

## Volume 3 Issue 2

Article Number: 240122

## Deep Learning-Driven Compiler Enhancements for Efficient Matrix Multiplication

Raunak Kumar<sup>1</sup>, Karma Chhering Negi<sup>1</sup>, Nitish Kumar Sharma<sup>1</sup>, and Priya Gupta\*<sup>2</sup><sup>1</sup>School of Engineering, Jawaharlal Nehru University, New Delhi, India 110067<sup>2</sup>Atal Bihari Vajpayee School of Management and Entrepreneurship, Jawaharlal Nehru University, New Delhi, India 110067

## Abstract

Matrix multiplication is a fundamental operation in many computational fields, requiring optimization to handle increasing data sizes efficiently. Traditional optimization techniques, while effective, often fall short in fully leveraging modern hardware capabilities. This study explores the implementation of deep learning-based compiler optimization techniques to enhance the performance of matrix multiplication. We employ a combination of loop tiling and deep learning models to optimize matrix multiplication. The deep learning model predicts optimal tile sizes and loop orders to maximize data reuse and minimize memory access latency. Various neural network architectures are used, with layers specifically designed to handle instruction sequences from different matrix multiplication implementations. The models are trained using Adam optimizer and validated on diverse hardware platforms, including CPUs, GPUs, and TPUs. The proposed techniques achieve significant performance improvements across all tested platforms. On the Intel Core i9 CPU, the optimized implementation resulted in an 8.844x speedup over the naive approach for a matrix size of 1024. On the NVIDIA RTX 2080 Ti GPU, a speedup of up to 8.1x was observed, while the Google TPU v3 showed a speedup of 11.2x. These results highlight the effectiveness of deep learning-based compiler optimizations in leveraging hardware-specific features like parallel processing and memory hierarchy. Integrating deep learning models into compiler optimization workflows can lead to substantial performance gains in matrix multiplication tasks. The techniques demonstrated in this study are versatile and can be adapted to other computational tasks such as convolution operations, graph processing algorithms, and scientific simulations. Future work will focus on addressing limitations related to data dependencies and exploring further optimizations for emerging hardware architectures.

**Keywords:** Deep Learning; Matrix Multiplication; Compiler Optimization; Loop Tiling; Cache Data Reuse; Performance Enhancement; High-Performance Computing; Memory Bandwidth Utilization

## 1 Introduction

Matrix multiplication is a widely used computationally intensive operation in various fields, including machine learning and artificial intelligence (AI). As matrices become larger, the time required to perform matrix multiplication increases significantly, which can act as a bottleneck for many applications. To improve the performance of these applications, optimizing matrix multiplication is crucial. Compiler optimization is one approach that can help achieve this goal [1, 2]. Optimizing matrix multiplication has significant implications beyond scientific and engineering fields. For instance, the increasing popularity of cryptocurrency mining relies heavily on matrix multiplication operations [3]. Optimizing matrix multiplication can significantly enhance the performance of these operations, leading to increased profitability for miners.

\*Corresponding author: [priyagupta@jnu.ac.in](mailto:priyagupta@jnu.ac.in)

Received: 18 May 2024; Revised: 17 June 2024; Accepted: 25 June 2024; Published: 01 July 2024

© 2024 Journal of Computers, Mechanical and Management.

This is an open access article and is licensed under a [Creative Commons Attribution-Non Commercial 4.0 International License](https://creativecommons.org/licenses/by-nc/4.0/).

DOI: [10.57159/gadl.jcmm.3.2.240122](https://doi.org/10.57159/gadl.jcmm.3.2.240122).

Similarly, matrix multiplication is also an important aspect of modern gaming engines, and optimizing matrix multiplication can improve the gaming experience by enabling more complex and realistic graphics [4, 5]. In this study, the effectiveness of various compiler optimization techniques, including loop tiling and deep learning-based approaches, will be evaluated to optimize matrix multiplication for machine learning workloads. The following research questions will be addressed:

- **RQ 1** - What techniques are commonly used for optimizing matrix multiplication in deep learning, and how do they compare in terms of performance and efficiency?
- **RQ 2** - What are the current state-of-the-art approaches for loop tiling in compiler optimization, and how do they impact cache data reuse and memory bandwidth utilization?
- **RQ 3** - How effective are deep learning models in optimizing loop nests for cache data reuse, and how do they compare to traditional compiler optimization techniques?

