



HAL
open science

CATREEN : Context-Aware Code Timing Estimation with Stacked Recurrent Networks

Abderaouf Amalou, Elisa Fromont, Isabelle Puaut

► **To cite this version:**

Abderaouf Amalou, Elisa Fromont, Isabelle Puaut. CATREEN : Context-Aware Code Timing Estimation with Stacked Recurrent Networks. ICTAI 2022 - 34th IEEE International Conference on Tools with Artificial Intelligence, Oct 2022, Virtually, China. pp.1-6, 10.1109/ICTAI56018.2022.00090 . hal-03890057v1

HAL Id: hal-03890057

<https://hal.science/hal-03890057v1>

Submitted on 8 Dec 2022 (v1), last revised 3 Jul 2023 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

CATREEN : Context-Aware Code Timing Estimation with Stacked Recurrent Networks

Abderaouf N Amalou

Univ. Rennes, INRIA, CNRS, IRISA

Rennes, France

abderaouf.amalou@irisa.fr

Elisa Fromont

Univ. Rennes, IUF, INRIA, CNRS, IRISA

Rennes, France

elisa.fromont@irisa.fr

Isabelle Puaut

Univ. Rennes, INRIA, CNRS, IRISA

Rennes, France

isabelle.puaut@irisa.fr

Abstract—Automatic prediction of the execution time of programs for a given architecture is crucial, both for performance analysis in general and for compiler designers in particular. In this paper, we present CATREEN, a recurrent neural network able to predict the steady-state execution time of each basic block in a program. Contrarily to other models, CATREEN can take into account the *execution context* formed by the previously executed basic blocks which allows accounting for the processor micro-architecture without explicit modeling of micro-architectural elements (caches, pipelines, branch predictors, etc.). The evaluations conducted with synthetic programs and real ones (programs from Mibench and Polybench) show that CATREEN can provide accurate prediction for execution time with 11.4% and 16.5% error on average, respectively and that we got an improvement of 18% and 27.6% respectively when comparing our tool estimations to the state-of-the-art LSTM-based model.

I. INTRODUCTION

The complexity of developing cycle-accurate simulators and integrating them into compiler infrastructures has led compiler designers to use much simpler ways to estimate execution times. Such estimation is useful to decide about the compiler optimizations to apply, the simplest ones being architecture-independent cost functions, e.g., simply counting the number of machine instructions. An alternative, reachable thanks to the recent advances in Machine Learning (ML), consists in predicting the execution times of code snippets to guide optimizations [1]. The benefit of using ML for timing prediction is threefold: (i) no detail of the processor microarchitecture is needed, because the behavior is learned from timing measurements; (ii) porting the ML-based execution timing predictor to a new microarchitecture does not need any deep expertise, only a new training step is required; (iii) even if training an ML model is a time-consuming task, prediction is in general fast, allowing ML-based timing predictions to be used in compilers.

Today’s processors feature increasing complexity (e.g., cache hierarchy, pipelines, branch predictors, instruction scheduling in out-of-order cores), which, combined with the lack of associated documentation, makes building a cycle-accurate simulator like [2] for each new architecture time-consuming and error-prone. This complexity is mainly due to the integration of several hardware accelerators, which aim to speed up the execution time of programs. The cohabitation of these components makes the timing modeling of a processor difficult to achieve. Moreover,

these components introduce dependencies between successive instructions, making the timing of a sequence of instructions dependent on its *execution context* (history of instructions executed before the sequence under study).

In this paper, we introduce CATREEN, for Context-Aware code Timing estimation with stacked REcurrEnt Networks. CATREEN infers the steady-state execution times of individual basic blocks (BBs) in programs (at the assembly code level). As compared to related works on estimating the execution time of BBs, the novelty of our work is to assess the execution time of a BB within its *execution context*. Instead of calculating the best-case execution time in isolation of a basic block b ; like it was done in ITHEMAL [3], CATREEN calculates the execution time of b after executing a sequence of other BBs. Experiments on an embedded architecture (STM32H6 board, featuring an ARM Cortex M7 processor) demonstrate that CATREEN produces more accurate predictions than state-of-the-art context-agnostic ML-based techniques, on both synthetic programs and real codes.

