sumbuild n (lam x f)) (sumbuild n (lam x g)))		✓
(sumbuild n (lam x (if c t f)))	(if c (sumbuild n (lam x t)) (sumbuild n (lam x f)))		✓
(sumbuild n (lam x c))	(mul n c)	$x \notin c \cap c$ is a constant	✓
(sumbuild o (lam oi (add e1 e2)))	(add (sumbuild o (lam oi e1)) (sumbuild o (lam oi e2)))		✓
(sumbuild o (lam oi (build n (lam ni e))))	(build n (lam ni (sumbuild o (lam oi e))))		✓
(add (deltaVec o oi e1) (deltaVec o oi e2))	(deltaVec o oi (add e1 e2))		✓
(add (tuple x y) (tuple a b))	(tuple (add x a) (add y b))		✓
(mul (tuple x y z)	(tuple (mul x z) (mul y z))		✓
(if True a b)	a		
(if False a b)	b		
(if p True False)	p		
(if p a a)	a		
(eq a a)	True		



**Algorithm 3** Monte Carlo Tree Search (MCTS)

---

```

function MCTS( $s_0, V, d, \alpha, c$ )
   $\pi_t(a|s) \leftarrow \arg \max_{a \in A} \left( X(s')/n(s') + \beta \sqrt{\frac{\ln n(s)}{n(s')+1}} \right)$ 
   $\pi_r(a|s) \leftarrow \text{softmax}_{a \in A}(R(s, a) + V(s'), \alpha)$ 
   $M \leftarrow \emptyset$ 
   $\forall s \in S : n(s) \leftarrow 0$  ▷ Initialize a visitation count
  while not term_condition() do
     $T_t \leftarrow \text{TREESEARCH}(s_0, d, \pi_t)$ 
     $s_l \leftarrow \text{tail}(T_t)$ 
     $T_r \leftarrow \text{ROLLOUTSEARCH}(s_l, d - \text{length}(T_t), \pi_r)$ 
     $T \leftarrow \text{concat}([s_0], T_t, T_r)$ 
     $x \leftarrow c(s_0) - \min_{s \in T_r} c(s)$ 
    ▷  $x$  is the max cost reduction found by rollout
    for all  $s \in T_t$  do ▷ MCTS backup
       $X(s) \leftarrow X(s) + x$ 
       $n(s) \leftarrow n(s) + 1$ 
    end for
     $M \leftarrow M \cup \{T\}$ 
  end while
  return  $M$ 
end function

function TREESEARCH( $s, d, \pi$ )
   $T \leftarrow []$ 
  while  $\text{length}(T) < d$  and  $n(s) > 0$  do
     $a \leftarrow \pi(a|s)$  ▷ Tree policy  $\pi$  is deterministic.
     $s \leftarrow a \circ s$ 
     $T \leftarrow \text{concat}(T, [s])$ 
  end while
  return  $T$ 
end function

```

---

$s'$  in  $M$ :  $X(s') = \sum_{T \in H: s' \in T} \max_{s_d \in T: d(s_d) \geq d(s')} c(s_0) - c(s_d)$ , and  $M$  is a set of trajectories sampled at current epoch for the current expression so far. The more the state  $s'$  is visited, the exploration bonus  $(\beta \sqrt{\frac{\ln n(s)}{n(s')+1}})$  is reduced. This way, the agent is encouraged to try a diverse set of actions.

We evaluated the performance of  $A^*$  search and MCTS for both training and test. The experimental setup is the same as the *Generalisation to Unseen Data* experiment in Section 5 except for the used search algorithm. For MCTS, we used  $\alpha = 5.0$  and  $\beta = 0.5$  for both training and test.

