



Neural Network Precision Tuning Using Stochastic Arithmetic

Quentin Ferro^{1(✉)}, Stef Graillat¹, Thibault Hilaire¹, Fabienne Jézéquel^{1,2},
and Basile Lewandowski¹

¹ Sorbonne Université, CNRS, LIP6, 75005 Paris, France

{quentin.ferro, stef.graillat, thibault.hilaire, fabienne.jezequel}@lip6.fr

² Université Paris-Panthéon-Assas, 75005 Paris, France

Abstract. Neural networks can be costly in terms of memory and execution time. Reducing their cost has become an objective, especially when integrated in an embedded system with limited resources. A possible solution consists in reducing the precision of their neurons parameters. In this article, we present how to use auto-tuning on neural networks to lower their precision while keeping an accurate output. To do so, we use a floating-point auto-tuning tool on different kinds of neural networks. We show that, to some extent, we can lower the precision of several neural network parameters without compromising the accuracy requirement.

Keywords: Precision · Neural networks · Auto-tuning · Floating-point · Stochastic arithmetic

1 Introduction

Neural networks are nowadays massively used and are becoming larger and larger. They often need a lot of resources, which can be a problem, especially when used in a critical embedded system with limited computing power and memory. Therefore it can be very beneficial to optimise the numerical formats used in a neural network. This article describes how to perform precision auto-tuning of neural networks. From a neural network application and an accuracy requirement on its results, it is shown how to obtain a mixed precision version using the PROMISE tool [15]. A particularity of PROMISE is the fact that it uses stochastic arithmetic [38] to control rounding errors in the programs it provides.

Minimizing the format of variables in a numerical simulation can offer advantages in terms of execution time, volume of data exchanged and energy consumption. During the past years several algorithms and tools have been proposed for precision auto-tuning. On the one hand, tools such as FPTuner [8], Salsa [10], Rosa/Daisy [11, 12], TAFFO [5], POP [2] rely on a static approach and are not intended to be used on very large code. On the other hand, dynamic tools such as CRAFT HPC [24], Precimonious [33], HiFPTuner [16], ADAPT [30], FloatSmith [25], PROMISE [15], have been proposed for precision auto-tuning

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2022

O. Isac et al. (Eds.): NSV 2022/FoMLAS 2022, LNCS 13466, pp. 164–186, 2022.

https://doi.org/10.1007/978-3-031-21222-2_10

in large HPC code. Moreover, tools have been recently developed for precision auto-tuning on GPUs: AMPT-GA [22], GPUMixer [23], GRAM [18]. A specificity of PROMISE lies in the fact that it provides mixed precision programs validated owing to stochastic arithmetic [38], whereas other dynamic tools rely on a reference result possibly affected by rounding errors. PROMISE has been used in various applications based on linear algebra kernels, but not yet for precision auto-tuning in deep neural networks.

While much effort has been devoted to the safety and robustness of deep learning code (see for instance [13, 27, 28, 32, 34]) a few studies have been carried out on the effects of rounding error propagation on neural networks. Verifiers such as MIPVerify [36] are designed to check properties of neural networks and measure their robustness. However, the impact of floating-point arithmetic both on neural networks and on verifiers is pointed out in [42]. Because of rounding errors, the actual robustness and the robustness provided by a verifier may radically differ. In [9, 41] it is shown how to control the robustness of different neural networks with greater efficiency using interval arithmetic based algorithms. In [26] a software framework is presented for semi-automatic floating-point error analysis in the inference phase of deep neural networks. This analysis provides absolute and relative error bounds owing to interval and affine arithmetics.

Precision tuning of neural networks using fixed-point arithmetic has been studied in [3]. Owing to the solution of a system of linear constraints, the fixed-point precision of each neuron is determined, taking into account a certain error threshold.

In this article, we consider floating-point precision tuning, which is also studied in [20]. Focusing on interpolator networks, i.e. networks computing mathematical functions, the authors propose an algorithm that takes into account a given tolerance δ on the relative error between the assumed correctly computed function and the function computed by the network. The main difference with the present article is the auto-tuning algorithm: in [20] the precision is optimized by solving a linear programming problem, while PROMISE uses a hypothesis-trial-result approach through the Delta-Debug algorithm [40]. Furthermore, the algorithm in [20] relies on a reference result that may be altered by rounding errors, while PROMISE uses stochastic arithmetic for the numerical validation of its results.

Stochastic arithmetic uses for rounding error estimation a random rounding mode: the result of each arithmetic operation is rounded up or down with the same probability. As a remark, another stochastic rounding often used in neural network training and inference uses a probability that depends on the position of the exact result with respect to the rounded ones (see for instance [14, 17, 29, 31, 35, 39]). This stochastic rounding does not aim at estimating rounding errors, it enables the update of small parameters and avoids stagnations that may be observed with round to nearest.

In this work, we consider tuning the precision of an already trained neural network. One of our contributions is a methodology for tuning the precision of a neural network using PROMISE in order to obtain the lowest precision for each

of its parameters, while keeping a certain accuracy on its results. We present and compare the results obtained for different neural networks: an approximation of the sine function, an image classifier processing the MNIST dataset (2D pictures of handwritten digits), another image classifier using this time convolutional layers and processing the CIFAR10 dataset (3D images of different classes), and the last one introduced in [4] and used in [26] that aims at approximating a Lyapunov function of a nonlinear controller for an inverted pendulum.

After a preliminary reminder on deep neural networks and stochastic arithmetic in Sect. 2, Sect. 3 describes our methodology and Sect. 4 presents our results considering the different neural networks previously mentioned.

2 Preliminary

2.1 Neural Networks

An artificial neural network is a computing system defined by several neurons distributed on different layers. Generally, we consider dense layers that take as an input a vector and in which the main computation is a matrix-vector product. In this case, from one layer to another, a vector of neurons $x^{(k)} \in \mathbb{R}^{n_k}$ with $k \in \mathbb{N}$ is transformed into a vector $x^{(k+1)} \in \mathbb{R}^{n_{k+1}}$ by the following equation

$$x^{(k+1)} = g^{(k)}(W^{(k)}x^{(k)} + b^{(k)}) \quad (1)$$

where $W^{(k)} \in \mathbb{R}^{n_{k+1} \times n_k}$ is a weight matrix, $b^{(k)} \in \mathbb{R}^{n_{k+1}}$ a bias vector and $g^{(k)}$ an activation function. The activation function is a non-linear and often monotonous function. Most common activation functions are described below.