The remainder of this paper is organized as follows. Section II surveys related works that use ML-based techniques for performance evaluation and optimization. An overview of CATREEN is given in Section III. The performance of CATREEN is evaluated in Section IV. Finally, Section V concludes with a summary of the contributions and directions for future work.

II. RELATED WORK

The complexity of evaluating and improving the performance of programs has motivated the use of machine learning techniques. These techniques can be classified into two categories: those which directly evaluate the execution time of programs (§ II-A) and those which target other metrics that are related to performance like speedup or energy consumption, albeit not directly evaluating execution times (§ II-B).

A. Machine learning for execution time prediction

ITHEMAL [3] leverages a hierarchical multi-scale Recurrent Neural Network (RNN) and, in particular, Long Short-Term Memory (LSTM) layers to predict the throughput of basic blocks. ITHEMAL captures the interactions between instructions from the same basic block. Each basic block is isolated

from the program and executed repetitively to reach its steady-state behavior and obtain its throughput (peak performance or *best-case* execution time). In contrast to ITHMAL, CATREEN relaxes the assumption that a basic block is executed in isolation and predicts the execution time of a basic block given its actual context of execution at steady-state.

[4] advocates the use of sparse polynomial regression to predict the execution time of programs given a set of static features. CATREEN differs from [4] by the features used during learning that allow capturing the execution context of basic blocks. It also differs by the type of method used to predict the throughput (in our case, a recurrent neural network).

Ritter and Hack [5] present a framework for inferring port usage of instructions based on an evolutionary algorithm that solves a linear program to predict the throughput of a basic block. The solution was designed to infer port mappings for an out-of-order processor but it is not clear how the approach could be extended to infer the throughput for a more complex trace execution, e.g., how data dependencies could be included in the analysis. In comparison, CATREEN was designed for in-order processor architecture due to the sequential processing of instructions, but it has the ability to learn instruction and data dependencies.

B. Machine learning for performance optimization

ML techniques were proposed in the past to improve the performance of programs [6]–[8].

The DeepTune [6] tool leverages an RNN LSTM-based model that can be used as a classifier to choose the best choice between running a program on the CPU or the GPU, or to predict the best number of threads among a set of values. In [7], a neural network (NN) model uses performance counters to predict the performance (instructions per cycle) of a thread migration from one core to another in S-NUCA architecture. Tousi and Lujan [8] study the feasibility of using classic machine learning techniques to predict performances (latency, speedup...) on the SPEC Benchmarks, while CATREEN only uses programs that can run without an operating system (bare-metal environment to avoid OS noises) which is not the case of SPEC Benchmarks.

If the methods presented before and their inputs are (somewhat) similar to ours, our end goal is entirely different. CATREEN, in contrast to the works mentioned above, focuses on the prediction of execution time for a given hardware and compiler settings.

III. EXECUTION TIME PREDICTION USING CATREEN

CATREEN leverages a stacked recurrent neural network architecture to predict the execution time of a basic block b (or the *basic block under analysis*) given its execution context (sequence of basic blocks executed before b). By considering the execution history of basic blocks, CATREEN naturally accounts for the state of the hardware when executing the basic block (pipeline, cache hierarchy, and branch prediction).

A. Overview of CATREEN

Recurrent Neural Networks (RNN) are particular neural network architectures that can be trained to predict outputs from a variable-length sequence of data. These networks can memorize (parts of) the sequence and the computed sequential information along time. Long Short-Term Memory (LSTM) are special kinds of RNN, capable of learning long-term dependencies [9] in the sequences. LSTM architectures have mechanisms, called gates, that are trained to choose whether or not to keep particular sequential information. These architectures are, for example, extensively used in domains such as natural language processing [10] [11], where the context of a word in a sentence is useful to learn its characteristics. Figure 1 represents the overall architecture of CATREEN. CATREEN estimates the execution time of a basic block within a sequence of basic blocks in a similar way that an LSTM would process a paragraph in natural language to predict, for example, the sentiment (positive or negative) of a given sentence in this paragraph. The analogy is as follows: an instruction represents a word. An instruction is composed of an opcode and one or more operands, which can be treated as letters that constitute this word. A set of instructions will form a basic block (a sentence), and a sequence of basic blocks will therefore represent a paragraph.