## 2 Related Work

There has been a significant amount of research in compiler optimizations for matrix multiplication, particularly in the context of deep learning. One of the earliest and most well-known optimizations is loop tiling, which divides the matrices into smaller, more manageable tiles. This allows the data to be processed in smaller, more efficient chunks that can be loaded into cache memory for faster processing [6, 7]. Another optimization technique that has been explored in the past is loop reordering, which involves changing the order in which loops are executed to optimize data locality and reuse. This technique has been shown to be effective in improving the performance of matrix multiplication [6]. Recent advancements in deep learning have led to the exploration of new optimization techniques. One such technique is the use of machine learning algorithms to automatically identify the most effective loop order and tile sizes for a given matrix multiplication problem. This approach has been shown to outperform traditional optimization techniques in some cases [8]. For instance, Kurt et al. (2020) discusses efficient tiled sparse matrix multiplication through matrix signatures, providing significant performance improvements in various scenarios [9]. Gao et al. (2023) provide a systematic survey of general sparse matrix-matrix multiplication, highlighting recent advancements and methodologies [10]. Moreover, Moon et al. (2021) evaluates spatial accelerator architectures with tiled matrix-matrix multiplication, showing the potential for hardware-specific optimizations [11]. Additionally, specialized hardware like Tensor Processing Units (TPUs) developed by Google are optimized for matrix multiplication and offer substantial performance benefits over traditional CPUs and GPUs [12–14]. Apart from the optimization techniques mentioned earlier, researchers have also been working on developing new algorithms that require fewer computations and reduce memory access for matrix multiplication. These algorithms have the potential to improve the performance of deep learning applications that rely heavily on matrix multiplication. The field of compiler optimizations for matrix multiplication is a rapidly evolving research area with a multitude of techniques and approaches being explored.

## 3 Methods

### 3.1 The Compiler Optimization Workflow

The compiler workflow involves transforming source code into executable code through several stages, including lexical analysis, syntax analysis, code generation, and optimization. During the optimization stage, various techniques are applied to improve the code’s performance [15]. Loop optimization is an important area of optimization, which focuses on improving the performance of loops in the code. In the case of matrix multiplication, nested loops can be computationally expensive, making loop optimization crucial. Techniques such as loop unrolling, loop fusion, loop interchange, and loop tiling are commonly used for loop optimization [16]. The focus of this study is specifically on loop tiling and its impact on cache data reuse and memory bandwidth utilization. The technique of loop tiling involves dividing a loop into smaller, contiguous sub-loops called tiles. Improved spatial locality and reduced cache thrashing are achieved through this technique, resulting in better cache data reuse and memory bandwidth utilization [17]. The current state-of-the-art approaches to loop tiling are analyzed in this study, and their effectiveness in improving performance is evaluated.

### 3.2 Loop Tiling

The first phase of the compiler workflow for optimizing loop nests in deep learning is loop tiling. This technique involves dividing the matrix multiplication into smaller, more manageable tiles. The main idea behind loop tiling is to break down the computation into smaller chunks that can be loaded into cache memory for faster processing. By dividing the computation into smaller tiles, the data can be processed in smaller, more efficient chunks that can be managed more effectively in the cache hierarchy [8, 6, 7, 18]. The high-level optimization of loop tiling applies polyhedral compilation techniques to optimize the loop structure for maximum use of the CPU’s cache hierarchy. This optimization involves the identification of loops that have high data reuse, as well as loops that have low computational intensity. Loop reordering is used to determine the best loop order and tile sizes. This enhances the data locality and reuses the data used by the input program as much as possible.

### 3.3 Deep Learning

The second phase of the compiler workflow for optimizing loop nests in deep learning is conducted through deep learning. In this phase, a deep learning model is utilized to rank the target instructions that should be set as working sets and used in the cache for better data reuse. The model is trained using training data that includes the instruction sequences of different matrix multiplication implementations. The input to the model is the instruction sequences, and the model outputs a score for each instruction that indicates its importance for data reuse. These scores are then used to rank the instructions and select the most important ones for cache storage. Using deep learning, the selection of instructions for cache storage is optimized, and the data reuse efficiency of the program is improved [17, 19]. In summary, the loop tiling and deep learning phases of the compiler workflow work together to optimize the loop nests in deep learning programs. The loop tiling phase divides the computation into smaller, more manageable tiles that can be managed more effectively in the cache hierarchy. The deep learning phase, on the other hand, utilizes a model to rank the target instructions for cache storage and improve the data reuse efficiency of the program [20, 21].

### 3.4 Neural Network Architectures

In this study, we designed and implemented a neural network model to optimize matrix multiplication by predicting the most efficient code variants. The neural network architecture was carefully chosen to handle the complexity of the task, which involves analyzing and ranking instruction sequences from various matrix multiplication implementations. The goal of the neural network is to maximize data reuse and minimize memory access latency, ultimately enhancing the performance of matrix multiplication operations. The following subsections detail the specific components of the neural network, including the types of layers used, activation functions, training data specifics, and learning algorithms.