- Sigmoid: computes $\sigma(x) = 1/(1 + e^{-x}) \forall x \in \mathbb{R}$
- Hyperbolic Tangent: applies $\tanh(x) \forall x \in \mathbb{R}$
- Rectified Linear Unit: $\text{ReLU}(x) = \max(x, 0) \forall x \in \mathbb{R}$
- Softmax: normalizes an input vector x into a probability distribution over the output classes. For each element x_i in x , $\text{softmax}(x_i) = e^{x_i} / \sum e^{x_j}$

Figure 1 shows a graphical representation of a neural network with two layers.

Dense layers are sometimes generalized to multidimensional arrays and involve tensor products. Other layers exist such as convolution layers often used on multi-dimensional arrays to extract features out of the data. To do so, a convolution kernel is applied to the input data to produce the output. Different kinds of layers, for instance pooling layers and flatten layers, do not require weight nor bias. Pooling layers reduce the size of the input data by summarizing it by zones given a function such as the maximum or the average. Flatten layers are intended to change the shape of data, from a 3D tensor to a vector for example and can be used to pass from a convolutional layer to a dense layer.




Fig. 1. Neural network with two layers

2.2 Floating-Point Arithmetic

A floating-point number x in base β is defined by:

$$x = (-1)^s \times m \times \beta^e \quad (2)$$

with s its sign being either 1 or 0, m its significand being an integer and e its exponent being also an integer. In this paper, we consider binary floating-point numbers, i.e. numbers in base $\beta = 2$ that adhere to the IEEE 754 Standard [1]. The IEEE 754 Standard defines different formats with a fixed number of bits for the significand and the exponent. The number of bits for the significand is the precision p , hence the significand can take values ranging from 0 to $\beta^p - 1$. The exponent e ranges from e_{min} to e_{max} with $e_{min} = 1 - e_{max}$ and $e_{max} = 2^{len(e)-1} - 1$, with $len(e)$ the exponent length in bits. The sizes of the three different formats used in this paper, commonly named half, single (or float), and double, are sum up in Table 1. As a remark, another 16-bit format called bfloat16 exists, for example on ARM NEON CPUs. Owing to their 8-bit-large exponent, bfloat16 numbers benefit from a wider range, but have a lower significand precision (8 bits instead of 11).

Table 1. Basic binary IEEE 754 formats

Name	Format	Length	Sign	Significand length ^a	Exponent length
Half	binary16	16 bits	1 bit	11 bits	5 bits
Single	binary32	32 bits	1 bit	24 bits	8 bits
Double	binary64	64 bits	1 bit	53 bits	11 bits

^a Including the implicit bit (which always equals 1 for normal numbers, and 0 for subnormal numbers). The implicit bit is not stored in memory.

2.3 Discrete Stochastic Arithmetic (DSA)

Discrete Stochastic Arithmetic (DSA) is a method for rounding error analysis based on the CESTAC method [7, 37]. The CESTAC method allows the estimation of round-off error propagation that occurs when computing with floating-point numbers. Based on a probabilistic approach, it uses a random rounding

mode: at each operation, the result is either rounded up or down with the same probability. Using this rounding mode, the same program is run N times giving us N samples R_1, \dots, R_N of the computed result R . The accuracy of the computed result (i.e. its number of exact significant digits) can be estimated using Student's law with the confidence level 95%. In practice the sample size is $N = 3$. Indeed, it has been shown that $N = 3$ is in some reasonable sense the optimal value. The estimation with $N = 3$ is more reliable than with $N = 2$ and increasing the sample size does not improve the quality of the estimation. Theoretical elements can be found in [6, 37].

The CADNA [6, 21, 37] (Control of Accuracy and Debugging for Numerical Applications) software¹ implements DSA in code written in C, C++ or Fortran. It introduces new variable types, the stochastic types. Each stochastic variable contains three floating-point values and one integer being the exact number of correct digits. CADNA can print each computed value with only its exact significant digits. In practice, owing to operator overloading, the use of CADNA only requires to change declaration of variables and input/output statements.

2.4 The PROMISE Software

The PROMISE software² aims at reducing the precision of the variables in a given program. From an initial code and a required accuracy, it returns a mixed precision code, lowering the precision of the different variables while keeping a result that satisfies the accuracy constraint. To do so, some variables are declared as custom typed variables that PROMISE recognizes. PROMISE will consider tweaking their precisions. Different variables can be forced to have the same precision by giving them the same custom type. It may be useful to avoid compilation errors or casts of variables.

PROMISE computes a reference result using CADNA and relies on the Delta-Debug algorithm [40] to test different type configurations, until a suitable one lowering the precision while satisfying the accuracy requirement is found. PROMISE provides a transformed program that can mix half, single and double precision. Half precision can be either native on CPUs that support it or emulated using a library developed by C. Rau³. PROMISE dataflow is presented in Fig. 2. After computing the reference result, PROMISE tries to lower the precision of the variables from double to single, then from single to half, using twice the Delta-Debug algorithm. The accuracy requirement may concern one or several variables (e.g. in an array). PROMISE checks that the number of common digits between the computed result(s) and the reference result(s) is at least the required accuracy. In the case of several variables, the requirement has to be fulfilled by all of them.

¹ <http://cadna.lip6.fr>.

² <http://promise.lip6.fr>.

³ <http://half.sourceforge.net>.