The architecture of CATREEN is composed of five layers, depicted in Figure 1 from top to bottom. The first layer is a *tokenization and embedding* layer that pre-processes the assembly language, extracted from the machine code, and generates inputs that are understandable by the next layers. The next three layers are the central layers of CATREEN. They are LSTM layers (called RNN in the following): one to process an instruction (Instruction layer in Figure 1), one to process a basic block (BB layer in Figure 1), and one to process a sequence of basic blocks (Sequence layer in Figure 1). After that, and before processing the last BB in the sequence (the basic block under analysis), a residual connection is created using this vector and it is directly concatenate with the formed context (the result of processing all BBs), this helps the model to converge faster as we aim to estimate the execution time of the last BB in the sequence. Finally, The formed context and the BB under analysis vector are sent to a dense layer that is used to predict the final timing output (timing a basic block in the context of the previously executed basic blocks). More details on each layer are given below.

B. Tokenization and embedding

In order for the RNN to properly interpret basic blocks, an encoding of machine code into a sequence of integers is needed, where each integer represents an index (token) in a predefined dictionary. This is performed as follows. We pre-process the raw data (assembly code) by encoding each basic element of an instruction (opcode and operand) with a unique number (token), with a special treatment given to the operands: all immediate operands are encoded with the same token, the same treatment is done for all addresses which are encoded with the same token. Finally, addressing modes using registers

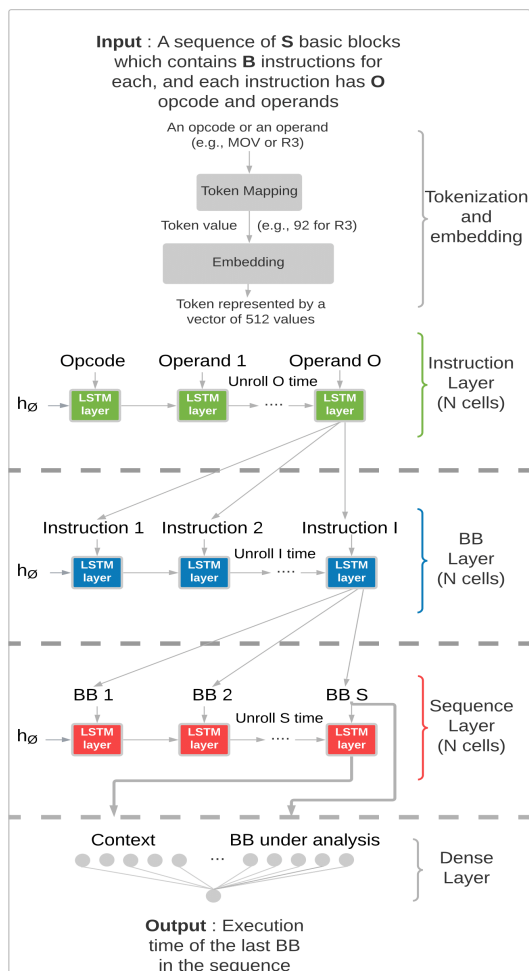


Fig. 1: Architecture of CATREEN. The input is a sequence of basic blocks consisting of a sequence of instructions which are themselves, sequences of operands/opcodes. CATREEN regresses a timing estimation for the last basic block in the input sequence.

are differentiated by assigning a distinct token value per pair (register number, addressing mode) with the addressing modes supported by the considered architecture (e.g., for the ARM targets, direct access, indirect access, and indirect access with offset).

In order to be suitable inputs for RNN layers, each token is again encoded into a distinct fixed-size vector of floats within the interval $[-1,1]$ using word2vec [12].