Figure 9a shows the results of running all possible combinations of search algorithms when used for training and test in various configurations. Overall, using  $A^*$  for both training and test achieved the best performance. In particular, when we fixed the algorithm used during test to  $A^*$  and varied the training algorithm between  $A^*$  and MCTS,  $A^*$  achieved a significantly lower the total minimum cost than MCTS. Similarly, when we fixed the algorithm used for training to  $A^*$  and compared the performance during testing  $A^*$  achieved significantly lower cost than MCTS again.

## APPENDIX E LIST OF EXPRESSIONS

Table 5: List of expressions for Generalization to Unseen Expressions.

Train expression set
(mul (div (mul 2.0 (add x y)) (sub (add x y) 1)) (mul x y))
(add (add (mul x 2.0) 1) (div (add (div x (sub y 1)) (mul x 2.0)) 2))
(div (add (mul 2.0 (mul x y)) (mul x y)) (div 1.0 (add x 1)))
(div (add (add (div x (sub y 1)) (log x)) 1) (sub (mul (div x (sub y 1)) 3) 1))
(add (div (div 1.0 (add x 1)) (add (add x y) (div 1.0 (add x 1)))) (div 1.0 (add x 1)))
(mul (add (log x) (mul x 2.0)) (mul (log x) 3))
(mul (add (add x y) (log x)) (mul (add x y) 3))
(div (add (add (mul x 2.0) 1) 1) (sub (add (add x y) (mul x 2.0)) 1))
(sub (mul (mul 2.0 (mul x y)) (mul x y)) (mul x 2.0))
(mul (add (div x (sub y 1)) 1) (add (mul x y) (div x (sub y 1))))
(mul (div (div 1.0 (add x 1)) (add (div x (sub y 1)) (div 1.0 (add x 1)))) (div 1.0 (add x 1)))
(add (add (log x) (div 1.0 (add x 1))) (div (mul (log x) 3) 2))
(add (add (add x y) (div x (sub y 1))) (div (mul (add x y) 3) 2))
(div (mul (mul x 2.0) (mul x 2.0)) (sub (mul x 2.0) (div 1.0 (add x 1))))
(mul (div (mul 2.0 (mul x y)) (sub (mul x y) 1)) (log x))
(add (add (div x (sub y 1)) 1) (div (add (add x y) (div x (sub y 1))) 2))
(div (add (mul 2.0 (div 1.0 (add x 1))) (div 1.0 (add x 1))) (mul x 2.0))
(div (add (add (log x) (mul x y)) 1) (sub (mul (log x) 3) 1))
(div (add (add (add x y) (mul x 2.0)) 1) (sub (mul (add x y) 3) 1))
(add (div (mul x 2.0) (add (mul x y) (mul x 2.0))) (mul x 2.0))
(mul (add (mul x y) (div x (sub y 1))) (mul (mul x y) 3))
(div (add (add (div x (sub y 1)) 1) 1) (sub (add (div 1.0 (add x 1)) (div x (sub y 1))) 1))
(sub (mul (mul 2.0 (div 1.0 (add x 1))) (div 1.0 (add x 1))) (log x))
(mul (add (log x) 1) (add (add x y) (log x)))
(mul (add (add x y) 1) (add (div 1.0 (add x 1)) (add x y)))
(mul (div (mul x 2.0) (add (log x) (mul x 2.0))) (mul x 2.0))
(add (add (mul x y) (add x y)) (div (mul (mul x y) 3) 2))
(div (mul (div x (sub y 1)) (div x (sub y 1))) (sub (div x (sub y 1)) (mul x 2.0)))
(mul (div (mul 2.0 (div 1.0 (add x 1))) (sub (div 1.0 (add x 1)) 1)) (mul x y))
(add (add (log x) 1) (div (add (div x (sub y 1)) (log x)) 2))
(add (add (add x y) 1) (div (add (mul x y) (add x y)) 2))
(div (add (mul 2.0 (mul x 2.0)) (mul x 2.0)) (div x (sub y 1)))
(div (add (add (mul x y) (div 1.0 (add x 1))) 1) (sub (mul (mul x y) 3) 1))
(add (div (div x (sub y 1)) (add (log x) (div x (sub y 1)))) (div x (sub y 1)))
(mul (add (div 1.0 (add x 1)) (add x y)) (mul (div 1.0 (add x 1)) 3))
(div (add (add (log x) 1) 1) (sub (add (mul x 2.0) (log x)) 1))
Test expression set
(div (add (mul 2.0 (add x y)) (add x y)) (mul x 2.0))
(div (add (add (mul x 2.0) (mul x y)) 1) (sub (mul (mul x 2.0) 3) 1))
(add (div (mul x y) (add (div x (sub y 1)) (mul x y))) (mul x y))
(mul (add (div x (sub y 1)) (div 1.0 (add x 1))) (mul (div x (sub y 1)) 3))
(div (add (add (div 1.0 (add x 1)) 1) 1) (sub (add (log x) (div 1.0 (add x 1))) 1))
(sub (mul (mul 2.0 (log x)) (log x)) (add x y))
(sub (mul (mul 2.0 (add x y)) (add x y)) (div 1.0 (add x 1)))
(mul (add (mul x 2.0) 1) (add (log x) (mul x 2.0)))
(mul (div (mul x y) (add (add x y) (mul x y))) (mul x y))
(add (add (div x (sub y 1)) (mul x 2.0)) (div (mul (div x (sub y 1)) 3) 2))
(div (mul (div 1.0 (add x 1)) (div 1.0 (add x 1))) (sub (div 1.0 (add x 1)) (mul x y)))
(mul (div (mul 2.0 (log x)) (sub (log x) 1)) (div x (sub y 1)))