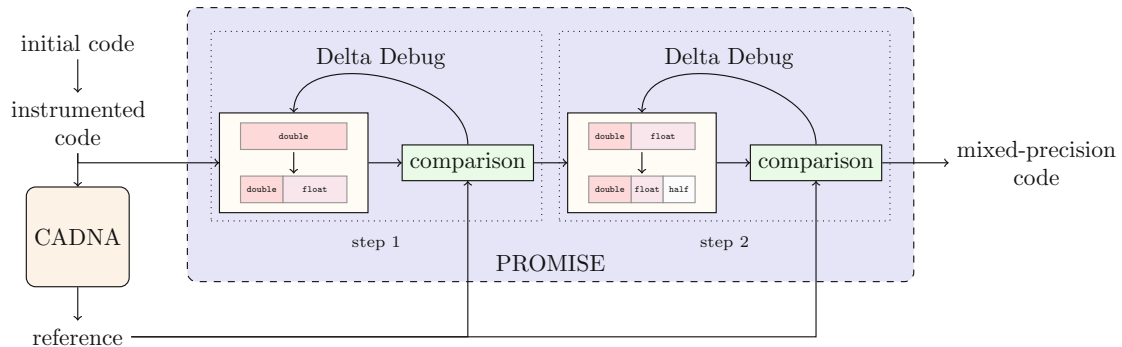


Fig. 2. PROMISE dataflow

3 Methodology

For neural network management, we use Python language with either Keras⁴ or PyTorch⁵. Keras and Pytorch are two open-source Python libraries that implement structures and functions to create and train neural network models. Both of them also allow us to save our model in HDF5 (Hierarchical Data Format)⁶, a file format designed to store and organize large data. HDF5 uses only two types of objects: datasets that are multidimensional arrays of homogeneous type, and groups, which contain datasets or other groups. HDF5 files can be read by Python programs using the h5py package. The associated data can be manipulated with Pandas⁷, a Python library that proposes data structures and operations to manage large amount of data.

Keras is used to develop, train and save our neural network models, except in the case of the inverted pendulum which uses PyTorch. As already mentioned in the introduction, precision tuning is performed on trained models, hence in the inference stage. The process path is summarized in Fig. 3. For each neural network, we first convert the HDF file to CSV files using a Python script. The script loads the HDF file, stores the parameters in Pandas DataFrames, and then saves the parameters in CSV files using the Pandas DataFrame *to_csv* function. For each layer that needs it, we create a CSV file with the weights of the layer and a CSV file with the bias of the layer. Indeed, some layers do not need weights nor bias, for example flattening layers that only change the data shape (from 2 dimensions to 1 dimension for example). Secondly, we use the data in the CSV files to create a C++ program, once again using a Python script that reads the CSV files and creates the necessary variables and computation. The translation scripts are based on the work done in the keras2c⁸ library. Once the C++ version created, we apply PROMISE on it.

⁴ <https://keras.io>.

⁵ <https://pytorch.org>.

⁶ <https://www.hdfgroup.org>.

⁷ <https://pandas.pydata.org>.

⁸ <https://f0uriest.github.io/keras2c/>.

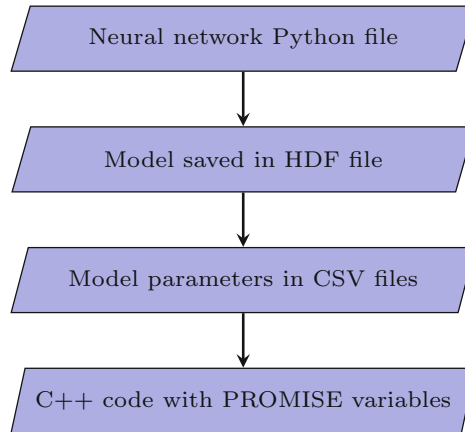


Fig. 3. Flowchart of the translation from a Python neural network to a C++ program with PROMISE variables

4 Experimental Results

Results obtained for four different neural networks are presented in this section. For the neural networks using a database, `test_data[i]` refers to $(i+1)^{th}$ test input provided by the database. PROMISE is applied to each neural network considering one type by neuron (half, single or double), then one type per layer, i.e. all the parameters of a layer have the same precision. In our analysis, the difference between the two approaches lies in the number of different type declarations in the code. However, it must be pointed out that, in dense layers, having one type per neuron implies independent dot products, whereas having one type per layer would enable one to compute matrix-vector products that could perform better.

In our experiments, the input is in double precision. In accordance with Fig. 2, for any neural network, the reference value is the value computed at the very first step of PROMISE. All the results presented in this section have been obtained on a 2.80 GHz Intel Core i5-8400 CPU having 6 cores with 16 GB RAM except indicated otherwise.

4.1 Sine Neural Network

To approximate the sine function, we use a classical densely-connected neural network with 3 layers. It is a toy problem, since using a neural network to compute sine is not necessary. However, this simple example validates our approach. The tanh activation function is used in the 3 different dense layers. The layers have respectively 20, 6 and 1 neuron(s) and the input is a scalar value x . Figure 4 presents the computation carried out by the neural network, considering one type per neuron. Colored variables are PROMISE variables, the precision of which can be tweaked. Variables with the same color have the same precision. The parameters of a neuron (weight(s) and bias) have the same color, hence the same type. The output type of each layer is also tuned. In this example, we assign types to 27 neurons and 3 outputs.