C. LSTM RNN layers of CATREEN

Since we do not, *a priori*, know what is the length of the sequences (i.e., how many basic blocks they contain) nor the basic blocks' length (i.e., how many instructions they contain) nor even the length of the instructions (i.e., how many operands they contain), we model these data with LSTM, which are suited to variable-length sequences. Three levels of LSTM are used for that purpose:

- The *Instruction Layer* (in green in Figure 1) is the first LSTM layer in CATREEN. It takes as input a sequence of embedded code operands/opcodes (i.e., an instruction). We loosely denote the length of the sequence of operands/opcodes by O (resp. I and B in the subsequent layers) even though it differs from one instruction to another. This LSTM layer processes the entire sequence of embedded tokens from instruction and produces a single final output representation for the entire instruction. To avoid direct influence between instructions, the hidden state is reset to its initial state h_\emptyset at each instruction (resp. at each BB and sequence), this helps to learn longer sequences where each representation is sent to the next stage to be better interpreted: the outputs (instructions) are treated one by one by the next LSTM layer (BB layer, for Basic Block layer). We loosely denote the dimension per layer by N even though it may differ in the three RNN layers. The selected hidden size per layer is given in Section IV-D.
- The *BB layer* (in blue in Figure 1) processes the sequence (of length I) of the treated instructions in a basic block and produces a new representation for each basic block. Again, all representations of a basic block are treated one element after one by the next LSTM layer (Sequence layer).
- Finally, the *Sequence layer* (in red in Figure 1) processes the sequence of S basic blocks within a sequence. Similarly to the other RNN layers, it produces at the end a representation of the sequence or what we call *Context*. Then, the value of *Context* is concatenated to the value of the *BB under analysis* (that we save from the previous layer BBs) and they are given as inputs to a fully connected linear layer connected to a single output which produces an estimate of the execution time of the last basic block of the sequence given the execution history.

IV. EXPERIMENTAL EVALUATION

Before evaluating CATREEN, we describe the data used to train and test our tool in Section IV-A, followed by the experimental setup in Section IV-B. In Section IV-C, we present the context-agnostic baselines used to assess our method. Then, we show that a careful selection of hyper-parameters is crucial for high precision results, for both CATREEN and the baseline predictors (Section IV-D). We then demonstrate that CATREEN outperforms the baseline context agnostic techniques, on both synthetic code (Section IV-E) and real code (MiBench2 and PolyBench programs, Section IV-F). A more detailed evaluation of CATREEN, that studies the inference time of CATREEN and provides further information on hyper-parameter selection can be found on a longer version of this paper [13]. Code and data can be found by following the link <https://gitlab.inria.fr/aamalou/CATREEN.git>.

A. Training and test data

The supervised training of our stacked LSTM network requires a large number of labeled data samples. In order to cover a large number of possible programs, we use synthetic codes. To test our model, both synthetic and real programs

are used. In both cases (training and testing), we use a hardware-assisted solution to obtain a timed execution trace from these programs (*ground truth* timing data used as a label for supervised learning). The execution trace (instruction-by-instruction representation of the program execution) is then pre-processed to constitute suitable inputs for our model.

1) *Synthetic data*: A code generator was developed to produce varied C source code programs. The code generator produces programs that randomly manipulate a selected number of statements and variables (the user provides parameters specifying the proportion of each element from a declared grammar). The code generator produces source code with loops, if-then-else constructs, and uses all the elementary types from the C language (*integers, floats, etc.*). Basic statements use the most common operations available in C (arithmetic and logical operations, shift and rotate operations, binary and unary operators or booleans, etc.). Arrays are also covered with various access modes: constant (access to a constant index), sequential, linear (affine loop indices), and random. The code generation ensures the absence of run-time errors by construction (out-of-bound array accesses, divide by zero).

2) *Real Data*: We evaluate our approach on a dataset of real programs, MiBench 2, a benchmark suite based on MiBench [14] and ported for IoT devices. Our tests were carried out on the *Automotive and Industrial Control* category of MiBench. We also use PolyBench v4.2 [15] a set of programs that has the specificity to manipulate nested loops.