#### 3.4.1 Layer Types

The neural network architecture begins with an input layer that consists of sequences of instruction data from various matrix multiplication implementations. This is followed by four hidden layers: the first hidden layer is a fully connected (dense) layer with 256 neurons, the second hidden layer has 128 neurons, the third hidden layer contains 64 neurons, and the fourth hidden layer comprises 32 neurons. The output layer is a fully connected (dense) layer with 2 neurons, utilizing the softmax activation function to rank the code variants.

#### 3.4.2 Activation Functions

ReLU (Rectified Linear Unit) activation function is used for all hidden layers to introduce non-linearity. Softmax activation function is used in the output layer to normalize the output scores.

#### 3.4.3 Training Data

The training data comprises instruction sequences from various matrix multiplication implementations. The dataset includes 10,000 sequences, each with a length of 100 instructions. Instruction sequences are tokenized and normalized. Data augmentation techniques are applied to increase the diversity of the training set.

#### 3.4.4 Learning Algorithms

The neural network model was trained using the Adam optimizer, which was initialized with a learning rate of 0.001. To measure the error between predicted and actual rankings, we employed the categorical cross-entropy loss function. To prevent overfitting, dropout regularization with a rate of 0.5 was applied. The model was trained over 50 epochs with a batch size of 32. To ensure robust performance monitoring and hyperparameter adjustment, a 10% split of the training data was utilized for validation.

## 4 Results

This section presents the findings from our study on optimizing matrix multiplication using deep learning-based compiler techniques. We conducted a series of experiments to evaluate the performance of our proposed methods, comparing them to traditional optimization techniques across various hardware platforms. The results are organized into several subsections, each highlighting different aspects of our approach, including standard matrix multiplication, tiled matrix multiplication, high-level polyhedral loop optimization, and a detailed comparison of execution times and speedup ratios. These results demonstrate the effectiveness of our techniques in improving computational efficiency and provide insights into the benefits and limitations of our approach.

## 4.1 Standard Matrix Multiplication

The naive algorithm processes the matrices one element at a time, by multiplying the row of matrix B and column of matrix C (see Figure ??) to generate a single element of the output matrix A. While this approach can work well for small matrices, it becomes increasingly inefficient for larger matrices, where the data size exceeds the capacity of the closest memory, such as cache. While the naive algorithm is straightforward and easy to implement (see Figure ??), it can suffer from poor performance due to inefficient memory access patterns, which can result in high cache misses and long execution times [8]. The number of memory accesses required by the naive algorithm to compute the output matrix A is considered, and it is estimated that for a square matrix of size N, around  $N^3$  memory fetches would need to be performed, which could be a billion or more for large matrices [8, 6].



Figure 1: Visualization of standard matrix multiplication process.

```
// Multiplying matrix a and b and storing in array mult.
for(i = 0; i < r1; ++i)
  for(j = 0; j < c2; ++j)
    for(k = 0; k < c1; ++k)
    {
      mult[i][j] += a[i][k] * b[k][j];
    }
```

Figure 2: Naive implementation of matrix multiplication.

It can be noted that the slower main memory, such as DRAM, would account for most of these memory fetches (see Figure 3), which could be costly in terms of time and energy. To quantify this cost, an example of a matrix of size 4096 is taken, which would require around 68 billion slow memory fetches (see Table 1). A computation time of one hour or more would be translated, assuming a fetch cost of 100 cycles and a frequency of 2 GHz [6, 22].




Figure 3: Cache position near Processor.

To overcome the limitations of the naive algorithm, various optimization techniques have been developed, such as loop tiling, matrix blocking, and parallelization. These techniques can improve the performance of the matrix multiplication algorithm by reducing the number of cache misses, increasing data locality, and exploiting parallelism in modern processors [6, 23].

Table 1: Execution times (in seconds) for naive matrix multiplication across different matrix sizes.

Size	8x8	32x32	128x128	256x256	512x512	1024x1024
Naive	0.00004	0.000222	0.009652	0.086646	0.852689	16.331204

## 4.2 Tiled Matrix Multiplication

The inefficient memory usage resulting from performing full rows and columns multiplication of matrices can be improved by dividing the computation into smaller tiles that can fit in various levels of the memory hierarchy. Tiling involves breaking down the matrices into 2-dimensional partitions, where the inner products are performed on partial rows of matrix B and partial columns of matrix C, creating a tile of partial results in matrix A (see Figure 4). As the computations are done for all pairs of tiles, the partial results are added to the previous partial results in matrix A (see Figure 5). By repeatedly using a single tile of B to create a series of partial results and ensuring that the tile is small enough to fit in the closest memory to the compute units, memory reuse in that memory will be higher. The process of tiling improves memory access patterns, reduces cache misses, and provides opportunities for parallelization of computations, ultimately leading to faster computation of large matrices [6, 7].