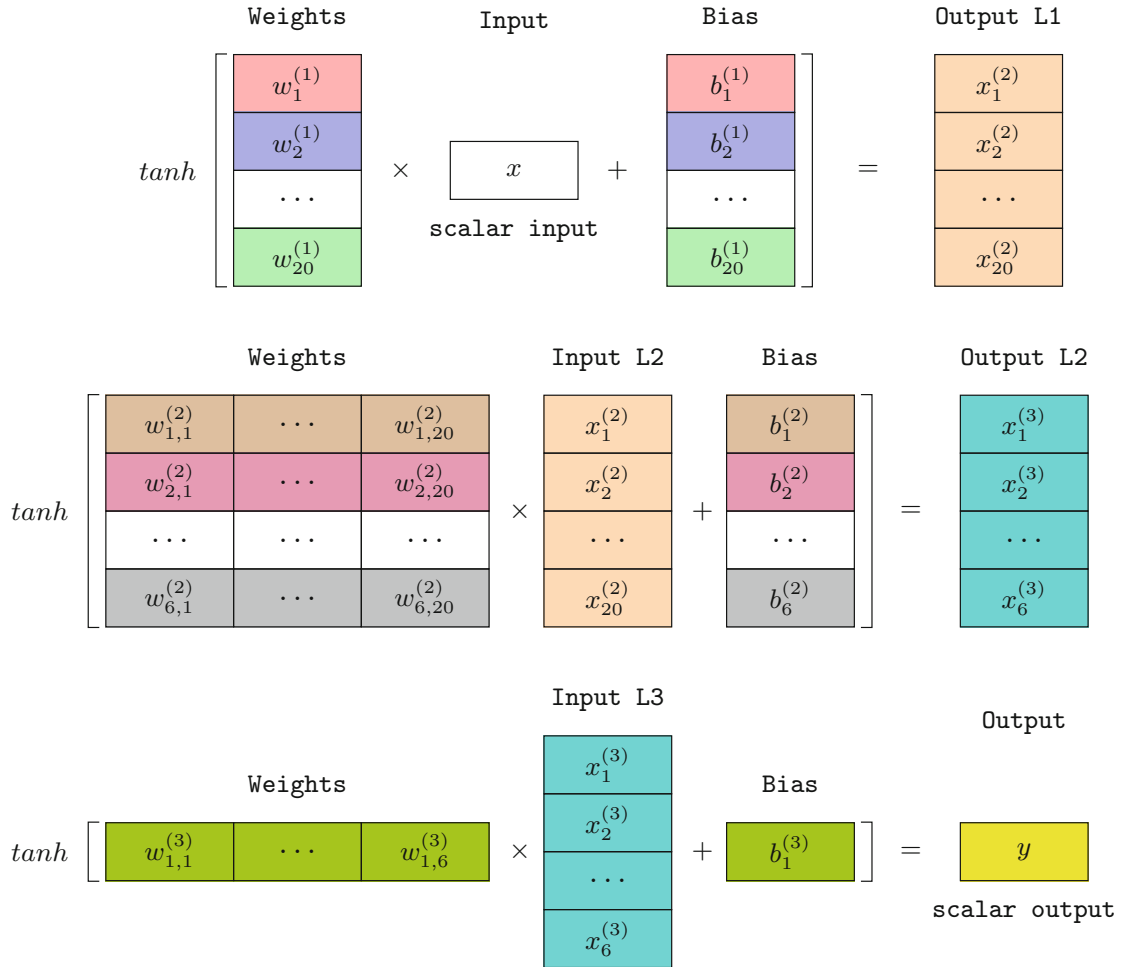


Fig. 4. Computation carried out in the sine neural network

Figure 5 displays the distribution of the different types with input value 0.5 considering one type per neuron. The x -axis presents the required accuracy on the results, i.e. the number of significant digits in common with the reference result computed using CADNA. We can notice the evolution of the distribution depending on the number of exact significant digits required on the result. As expected, first we only have half precision variables, then some of them start to be in single, then in double precision, until eventually all of them are in double precision. Therefore, requiring the highest accuracy is not compatible with lowering the precision in this neural network. But still, a good compromise can be found, since we only have single and half precision variables for a required accuracy up to 7 digits, and still have 1/3 of single precision variables for a required accuracy up to 9 digits.

Figure 5 also presents the computation time of PROMISE in seconds for each required accuracy. It consists of the time to compute the reference result, and the time to apply the Delta-Debug algorithm twice (from double to single precision then from single to half precision), compiling and executing the tested distribution each time. It can be noticed that the computation time (less than 2 min) remains reasonable given the 3^{30} possibilities.




Fig. 5. Number of variables of each type and computation time for the sine neural network with input value 0.5

Figure 6 shows the type distribution considering one precision per layer. This approach per layer enables one to decrease the execution time of PROMISE, but it does not really help lowering the precision of the network parameters. Indeed, each time a parameter in a layer requires a higher precision, all the parameters of the same layer pass in higher precision. But still, it can be noticed that the first layer (that represents 2/3 of the neurons) stays in half precision for a required accuracy up to 4 digits.




Fig. 6. Precision of each layer for the sine neural network with input value 0.5

Figures in Appendix A.1 and Appendix A.2 show that the input value can have a slight impact on the type distribution. We compare the results for two input values randomly chosen: 0.5 and 2.37. From Fig. 5, with input value 0.5, PROMISE provides a type distribution with 8 half precision neurons and 22 single precision neurons for a required accuracy of 5 digits. From Appendix A.1, with input value 2.37, only 4 half precision neurons are obtained for a required accuracy of 5 digits. Actually, if 3, 4, or 5 digits are required, less neurons are in half precision than with input 0.5. But if 8 digits are required, one neuron remains in half precision with input 2.37, while no neuron is in half precision with input 0.5. Nevertheless, the type distribution with respect to the required accuracy remains globally the same, parameters all starting as half precision variables, then passing to higher precision.

4.2 MNIST Neural Network

Experiments have been carried out with an image classification neural network processing MNIST data of handwritten digits⁹. This neural network also uses classical dense layers. The main difference in this case is that the entry is a vector of size 784 (flatten image) and the output is a vector of size 10. This neural network consists of two layers: the first one with 64 neurons and the activation function ReLU, and the second one with 10 neurons and the activation function softmax which provides the probability distribution for the 10 different classes. Considering as previously one type per neuron, plus one type for the output of each layer, 76 different types have to be set either to half, single, or double precision.

Figure 7 shows the type distribution considering one type per neuron with the input image test_data[61]: the 62nd test data out of the 10,000 provided by MNIST. The x -axis represents the required accuracy on the output consisting in a vector of size 10. The maximum accuracy on the output is 13 digits, higher expectation could not be matched. Such high accuracy is nonetheless not realistic because not necessary for a classifier. However, exhaustive tests have been performed: all possible accuracies in double precision have been successively required.

The main difference with the sine neural network lies in the fact that a significant number of neurons stay in half precision no matters the required accuracy. Depending on the input, around 50% of neurons can stay in half precision and sometimes nearly 60% as shown in Appendix B.2. Thus, applying PROMISE to this neural network, even when requiring the highest accuracy, can help lowering the precision of its parameters. Neurons that keep the lowest precision are not the same depending on the input, but they always belong to the first layer. Hence, the first layer seems to have less impact on the output accuracy than the second one.

⁹ <http://yann.lecun.com/exdb/mnist>.

The computation times also reported in Fig. 7 are much higher than for the sine approximation network. For the majority of the accuracy requirements, more than 15 min are necessary to obtain a mixed precision version of the neural network. The MNIST neural network has one layer less than the sine approximation network, but more neurons. Since we consider one precision per neuron, the number of possible type configurations (3^{76}) is much higher, hence the computation time difference. However, the execution time of PROMISE remains reasonable and performing such a tuning by hand would have been much more time consuming.