3) *Ground truth timing generation*: Training and validating CATREEN have to be performed with accurate timing values. Moreover, the way the timing values are obtained should not change the timing of the code under execution (well-known as *probe-effect*). In CATREEN, we opted for a hardware-based solution, using the Joint Test Action Group (JTAG) interface. The J-Trace Pro trace solution from Segger [16] is used to connect to the JTAG interface of the target processor, alongside Ozone [17], a cross-platform debugger and performance analyzer. Ozone generates execution traces with a format of one line per machine instruction. Each line contains information about the value of the cycle counter after executing the instruction, the address of the instruction, its opcode and operands, and the corresponding assembly code.

4) *Data pre-processing*: To create the final data (both for synthetic and real programs), the execution trace generated by Ozone is processed as follows. First, BBs are extracted and the execution time of each BB is calculated. In order to eliminate outliers, the ground truth timing of each BB is obtained by using multiple executions and keeping the median value. CATREEN is trained by using a normalized timing value equal to the median value of the measurements divided by the number of instructions in the BB. Sequences that are too short to have enough contextual information (in our experiments, sequences shorter than 5 BBs) are filtered out.

B. Experimental setup

Experiments were performed on an STM32H7 board. The board features an ARM Cortex M7 processor, which has a

6-stage in-order pipeline, 16 KB instruction and data caches, and a branch predictor. We generated 1000 synthetic code snippets that we executed on the bare-metal target processor to eliminate any possible interference. After the pre-processing data phase to get sequences (as explained in Section IV-A4), 10000 different sequences were available for training, 2000 different sequences for validation, and 1000 different sequences for testing. Each subset of data (training, validation, test) comes from a separate set of programs to remove the bias it would otherwise introduce. We performed 1000 execution time measurements for each basic block, with cache warming before sampling (cache warming is performed by executing the program 20 times). All learning models were implemented in PyTorch [18] and were trained on NVIDIA GeForce RTX 2080 Ti. CATREEN training lasted four days for each setting (set of values of hyperparameters).

C. Baseline ML-based execution time predictors

The performance of CATREEN is compared with two context-agnostic execution time predictors. The most simple, albeit not naive one, is a Multi-Layer Perceptron regressor (denoted as *MLPr* in the following). An MLPr is a feed-forward neural network that does not take into account sequential information and requires a fixed-size input. Such a baseline has been successfully used for timing prediction in [19]. Specifically, the MLPr we have implemented takes as an input 233 static features of the basic blocks: the proportions of each type of machine instruction (e.g., MOV, ADD, LDR); the proportion of the different addressing modes (e.g., immediate, direct access, register indirect access, register indirect access with offset). We used a grid search algorithm to select proper MLPr hyperparameters (number of hidden layers, optimizer, learning rate, and the loss function). The result of this search gave us the following ideal (on the validation dataset) parameters: {hidden_layer_sizes=(256, 256), learning_rate_init=0.001, solver='adam', loss_function='mean squared error'}. Our second baseline is ITHEMAL [3], that similarly to CATREEN, uses RNNs for execution time prediction. More precisely, we compare CATREEN with two versions of ITHEMAL. The first one is a direct re-implementation of ITHEMAL from the original paper (denoted as *ITHEMAL-untuned* in the following), using the hyperparameters suggested by the authors. The re-implementation consisted in porting the tokenization/embedding step of ITHEMAL to the ARM instruction set and using a GPU for training instead of a parallel CPU. For the second version of ITHEMAL (called *ITHEMAL-tuned* hereafter), we have tuned the hyperparameters of the model to better fit the new data.