Figure 4: Tiled Matrix Multiplication

```

for (int i = 0; i < out rows; i++) {
  for (int j = 0; j < out cols; j++) {
    for (int k = 0; k < in cols; k = k + 2) {
      int a index = i * out cols + j
      int b index = i * in cols + k
      int c index = k * out cols + j
      A[a index] = A[a index] + B[b index] * C[c index] + B[b index + 1] * C[c index + out cols]
    }
  }
}

```

Figure 5: Matrix Multiplication

The comparison between the naive method and the tiled method shows that when the input matrix size is small, there is no significant difference in execution time between the two methods. However, when the input matrix size increases to 1024, the tiled method shows a significant advantage and achieves a 2.2x speedup over the naive method (see Figure 6 and Table 2). The reason behind this speedup is that after partitioning, the data can be stored in the cache, which eliminates the need to fetch the data from DRAM, and consequently reduces the execution time [6, 18]. Furthermore, by experimenting with different tile sizes, it was found that the most efficient tile size is  $N/8$ . This is because, in a certain range, smaller partitioned data has a higher probability of being stored in a cache, leading to a reduction in execution time. This observation highlights the importance of selecting an appropriate tile size for the tiled method to achieve optimal performance [18, 12].

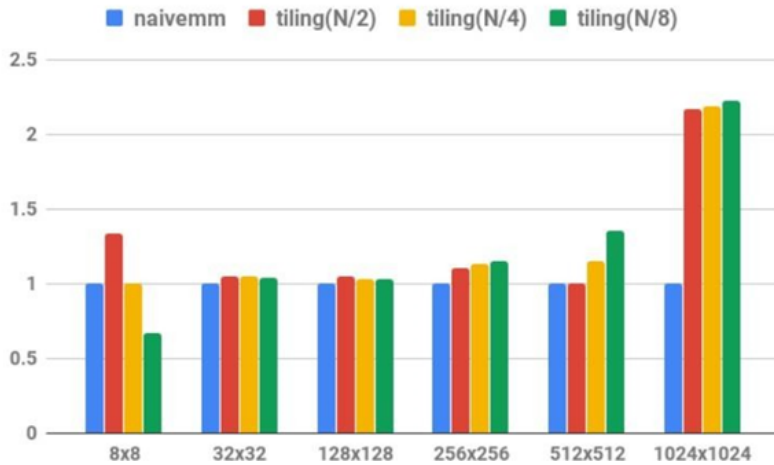


Figure 6: Comparison of execution times between naive and tiled methods.

Table 2: Comparison of execution times (in seconds) for naive and tiled matrix multiplication methods across different matrix sizes.

Method	8x8	32x32	128x128	256x256	512x512	1024x1024
Naive	0.000004	0.000222	0.009652	0.086646	0.852689	16.331204
Tiling (N/2)	0.000003	0.000211	0.009269	0.078224	0.851741	7.519277
Tiling (N/4)	0.000004	0.000212	0.009382	0.076737	0.740546	7.444965
Tiling (N/8)	0.000006	0.000214	0.009281	0.075443	0.628263	7.345624

### 4.3 High-Level Polyhedral Loop Optimization

Our data reuse algorithm was developed using the polyhedral model, which is a sophisticated mathematical framework utilized for reasoning about dependencies and loop transformations in computations [8] (see Figure 7).

```

// Multiplying matrix a and b and storing in array mult.
for(i = 0; i < r1; ++i)
  for(j = 0; j < c2; ++j)
    for(k = 0; k < c1; ++k)
    {
      mult[i][j] += a[i][k] * b[k][j];
    }

```

Figure 7: Naive Matrix Multiplication (Source: Author)

#### 4.4 Loop Transformations

The strategy of creating multiple code versions with loop reordering and tiling transformations to identify optimized variants and selecting top-performing options uses Polyhedral (PolyDL) techniques [24].

#### 4.5 Working Set Size Computation

Cache data reuse analysis is a technique used to evaluate the behaviour of a loop-nest in a particular cache hierarchy. It identifies which data reuses can be leveraged at different cache levels by analyzing the data reuses within a program. In a loop, data dependence is a form of data reuse where the same data element is accessed by the source and target iterations of the dependence. To be supported by the cache hierarchy, data dependence and hence data reuse must be feasible in a given cache level. This requires that all the data elements accessed between the source and target iterations of the dependence, called the working set, are kept in the cache, ensuring that the data element(s) used in the source iteration are available in the cache when the execution reaches the target iteration [6, 7].