With both the sine neural network and MNIST neural network, PROMISE execution time tends to increase with the accuracy requirement. This can be explained by the Delta-Debug algorithm in PROMISE. As previously described, PROMISE firstly checks whether the accuracy requirement can be satisfied with double precision. Then, owing to the Delta-Debug algorithm, PROMISE tries to lower the precision of most variables from double to single, and this can be very fast if single precision is enough to match the required accuracy. Finally, PROMISE tries to transform the single precision declarations into half precision ones, and this can be fast if half precision is suitable for all these declarations. The number of programs compiled and executed by PROMISE tends to increase with the required accuracy on the results. For instance, in the case of MNIST neural network, if 1 or 2 digits are required, 18 type configurations are tested by PROMISE, whereas if 7 digits are required, 260 configurations are tested.

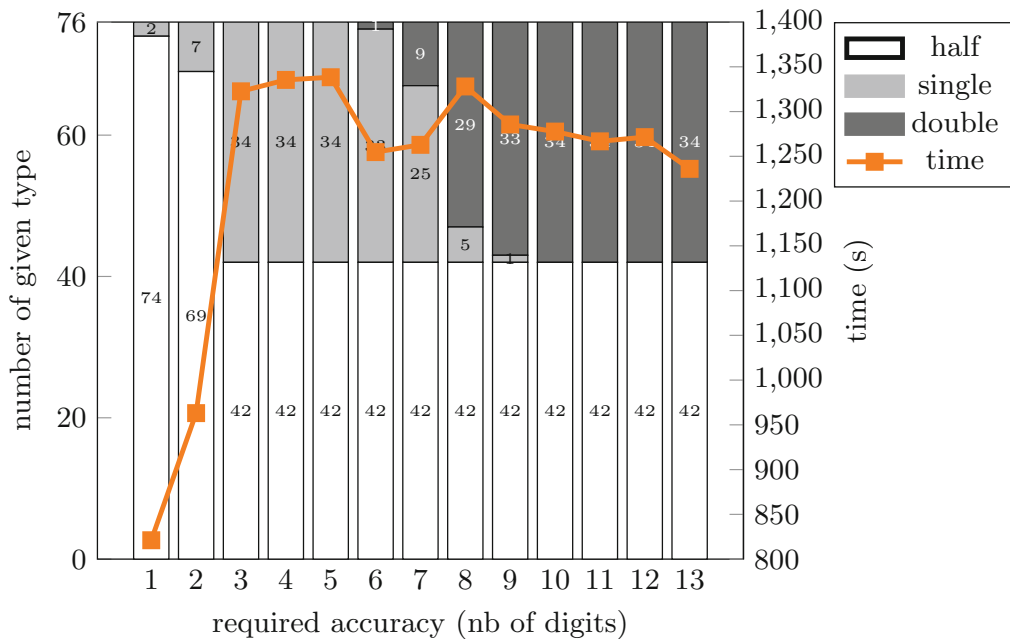


Fig. 7. Number of variables of each type and computation time for MNIST neural network with test_data[61] input

Results obtained considering one precision per layer are presented in Appendix B.1. The analysis is similar to the one previously given for the sine neural network. With the approach per layer, the execution time of PROMISE is lower than with the approach per neuron. But this approach forces some variables to be declared in higher precision. It can be noticed that with the approach per layer, both layers have the same precision. All the network parameters share the same type. Appendix B.2 and Appendix B.3 present the results obtained with another input. Like with the sine neural network, changing the input induces slight changes in the type configurations provided by PROMISE. However, the same trend can be observed.

4.3 CIFAR Neural Network

The neural network considered here is also an image classifier, but this time processing the CIFAR10 dataset. CIFAR¹⁰ is a dataset having 100 classes of colored images and the CIFAR10 dataset is reduced to 10 classes. Because images are of size $32 \times 32 \times 3$ (32 pixels width, 32 pixels height, 3 color channels), the network input is a 3D tensor of shape (32, 32, 3). The neural network consists of 5 layers: a convolutional layer having 32 neurons with activation function ReLU, followed by a max pooling of size 2×2 , a convolutional layer having 64 neurons with activation function ReLU, a flatten layer, and finally a dense layer of 10 neurons with activation function softmax. Taking into account one type per neuron and one type for each layer output, 111 types can be set.

Results presented here have been obtained on a 2.80 GHz Intel Core i9-10900 CPU having 20 cores and 64GB RAM. The maximum possible accuracy on the results is 13 digits. Although such a high accuracy is not necessary in a classifier network, exhaustive tests have been performed, like for the MNIST neural network. Results reported here refer to two input images out of the 10,000 provided by CIFAR10. Figure 8 and Appendix C.2 present the type configurations given by PROMISE with respectively `test_data[386]` and `test_data[731]`, considering one type per neuron. Appendix C.1 and Appendix C.3 show the results obtained with the same input images considering one type per layer.

PROMISE computation time tends to increase with the required accuracy. As already mentioned in Sect. 4.2, this can be explained by the Delta Debug algorithm. As previously observed, considering one type per layer results in lower PROMISE computation times, but often in uniform precision programs. Again, the input image slightly impacts the type configurations provided by PROMISE and the same trend can be observed.

¹⁰ <https://www.cs.toronto.edu/~kriz/cifar.html>.