Figure 9: Comparison of  $A^*$  search and Monte Carlo Tree Search.

Table 6: List of expressions for Generalization scenario.

Train/Test expression set
(div (div 1.0 x) (add 1.0 (div 1.0 x)))
(add (div (div 1.0 x) (add (div 1.0 x) 1.0)) (div 1.0 x))
(add (div (div 1.0 x) (add (div 1.0 x) 2.0)) (div 1.0 x))
(mul (div x y) (div x y))
(div (mul (div x y) x) y)
(add (div (mul x y) (add 1.0 (mul x y))) (mul x y))
(add (div 1.0 (add 1.0 (mul x y))) (mul x y))
(div (mul x y) (add 1.0 (mul x y)))

Figure 10: List of expressions in training set for Linear Algebra Primitives.

```

(def
  float_matrix_multiply (Vec n (Vec m Float))
  ((alpha : Float) (mat_x : Vec n (Vec m Float)))
  (build n
    (lam (i : Integer) (build m
      (lam (j : Integer) (mul alpha (index i
        (index j mat_x))))))))))

(def
  matrix_matrix_multiply (Vec n (Vec 1 Float))
  ((mat_x : Vec n (Vec m Float)) (mat_y : Vec m (Vec 1 Float)))
  (build n
    (lam (i : Integer) (build 1
      (lam (k : Integer) (sumbuild m
        (lam (j : Integer) (mul (index i
          (index j
            mat_x))
            (index j
              mat_y)))))))))))

(def
  vector_vector_add (Vec n Float)
  ((vec_x : Vec n Float) (vec_y : Vec n Float))
  (build n (lam (i : Integer) (add (index i vec_x) (index i vec_y)))))

(def
  matrix_matrix_add (Vec n (Vec m Float))
  ((mat_x : Vec n (Vec m Float)) (mat_y : Vec n (Vec m Float)))
  (build n
    (lam (i : Integer) (build m
      (lam (j : Integer) (add (index i (index
        j mat_x))
          (index i (index j
            mat_y))))))))))

(def
  outer_product (Vec n (Vec m Float))
  ((vec_x : Vec n Float) (vec_y : Vec m Float))
  (build n
    (lam (i : Integer) (build m
      (lam (j : Integer) (mul (index i vec_x)
        (index j vec_y)))))))

(def
  transpose (Vec m (Vec n Float))
  ((mat_x : Vec n (Vec m Float)))
  (build m
    (lam (i : Integer) (build n
      (lam (j : Integer) (index j (index i
        mat_x)))))))

(def
  scal (Vec n Float)
  ((a : Float) (vec_x : Vec n Float))
  (build n (lam (i : Integer) (mul a (index i vec_x)))))

(def
  axpy (Vec n Float)
  ((alpha : Float) (vec_x : Vec n Float) (vec_y : Vec n Float))
  (build n
    (lam (i : Integer) (add (index i (scal alpha vec_x))
      (index i vec_y)))))

(def
  dot Float
  ((vec_x : Vec n Float) (vec_y : Vec n Float))
  (sumbuild n
    (lam (i : Integer) (mul (index i vec_x) (index i vec_y)))))

(def
  rscl (Vec n Float)
  ((vec_x : Vec n Float))
  (build n (lam (i : Integer) (div (index i vec_x) (sum vec_x)))))

(def
  rot (Tuple (Vec n Float) (Vec n Float))
  ((x : Vec n Float) (y : Vec n Float) (c : Float) (s : Float))
  (tuple (build n
    (lam (i : Integer) (add (mul c (index i x)) (mul s (index
      i y))))))
    (build n
      (lam (i : Integer) (add (mul (mul -1.0 s) (index i x))
        (mul c (index i y)))))))