D. Hyperparameters tuning

Finding suitable hyperparameters is of utmost importance to guarantee the training convergence of the learned models on new data. Therefore, we show in the following how we have tuned (on validation data) our learning hyperparameters and how this can drastically improve the performance (in terms of Mean Absolute Percent Error, MAPE) of both CATREEN

and our closest competitor, ITHEMAL. The impact of three learning hyperparameters were studied:

- The *optimization algorithm*: Stochastic Gradient Descent (SGD), used in ITHEMAL [3] and ADAM optimizer widely used in deep learning.
- The *learning rate*: three learning rates; two constant values 10^{-3} and 10^{-4} , and an adaptive (that starts from 10^{-2} and decreases at each epoch by a factor of 10 until 10^{-4}).
- The *loss function*: the MAPE loss -used in ITHEMAL- and the symmetric Mean Absolute Percentage Error loss function (sMAPE) which is neutral regarding under or over-forecasting:

$$MAPE_{loss} = \frac{1}{n} * \sum_{i=0}^n \frac{|predict_i - actual_i|}{actual_i} \quad (1)$$

$$sMAPE_{loss} = \frac{2}{n} * \sum_{i=0}^n \frac{|predict_i - actual_i|}{predict_i + actual_i} \quad (2)$$

We kept the number of training epochs (20), the token embedding size 512, and the number of LSTM cells 512 the same during this study since this was the largest that we could do with our computation power.

The sensitivity of CATREEN and ITHEMAL to hyperparameter tuning is shown in Table I. Several observations can be made from the results. First, we observe that ADAM drastically reduces the generalization error of the learned models. Second, regarding the learning rate, the best results are obtained in all cases with a learning rate of 10^{-4} except when sMAPE+SGD is used, where 10^{-3} gives the best results (12% and 20%) both from CATREEN and ITHEMAL. The adaptive version (adapt) gives the worst results in most of the tests (except for ITHEMAL, when MAPE+SGD is used, in this case, it is the constant value 10^{-3} that gives the worst error percentage). Finally, we observe that sMAPE generally gives better results whatever the other parameters. MAPE is asymmetric by definition and puts a heavier penalty on overestimation errors than on underestimation errors. As a result, MAPE will favor models that are under-forecast. Overall, the results reported in Table I show that the hyperparameter selection has a huge impact on the model performance. They further show that the hyperparameter selection, as done in the original version of ITHEMAL, is sub-optimal for our data. In the next sections, we use CATREEN and ITHEMAL with their best-found hyperparameters.

E. Prediction performance on synthetic data

We compare the performance of MLPr, *ITHEMAL-untuned* [3], *ITHEMAL-tuned* and CATREEN on the 1000 test sequences. The results are reported in Table II and they are given as a MAPE, where the mean is computed over the 1000 test samples and the MAPE is expressed as a percentage. We also provide a Spearman score (rank correlation) and Pearson score (linear correlation) test for each method. The best results in the table are highlighted in bold. We observe that CATREEN gives the best results, with a MAPE of 11.4% an improvement of 18% compared to ITHEMAL tuned (which was trained on

the same data set for a fair comparison). This means that the model has benefited from the execution history of a given basic block to predict its timing, as further detailed below. In comparison, the other techniques that are context-agnostic provide more modest results; the simplest technique (MLPr) gives the worst results, which shows that handcrafted features (even though we introduced a large features number, here 233) are less discriminant than the learned ones.

F. Timing prediction on benchmarks

We evaluate our approach on real programs, the Automotive program set from MiBench 2 and PolyBench. The execution time is estimated for the entire program. For PolyBench programs, we choose to use small dataset option because of the flash memory size limitation for programs when running them on bare-metal (we also had to drop some programs that do not fit on the flash memory). The same experimental process as for the synthetic dataset (see Section IV-A) was used to (i) generate the sequences of basic blocks serving as inputs for the timing predictions, (ii) obtain the ground truth timing values. Table III reports the Absolute Percentage Error ($APE = 100 * \frac{|prediction - actual|}{actual}$) for the different programs, for CATREEN, ITHEMAL-tuned and MLPr. The results of ITHEMAL-untuned were not good enough to deserve their reporting in the table. The best APE values are highlighted in bold. The results show that CATREEN outperforms ITHEMAL-tuned and MLPr, with a lower APE on all benchmarks (16.5% on average), except for the mvt program. In general, we can observe that ITHEMAL-tuned tends to have good results in PolyBench programs which can be explained by the nested loops that constitute them. These loops make the execution time of the BBs steady (in best case), which can hide the context effect on these BBs. The use of a completely context-agnostic technique like MLPr gives, as expected, the worst predictions. The average prediction error of CATREEN on the tested benchmarks is comparable (although slightly higher) to that obtained on the synthetic codes. This shows that our synthetic dataset is globally representative of real code. The slightly higher error may come from two factors: (i) Not sufficiently representative code generation regarding array accesses. This can be addressed by improving our dataset, for instance, by synthesizing training data from real code; (ii) Too compact tokenization of memory accesses (CATREEN does not consider different memory addresses as distinct tokens, in order to avoid an explosion of tokens number).