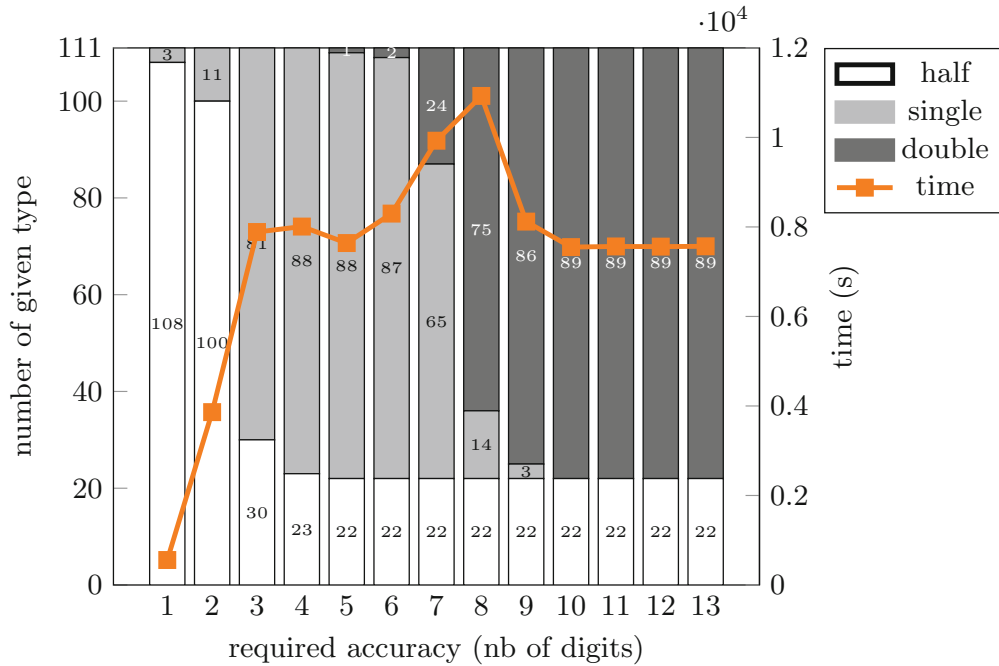


Fig. 8. Number of variables of each type and computation time for CIFAR neural network with test_data[386] input

Experiments have been carried out with neural networks also processing CIFAR10, but with more layers (up to 8 layers). Again, PROMISE could provide suitable type configurations taking into account accuracy requirements. However, PROMISE execution (that includes the compilation and execution of various programs) makes exhaustive tests more difficult with such neural networks. Possible PROMISE improvements described in Sect. 5 would enable precision tuning in larger neural networks that are themselves time consuming.

4.4 Inverted Pendulum

We present here results obtained with a neural network introduced in [4] in the context of reinforcement learning for autonomous control. In [4] methods are proposed for certified approximation of neural controllers and this neural network, related to an inverted pendulum, is used in a program that provides an approximation of a Lyapunov function. This neural network consists of 2 dense layers and uses the tanh activation function. The input is a state vector $x \in \mathbb{R}^2$ and the output, a scalar value in \mathbb{R} is an approximated value of the Lyapunov function for the state vector input. The first layer has 6 neurons and the second layer only one. Given the proximity between the two neural network models, results are expected to be close to the ones obtained for the sine approximation.

Figure 9 presents both the distribution of the different precisions and the execution time of PROMISE with respect to the accuracy requirement for input (0.5, 0.5). We consider here one type per neuron. As expected, the trend observed for the type configurations is the same as with the sine approximation. As

the required accuracy increases, the precision of the network parameters also increases. If one digit is required, all the parameters can be declared in half precision, and if at least 11 digits are required all the parameters must be in double precision. The computation time remains reasonable whatever the required accuracy. As previously observed, the computation time tends to increase with the required accuracy because of the number of type configurations tested by PROMISE.




Fig. 9. Number of variables of each type and computation time for the pendulum neural network with input (0.5, 0.5)

Appendix D.1 presents the results obtained considering one type per layer. Again, with this approach, the execution time of PROMISE is lower, but most configurations are actually in uniform precision. Results in Appendix D.2 and Appendix D.3 refer to an input consisting of two negative values ($-3, -6$). As previously observed, changing the input induces no significant difference.

5 Conclusion and Perspectives

We have shown with different kinds of neural networks having different types of layers how to lower the precision of their parameters while still satisfying a desired accuracy. Considering one type per neuron, mixed precision programs can be provided by PROMISE. Considering one type per layer enables one to reduce PROMISE execution time, however this approach often leads to uniform precision programs. It has been observed that with both approaches input values have actually a low impact on the type configurations obtained.

We plan to analyse the execution time of the mixed precision programs obtained with PROMISE on a processor with native half precision. Other perspectives consist in improving PROMISE. Another accuracy test more adapted to image classification networks could be proposed. We could improve the Delta-Debug algorithm used in PROMISE. Optimizations of the Delta-Debug algorithm are described in [19], including the parallelization potential. We could also consider the parallelization of PROMISE itself, i.e. applying PROMISE to different parts of a code in parallel. The extension of PROMISE and CADNA to other floating-point formats such as bfloat16 is another perspective. Taking benefit from the GPU version of CADNA, PROMISE could also be extended to GPUs. Floating-point auto-tuning in arbitrary precision is also a possible perspective that would enable the automatic generation of programs with a suitable type configuration for architectures such as FPGAs.

Acknowledgements. This work was supported by the InterFLOP (ANR-20-CE46-0009) project of the French National Agency for Research (ANR).

Appendices

Appendix A Sine Neural Network

Appendix A.1 Type Distribution for Sine Approximation with Input Value 2.37 Considering One Type per Neuron

See Fig. 10.




Fig. 10. Number of variables of each type and computation time for the sine neural network with input value 2.37

Appendix A.2 Type Distribution for Sine Approximation with Input Value 2.37 Considering One Type per Layer

See Fig. 11.




Fig. 11. Precision of each layer for the sine neural network with input value 2.37

Appendix B MNIST Neural Network

Appendix B.1 Type Distribution for MNIST with Test_data[61] Input Considering One Type per Layer

See Fig. 12.

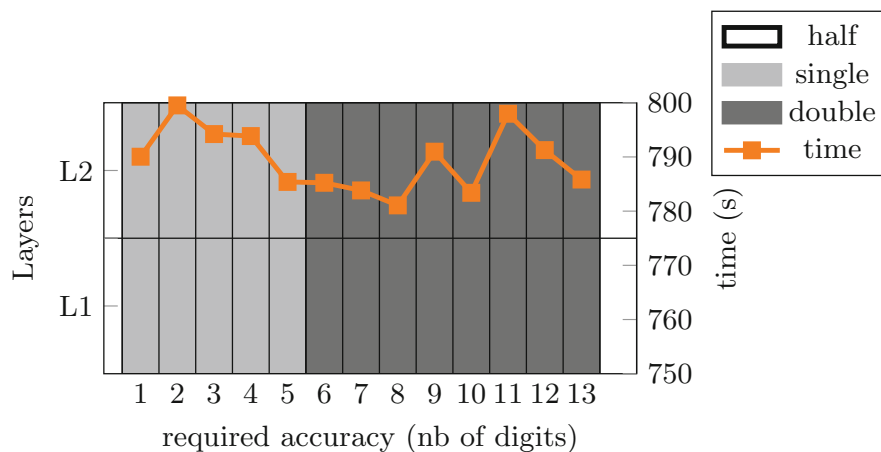


Fig. 12. Precision of each layer for MNIST neural network with test_data[61] input

Appendix B.2 Type Distribution for MNIST with Test_data[91] Input Considering one type per neuron

See Fig. 13.