```

Figure 11: Test expression for Linear Algebra Primitives (General Matrix Multiplication).

```

(def
  gemm (Vec n (Vec l Float))
  ((var0 : (Tuple Float
            (Vec n (Vec m Float))
            (Vec m (Vec l Float))
            Float
            (Vec n (Vec l Float))))))
  (let (beta (get$4$5 var0))
    (let (mat_b (get$3$5 var0))
      (let (mat_a (get$2$5 var0))
        (let (alpha (get$1$5 var0))
          (let (mat_c (get$5$5 var0))
            (let (mat_x (build n (lam (var4 : Integer)
                                     (build l (lam (k : Integer)
                                                  (sumbuild m (lam (var5 : Integer)
                                                                (mul (index var4 (index var5 mat_a))
                                                                    (index var5 (index k mat_b))))))))))
              (let (mat_x_6 (build n (lam (var2 : Integer)
                                          (build m (lam (var3 : Integer)
                                                         (mul alpha (index var2 (index var3 mat_x)))))))
                  (let (mat_y (build n (lam (var6 : Integer)
                                           (build m (lam (var1 : Integer)
                                                         (mul beta (index var6 (index var1 mat_c)))))))
                      (build n (lam (i : Integer)
                                     (build m (lam (j : Integer)
                                                    (add (index i (index j mat_x_6))
                                                         (index i (index j mat_y))))))))))))))))))

