BRNO UNIVERSITY OF TECHNOLOGY FACULTY OF INFORMATION TECHNOLOGY

Excel@FIT 2023
http://excel.fit.vutbr.cz

Estimation of Algorithm Execution Time Using Machine Learning

Martin Buchta

Abstract
This work aims to predict the execution time of k-Wave ultrasound simulations on supercomputers based on a given domain size. The program uses MPI and can be run on multiple nodes. Prediction models were developed using symbolic regression and neural networks, both of which trained on captured data and compared against each other.
The results demonstrate that the models outperform existing solutions. Specifically, the symbolic regression model achieved an average error of 5.64 % for suitable tasks, while the neural network model achieved an average error of 8.25 % on unseen domain sizes and across all tasks, including those not optimized for k-Wave simulations.
This work contributes a new, more accurate model for predicting execution time, and compares the effectiveness of neural networks and symbolic regression for this specific type of regression problem. Overall, these findings suggest that new models will have important practical applications in the field of k-Wave ultrasound simulations.
*xbucht28@stud.fit.vutbr.cz, *Faculty of Information Technology, Brno University of Technology*

1. Introduction

The k-Wave toolbox is a popular tool for performing ultrasound simulations. To speed up the simulations, the program can use multiple processors across nodes that communicate with each other using MPI. The execution time of the simulation depends on several factors, including the simulation domain size, number of nodes on the supercomputer, and number of computing units in MPI (ranks).

The ability to accurately predict the execution time of a simulation is crucial for several reasons. Firstly, it helps new users to estimate the time required for the simulation, which is particularly important when submitting a job on a supercomputer that requires users to specify the maximum execution time. Secondly, predicting the runtime can optimize the flow of multiple simulations on a pre-defined number of nodes[1], leading to more efficient resource allocation. Additionally, a good regression model can reduce the cost of the simulation by identifying the optimal set of resources required to run the simulation.

The aim of this work is to develop a model for predicting the execution time of k-Wave ultrasound simulations using symbolic regression and neural networks. Both methods are compared and evaluated on unseen domain sizes. In chapter 2 existing solutions are reviewed. In chapter 3 the data are described and a solution to address the issue is proposed. The section 4 discusses the accuracy of the models. The section 5 provides a brief conclusion.

2. Related works

The estimation of runtime for k-Wave simulations is often done manually, based on the user's experience or the closest simulation result found. In some cases, the execution time is predicted to be the maximum allowable runtime of 24 hours, which is not an accurate solution usually.

Marta Jaroš proposed a method for estimating the k-Wave runtime using interpolation and splines [2]. However, this method resulted in an average error of up to 15 % on unseen domain sizes.

3. Proposed solution

To accurately predict the runtime of k-Wave ultrasound simulations, this work proposes two models:

symbolic regression and neural networks. This chapter provides a brief overview of the data used to train the models and the models themselves.

3.1 Data

The dataset used in this work consists of 2 080 k-Wave runs, including the number of grid points in the X, Y, and Z axes, the number of nodes, the number of ranks, and the execution time. (See the Image 2.) In addition to these fields, additional features were added to the dataset, such as the portion of ranks to all available processors, the maximum prime factor of the axis size, and the share of ranks with not sufficient amount of data.

Two datasets were created, with the aim of distinguishing tasks suitable for k-Wave. The *Filtered dataset* comprises of domains with prime factors of each axis that are less than or equal to 7 (due to MPI), while the *Non-Filtered dataset* includes all the measured data, with domain sizes that may not be convenient as for prime factors. Generally, the simulation time of the domains with prime factors greater than 7 is less predictable and takes more time. Each dataset was divided further into train, validation, and test portions, based on the simulation domain size, such that the test set contains simulation domains that are not present in the train or validation sets.

3.2 Models

To estimate the runtime of k-Wave simulations, two models were trained on each dataset: symbolic regression and neural networks.

The symbolic regression models were trained using genetic algorithms in HeuristicLab, with the number of generations limited to 1 000 000 and each generation consisting of 100 individuals. The expressions (trees) were limited to the size of 42 and to the height of 16, with a mutation probability set to 10 %.

The neural networks were trained using TensorFlow, with an architecture consisting of six hidden fully connected layers, each with 256 neurons and ReLU activation. After each layer, a dropout layer was added. The training of the neural networks was stopped by an early stopping condition with patience of 500 epochs.

4. Prediction accuracy

In this chapter, the accuracy of the models is evaluated. The evaluation was performed by computing the average relative error of the models on the test split.

4.1 Filtered Dataset

On the filtered dataset, the average relative error was found to be 5.64 % with the symbolic regression model and 12.46 % with the neural networks model. This indicates that the symbolic regression model performed better than the neural network model on this dataset. (See Table 1 and Image 4.)

4.2 Non-Filtered Dataset

On the non-filtered dataset, the symbolic regression model had an average relative error of 13.55 %, while the neural networks model had an average relative error of 8.25 % on the test set. It is interesting to note that the neural networks model outperformed the symbolic regression model on this dataset. (See Table 2 and Image 5.)

5. Conclusions

Overall, the symbolic regression model performed better than the neural networks model on the *Filtered Dataset*, while the neural networks model performed better on the *Non-filtered dataset*.

The developed models provide a more accurate estimation of the execution time than manual estimation or simply predicting the longest runtime. They can save valuable computing resources and reduce the time required to run simulations. This work demonstrates the feasibility of using symbolic regression and neural networks for predicting the execution time of k-Wave simulations. Future works can explore other machine learning techniques or different features to further improve the accuracy of the models.

References

[1] Marta Jaroš, E. Bradley Treeby, Jiří Jaroš, and Panyiotis Georgiou. k-dispatch: A workflow management system for the automated execution of biomedical ultrasound simulations on remote computing resources. In *Proceedings of the Platform for Advanced Scientific Computing Conference, PASC 2020*, pages 1–10. Association for Computing Machinery, 2020.

[2] Marta Jaroš, Tomáš Sasák, E. Bradley Treeby, and Jiří Jaroš. Estimation of execution parameters for k-wave simulations. In *High Performance Computing in Science and Engineering. HPCSE 2019*, Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), pages 116–134. Springer Nature Switzerland AG, 2021.