Fig. 13. Number of variables of each type and computation time for MNIST neural network with test_data[91] input

Appendix B.3 Type Distribution for MNIST with Test_data[91] Input Considering One Type per Layer

See Fig. 14.




Fig. 14. Precision of each layer for MNIST neural network with test_data[91] input

Appendix C CIFAR Neural Network

Appendix C.1 Type Distribution for CIFAR with Test_data[386] Input Considering One Type per Layer

See Fig. 15.




Fig. 15. Precision of each layer for CIFAR neural network with test_data[386] input

Appendix C.2 Type Distribution for CIFAR with Test_data[731] Input Considering One Type per Neuron

See Fig. 16.




Fig. 16. Number of variables of each type and computation time for CIFAR neural network with test_data[731] input

Appendix C.3 Type Distribution for CIFAR with Test_data[731] Input Considering One Type per Layer

See Fig. 17.




Fig. 17. Precision of each layer for CIFAR neural network with test_data[731] input

Appendix D Inverted pendulum

Appendix D.1 Type Distribution for the Inverted Pendulum with Input (0.5,0.5) Considering One Type per Layer

See Fig. 18.




Fig. 18. Precision of each layer for the pendulum neural network with input (0.5, 0.5)

Appendix D.2 Type Distribution for the Inverted Pendulum with Input $(-3, -6)$ Considering One Type per Neuron

See Fig. 19.




Fig. 19. Number of variables of each type and computation time for the pendulum neural network with input $(-3, -6)$

Appendix D.3 Type Distribution for the Inverted Pendulum with Input $(-3, -6)$ Considering One Type per Layer

See Fig. 20.




Fig. 20. Precision of each layer for the pendulum neural network with input $(-3, -6)$

References

1. IEEE Standard for Floating-Point Arithmetic, IEEE Std 754-2019 (Revision of IEEE 754-2008), pp. 1–84 (2019)
2. Adjé, A., Khalifa, D.B., Martel, M.: Fast and efficient bit-level precision tuning. [arXiv:2103.05241](https://arxiv.org/abs/2103.05241) [cs] (2021)
3. Benmagnhia, H., Martel, M., Seladji, Y.: Fixed-point code synthesis for neural networks. *Artificial Intelligence, Soft Computing and Applications*, pp. 11–30 (2022). [arXiv:2202.02095](https://arxiv.org/abs/2202.02095)
4. Chang, Y.-C., Roohi, N., Gao, S.: Neural Lyapunov control. In: 33rd Conference on Neural Information Processing Systems (NeurIPS 2019) (2020). [arXiv:2005.00611](https://arxiv.org/abs/2005.00611)
5. Cherubin, S., Cattaneo, D., Chiari, M., Bello, A.D., Agosta, G.: TAFFO: tuning assistant for floating to fixed point optimization. *IEEE Embed. Syst. Lett.* **12**, 5–8 (2020)
6. Chesneaux, J.-M.: L’arithmétique stochastique et le logiciel CADNA, Habilitation à diriger des recherches. Université Pierre et Marie Curie, Paris, France (1995)
7. Chesneaux, J.-M., Graillat, S., Jezequel, F.: Numerical validation and assessment of numerical accuracy. Oxford e-Research Center (2009)
8. Chiang, W.-F., Baranowski, M., Briggs, I., Solovyev, A., Gopalakrishnan, G., Rakamarić, Z.: Rigorous floating-point mixed-precision tuning. In: Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, pp. 300–315. ACM, New York (2017)
9. Csendes, T.: Adversarial example free zones for specific inputs and neural networks. In: Proceedings of ICAI, pp. 76–84 (2020)
10. Damouche, N., Martel, M.: Mixed precision tuning with Salsa. In: Proceedings of the 8th International Joint Conference on Pervasive and Embedded Computing and Communication Systems, Porto, Portugal, pp. 47–56. SCITEPRESS - Science and Technology Publications (2018)
11. Darulova, E., Izycheva, A., Nasir, F., Ritter, F., Becker, H., Bastian, R.: Daisy - framework for analysis and optimization of numerical programs (tool paper). In: Beyer, D., Huisman, M. (eds.) TACAS 2018. LNCS, vol. 10805, pp. 270–287. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-89960-2_15
12. Darulova, E., Kuncak, V.: Towards a compiler for reals. *ACM Trans. Programm. Lang. Syst. (TOPLAS)* **39**, 8:1–8:28 (2017)
13. Dutta, S., Jha, S., Sankaranarayanan, S., Tiwari, A.: Output range analysis for deep feedforward neural networks. In: Dutle, A., Muñoz, C., Narkawicz, A. (eds.) NFM 2018. LNCS, vol. 10811, pp. 121–138. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-77935-5_9
14. Essam, M., Tang, T.B., Ho, E.T.W., Chen, H.: Dynamic point stochastic rounding algorithm for limited precision arithmetic in deep belief network training. In: 2017 8th International IEEE/EMBS Conference on Neural Engineering (NER), Shanghai, China, pp. 629–632. IEEE, May 2017
15. Graillat, S., Jézéquel, F., Picot, R., Févotte, F., Lathuilière, B.: Auto-tuning for floating-point precision with discrete stochastic arithmetic. *J. Comput. Sci.* **36**, 101017 (2019)
16. Guo, H., Rubio-González, C.: Exploiting community structure for floating-point precision tuning. In: Proceedings of the 27th ACM SIGSOFT International Symposium on Software Testing and Analysis, Amsterdam Netherlands, pp. 333–343. ACM, July 2018