V. CONCLUSION

We have presented CATREEN, an ML-based program timing predictor. CATREEN leverages recurrent neural networks to predict the steady-state execution time of basic blocks in a program while taking into account the *execution context* formed by the previously executed basic blocks. Experimental results have shown that CATREEN's timing predictions are 18% (27.6%) better than those estimated by state-of-the-art context-agnostic ML-based tools on synthetic (respectively real

TABLE I: MAPE performance of CATREEN and ITHEMAL, for different learning hyperparameters (loss function, optimizer, learning rate)

Loss function	MAPE						sMAPE					
	SGD			ADAM			SGD			ADAM		
Optimizer												
Learning Rate	10^{-3}	10^{-4}	adapt	10^{-3}	10^{-4}	adapt	10^{-3}	10^{-4}	adapt	10^{-3}	10^{-4}	adapt
CATREEN	18%	18%	19%	11%	13%	17%	12%	16%	17%	11%	10%	13%
ITHEMAL	39%	24%	38%	18%	17%	21%	20%	27%	28%	25%	25%	25%

TABLE II: MAPE of timing predictions for 1000 basic blocks by four different ML-based methods.

ML technique	MAPE	Spearman	Pearson
MLPr	28.8%	0.974	0.966
ITHEMAL-untuned	27.4%	0.973	0.968
ITHEMAL-tuned	13.9%	0.972	0.975
CATREEN	11.4%	0.983	0.977

TABLE III: Execution time APE on MiBench (*automotive*) and PolyBench for CATREEN, ITHEMAL-tuned and MLPr

Model	CATREEN	ITHEMAL-tuned	MLPr
basicmath	2.1%	13.4%	67.1%
bitcount	9.3%	29.3%	68.3%
qsort	19.0%	37.5%	68.6%
susan:corner+edges	18.9%	35.3%	107.1%
susan:smoothing	2.1%	26.1%	67.1%
covariance	15.8%	18.4%	81.7%
gemm	19.8%	21.1%	30.2 %
gemver	20.9%	22.5%	79.9 %
gesummv	15.0%	15.3%	28.8 %
symm	17.8%	18.4%	35.0 %
syrk	18.3%	18.6%	30.0%
trmm	22.6%	24.8%	26.1%
2mm	17.2%	21.3%	25.3 %
3mm	17.7%	21.3%	27.2 %
atax	23.2%	24.8%	25.5%
bicg	18.7%	20.2%	28.1 %
dotigen	21.5%	25.5%	26.2 %
mvt	25.7%	25.3%	42.1 %
cholesky	20.6%	22.4%	24.2 %
gramschmidt	16.9%	20.1%	21.8%
lu	18.7%	22.1%	106.5%
ludcmp	11.0%	16.9%	144.0%
trisolv	9.1%	11.9%	17.2%
floyd_warshall	11.5%	21.6%	268.7%
Avg.	16.5%	22.8%	62.2%

programs). [20] has observed that LSTM models can use a maximum of 200 context tokens which prevent them from learning longer-term information (i.e., they can remember sequences of hundreds but not thousands). An area for improvement, is to use *Transformers* [21] which have demonstrated a great ability to address text sequence learning. Using attention mechanisms, they can increase the quality of the contextual information drawn from the sequence under analysis and its size. We thus plan to consider Transformers for timing estimation.

ACKNOWLEDGMENTS

The authors are very grateful to Pierre Michaud and Hugo Reymond for their comments on earlier drafts of the paper.