```

Figure 12: List of expressions in train set for Convolutional Network.

```

(def
  rev$matrix_multiply (Tuple (Vec o (Vec k Float)) (Vec k Float))
  ((w : Vec o (Vec k Float))
   (image : Vec k Float)
   (d$r : Vec o Float))
  (sumbuild o (lam (oi : Integer)
    (let ((a_6 (index oi d$r))
          (a_5 (build k (lam (sum$i : Integer) a_6))))
      (tuple (sumbuild k (lam (ki : Integer)
        (deltaVec o oi
          (deltaVec k ki
            (mul (index ki image) (index ki a_5))))))
        (build k (lam (ki : Integer)
          (mul (index ki (index oi w)) (index ki a_5))))))))))
  (def
    rev$matrix_multiply_transpose (Tuple (Vec k (Vec o Float))
                                          (Vec k Float))
    ((w : Vec k (Vec o Float))
     (image : Vec k Float)
     (d$r : Vec o Float))
    (sumbuild o (lam (oi : Integer)
      (let ((a_6 (index oi d$r))
            (a_5 (build k (lam (sum$i : Integer) a_6))))
          (tuple (build k (lam (ki : Integer)
            (deltaVec o oi
              (mul (index ki image) (index ki a_5))))
              (build k (lam (ki : Integer)
                (mul (index oi (index ki w)) (index ki a_5))))))))))
    (def
      rev$max_ (Tuple Float Float)
      ((x1 : Float) (x2 : Float) (d$r : Float))
      (if (gt x1 x2) (tuple d$r 0.0) (tuple 0.0 d$r)))
    (def
      rev$maxpool (Vec n Float)
      ((image : Vec n Float) (d$r : Vec m Float))
      (sumbuild (div n 2) (lam (ni : Integer)
        (let ((t1698 (mul 2 ni))
              (t1702 (add 1 t1698))
              (t1705 (rev$max_ (index t1698 image)
                               (index t1702 image)
                               (index ni d$r))))
          (add (deltaVec n t1698 (select t1705 0))
              (deltaVec n t1702 (select t1705 1))))))
      (def
        rev$expensive (Tuple Float Float)
        ((x1 : Float) (x2 : Float) (d$r : Float))
        (tuple 0.0 0.0))
      (def
        rev$expensivepool (Vec n Float)
        ((image : Vec n Float) (d$r : Vec m Float))
        (sumbuild (div n 2) (lam (ni : Integer)
          (let ((t1720 (mul 2 ni))
                (t1724 (add 1 t1720))
                (t1727 (rev$expensive (index t1720 image)
                                       (index t1724 image)
                                       (index ni d$r))))
              (add (deltaVec n t1720 (select t1727 0))
                  (deltaVec n t1724 (select t1727 1))))))

```

Figure 13: Test expression for Convolutional Network.

```

(def
  rev$convld (Tuple (Vec k (Vec l (Vec kn Float)))
                    (Vec l (Vec n Float)))
  ((var0 : (Tuple (Vec k (Vec l (Vec kn Float)))
                  (Vec l (Vec n Float)))
   (Vec k (Vec n Float))))
  (let
    ((kernels (get$1$3 var0))
     (image (get$2$3 var0))
     (d$r (get$3$3 var0)))
    (sumbuild k (lam (ki : Integer)
                     (let (a_6 (index ki d$r))
                         (sumbuild n (lam (ni : Integer)
                                           (let (a_8 (index ni a_6))
                                               (let (a_7 (build kn (lam (var1 : Integer) a_8)))
                                                   (sumbuild kn (lam (kni : Integer)
                                                                     (let (a_10 (index kni a_7))
                                                                         (let (a_11 (build l (lam (sum$i : Integer) a_10)))
                                                                             (sumbuild l (lam (li : Integer)
                                                                                               (let (noi (sub (add ni (div kn 2)) kni))
                                                                                                   (let (outside_image (or (gt 0 noi) (gte noi n)))
                                                                                                       (add (if outside_image
                                                                                                             (tuple (constVec k (constVec l (
                                                                                                               constVec kn 0.0)))
                                                                                                             (constVec l (constVec n 0.0)))
                                                                                                       (tuple (constVec k (constVec l (
                                                                                                               constVec kn 0.0)))
                                                                                                       (deltaVec l li
                                                                                                       (deltaVec n noi
                                                                                                       (mul (index kni (index li (
                                                                                                         index ki kernels)))
                                                                                                       (index li a_11))))))
                                                                                                   (tuple (deltaVec k ki
                                                                                                       (deltaVec l li
                                                                                                       (deltaVec kn kni
                                                                                                       (mul (if outside_image
                                                                                                             0.0
                                                                                                       (index noi (index li
                                                                                                         image)))
                                                                                                       (index li a_11))))))
                                                                                                   (constVec l (constVec n 0.0))))))))
                                                                                                       ))))))))
    ))))))))

```