17. Gupta, S., Agrawal, A., Gopalakrishnan, K., Narayanan, P.: Deep learning with limited numerical precision. In: Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37, ICML 2015 (2015). [arXiv:1502.02551](https://arxiv.org/abs/1502.02551)
18. Ho, N.-M., Silva, H.D., Wong, W.-F.: GRAM: a framework for dynamically mixing precisions in GPU applications. *ACM Trans. Archit. Code Optim.* **18**, 1–24 (2021)
19. Hodovan, R., Kiss, .A.: Practical improvements to the minimizing delta debugging algorithm. In: Proceedings of the 11th International Joint Conference on Software Technologies, Lisbon, Portugal, pp. 241–248. SCITEPRESS - Science and Technology Publications (2016)
20. Ioualalen, A., Martel, M.: Neural network precision tuning. In: Parker, D., Wolf, V. (eds.) QEST 2019. LNCS, vol. 11785, pp. 129–143. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-30281-8_8
21. Jezequel, F., Hoseininasab, S., Hilaire, T.: Numerical validation of half precision simulations. In: Rocha, .A., Adeli, H., Dzemyda, G., Moreira, F., Ramalho Correia, A.M. (eds.) WorldCIST 2021. AISC, vol. 1368, pp. 298–307. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-72654-6_29
22. Kotipalli, P.V., Singh, R., Wood, P., Laguna, I., Bagchi, S.: AMPT-GA: automatic mixed precision floating point tuning for GPU applications. In: Proceedings of the ACM International Conference on Supercomputing, Phoenix, Arizona, pp. 160–170. ACM, June 2019
23. Laguna, I., Wood, P.C., Singh, R., Bagchi, S.: GPUMixer: performance-driven floating-point tuning for GPU scientific applications. In: Weiland, M., Juckeland, G., Trinitis, C., Sadayappan, P. (eds.) ISC High Performance 2019. LNCS, vol. 11501, pp. 227–246. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-20656-7_12
24. Lam, M.O., Hollingsworth, J.K., de Supinski, B.R., Legendre, M.P.: Automatically adapting programs for mixed-precision floating-point computation. In: Proceedings of the 27th International ACM Conference on International Conference on Supercomputing, ICS 2013, pp. 369–378. ACM, New York (2013)
25. Lam, M.O., Vanderbruggen, T., Menon, H., Schordan, M.: Tool integration for source-level mixed precision. In: 2019 IEEE/ACM 3rd International Workshop on Software Correctness for HPC Applications (Correctness), pp. 27–35 (2019)
26. Lauter, C., Volkova, A.: A framework for semi-automatic precision and accuracy analysis for fast and rigorous deep learning. [arXiv:2002.03869](https://arxiv.org/abs/2002.03869) [cs] (2020)
27. Lin, W., et al.: Robustness verification of classification deep neural networks via linear programming. In: Conference on Computer Vision and Pattern Recognition (2019)
28. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks. In: 6th International Conference on Learning Representations, ICLR (2019). [arXiv:1706.06083](https://arxiv.org/abs/1706.06083)
29. Mellempudi, N., Srinivasan, S., Das, D., Kaul, B.: Mixed precision training with 8-bit floating point. [arXiv:1905.12334](https://arxiv.org/abs/1905.12334) [cs, stat] (2019)
30. Menon, H., et al.: ADAPT: algorithmic differentiation applied to floating-point precision tuning. In: SC18: International Conference for High Performance Computing, Networking, Storage and Analysis, Dallas, TX, USA, pp. 614–626. IEEE, November 2018
31. Na, T., Ko, J.H., Kung, J., Mukhopadhyay, S.: On-chip training of recurrent neural networks with limited numerical precision. In: 2017 International Joint Conference on Neural Networks (IJCNN), Anchorage, AK, USA, pp. 3716–3723. IEEE, May 2017

32. Rakin, A.S. et al.: RA-BNN: constructing robust & accurate binary neural network to simultaneously defend adversarial bit-flip attack and improve accuracy. [arXiv:2103.13813](https://arxiv.org/abs/2103.13813) [cs, eess] (2021)
33. Rubio-González, C. et al.: Precimonious: tuning assistant for floating-point precision. In: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, SC 2013, pp. 27:1–27:12. ACM, New York (2013)
34. Singh, G., Gehr, T., Mirman, M., Püschel, M., Vechev, M.: Fast and effective robustness certification. In: Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems, NeurIPS, pp. 10825–10836 (2018)
35. Su, C., Zhou, S., Feng, L., Zhang, W.: Towards high performance low bitwidth training for deep neural networks, *J. Semicond.* **41**, 022404 (2020). <https://iopscience.iop.org/article/10.1088/1674-4926/41/2/022404>
36. Tjeng, V., Xiao, K., Tedrake, R.: Evaluating robustness of neural networks with mixed integer programming. [arXiv:1711.07356](https://arxiv.org/abs/1711.07356) [cs] (2019)
37. Vignes, J.: A stochastic arithmetic for reliable scientific computation. *Math. Comput. Simul.* **35**, 233–261 (1993)
38. Vignes, J.: Discrete stochastic arithmetic for validating results of numerical software. *Numer. Algorithms* **37**, 377–390 (2004)
39. Wang, N., Choi, J., Brand, D., Chen, C.-Y., Gopalakrishnan, K.: Training deep neural networks with 8-bit floating point numbers. In: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) Advances in Neural Information Processing Systems 31, pp. 7686–7695. Curran Associates Inc. (2018). [arXiv:1812.08011](https://arxiv.org/abs/1812.08011). <http://papers.nips.cc/paper/7994-training-deep-neural-networks-with-8-bit-floating-point-numbers.pdf>
40. Zeller, A., Hildebrandt, R.: Simplifying and isolating failure-inducing input. *IEEE Trans. Softw. Eng.* **28**, 183–200 (2002)
41. Zombori, D.: Verification of artificial neural networks via MIPVerify and SCIP, SCAN (2020)
42. Zombori, D., Bánhelyi, B., Csendes, T., Megyeri, I., Jelasity, M.: Fooling a complete neural network verifier. In: The 9th International Conference on Learning Representations (ICLR) (2021)