REFERENCES

- [1] Z. Wang and M. O’Boyle, “Machine learning in compiler optimization,” *Proceedings of the IEEE*, vol. 106, 2018.
- [2] A. Ltd, “Cycle Models.” [Online]. Available: <https://www.arm.com/products/development-tools/simulation/cycle-models>
- [3] C. Mendis, A. Renda, S. Amarasinghe, and M. Carbin, “Ithemal: Accurate, portable and fast basic block throughput estimation using deep neural networks,” in *Int. Conference on machine learning*. PMLR, 2019.
- [4] L. Huang, J. Jia, B. Yu, B.-G. Chun, P. Maniatis, and M. Naik, “Predicting execution time of computer programs using sparse polynomial regression,” *Advances in neural information processing systems*, vol. 23, 2010.
- [5] F. Ritter and S. Hack, “Pmevo: portable inference of port mappings for out-of-order processors by evolutionary optimization,” in *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*, 2020, pp. 608–622.
- [6] C. Cummins, P. Petoumenos, Z. Wang, and H. Leather, “End-to-end deep learning of optimization heuristics,” in *26th Int. Conference on Parallel Architectures and Compilation Techniques*. IEEE, 2017.
- [7] M. Rapp, A. Pathania, T. Mitra, and J. Henkel, “Neural network-based performance prediction for task migration on s-nuca many-cores,” *IEEE Transactions on Computers*, vol. 70, no. 10, pp. 1691–1704, 2020.
- [8] A. Tousei and M. Luján, “Comparative analysis of machine learning models for performance prediction of the spec benchmarks,” *IEEE Access*, vol. 10, pp. 11 994–12 011, 2022.
- [9] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [10] J. Wang, L.-C. Yu, K. R. Lai, and X. Zhang, “Investigating dynamic routing in tree-structured LSTM for sentiment analysis,” in *Int. Conference on Empirical Methods in Natural Language Processing and Int. Joint Conference on NLP*, 2019.
- [11] Y. Samih, S. Maharjan, M. Attia, L. Kallmeyer, and T. Solorio, “Multilingual code-switching identification via lstm recurrent neural networks,” in *Proceedings of the Second Workshop on Computational Approaches to Code Switching*, 2016.
- [12] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, “Distributed representations of words and phrases and their compositionality,” *Advances in neural information processing systems*, vol. 26, 2013.
- [13] A. N. Amalou, É. Fromont, and I. Puaut, “CATREEN: Context-Aware Code Timing Estimation with Stacked Recurrent Networks,” Sep. 2022, preprint. [Online]. Available: <https://hal.inria.fr/hal-03776508>
- [14] M. R. Guthaus, J. S. Ringenberg, D. Ernst, T. M. Austin, T. Mudge, and R. B. Brown, “Mibench: A free, commercially representative embedded benchmark suite,” in *4th IEEE international workshop on workload characterization*, 2001.
- [15] L.-N. Pouchet, “The polyhedral benchmark suite,” *On-line: http://www.cse.ohiostate.edu/pouchet/software/polybench*, 2012.
- [16] Segger, “J-Trace PRO – The Leading Trace Solution.” [Online]. Available: <https://www.segger.com/products/debug-probes/j-trace/>
- [17] S. M. GmbH, “Ozone User Guide & Reference Manual,” p. 348. [Online]. Available: <https://www.segger.com/>
- [18] PyTorch, “PyTorch.” [Online]. Available: <https://www.pytorch.org>
- [19] A. N. Amalou, I. Puaut, and G. Muller, “WE-HML: hybrid WCET estimation using machine learning for architectures with caches,” in *27th IEEE Int. Conference on Embedded and Real-Time Computing Systems and Applications, RTCSA*, 2021.
- [20] U. Khandelwal, H. He, P. Qi, and D. Jurafsky, “Sharp nearby, fuzzy far away: How neural language models use context,” *arXiv preprint arXiv:1805.04623*, 2018.
- [21] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Advances in neural information processing systems*, 2017.