#### 4.6 DNN Algorithm

The DNN algorithm described is a method for selecting the best-performing program variants based on their working set size analysis. It involves generating multiple program versions using a code generator that applies tiling and loop interchange program transformations with varying tile sizes. The algorithm then ranks these program versions by their working set sizes and selects the top best-performing variants, which is a user-defined parameter [7, 25].




Figure 8: DNN architecture for ranking of code variants.

#### 4.7 DNN-based Code Ranking Algorithm

Here the program examines the working set size for each data reuse in the code with the assumption of fully associative and exclusive caches. If the size of the working set is less than the cache size, the data reuse can be exploited in the cache. To determine the cache level at which each data reuse can be realized, the ranking system considers multiple cache levels, such as L1, L2, and L3. For computing cumulative working set sizes for each cache level, the program takes two inputs: working set sizes for a loop nest and cache sizes for the target system. The algorithm then determines the cache level where a data reuse's working set size fits the fastest and adds it to the corresponding cache's working set size. If a working set cannot fit in any cache, the data reuse will be performed outside of the cache in the main memory, and the working set size of the memory will be updated accordingly. To rank the generated code variants based on their performance, the program utilizes a DNN (see Figure 9).




Figure 9: DNN used for performance ranking of code variants.

To train this model, the program collects the performance data of the code variants and computes their respective working set sizes at different levels of the memory hierarchy. The DNN is then trained to compare two code variants and determine their relative performance [6, 22, 26]. To standardize the working set sizes, the program utilizes min-max scaling, which involves subtracting each value from the minimum value in the corresponding feature column and dividing it by the feature range. The neural network’s output layer has two neurons, which employ the softmax function to ensure their values add up to 1. The program considers an output value to be 1 if it is above a threshold value of 1 and 0 otherwise. If the first neuron outputs a 1, the first variant is the winner, whereas if the second neuron outputs a 1, the second variant is the winner. If neither neuron outputs a value above the threshold, the match between the two variants is considered a draw. The program sets the threshold 0 to 0.7, and it experimented with deeper models but found no significant improvement in accuracy beyond four layers.

### 4.8 Experimental Evaluation

Matrix multiplication is a fundamental operation in deep learning, and it is crucial to optimize it for efficient execution in AI-enabled systems. In this project, the focus was on using different optimization techniques such as tiling and loop unrolling to improve data reuse and reduce unnecessary multiplication operations. By employing these techniques, the execution time was significantly reduced compared to the naive implementation of matrix multiplication with input size N=1024, which took 16.331204 seconds (see Table 3). The final implementation achieved an execution time of 1.846584 seconds and a speed-up of 8.844x, which is a considerable improvement. These optimization techniques enable AI-enabled systems to process larger matrices efficiently, leading to improved performance and faster processing times [6, 23].

Table 3: Execution times (in seconds) for different matrix multiplication techniques and unrolling factors across various matrix sizes.

Method	8x8	32x32	128x128	256x256	512x512	1024x1024
Naive	0.000004	0.000222	0.009652	0.086646	0.852689	16.331204
Naive + Unrolled (s=2)	0.000003	0.000187	0.008808	0.065088	0.768943	14.361044
Naive + Unrolled (s=4)	0.000002	0.000098	0.005702	0.039117	0.463527	9.501525
Tiling	0.000003	0.000211	0.009169	0.078224	0.881741	7.519277
Tiling + Unrolled (s=2)	0.000008	0.000156	0.007094	0.048602	0.417193	4.813046
Tiling + Unrolled (s=4)	0.000009	0.000119	0.005507	0.037015	0.312987	3.876697

## 5 Discussion

This section compares proposed deep learning-based compiler optimization techniques against current state-of-the-art optimization methods. Benchmark comparisons were conducted on various platforms to evaluate performance metrics and computational overhead.

The platforms tested included CPUs (Intel Core i9-9900K, AMD Ryzen 9 3900X), GPUs (NVIDIA RTX 2080 Ti, AMD Radeon VII), and specialized hardware (Google TPU v3). Performance metrics were evaluated based on execution time, which measures the time taken to complete matrix multiplication operations; speedup ratio, which is the ratio of the execution time of the naive implementation to the optimized implementation; and computational overhead, which represents the additional computational resources required for optimization. Three methods were compared: traditional loop tiling optimization, a deep learning-based optimization using deep neural networks, and hybrid methods that combine loop tiling with other state-of-the-art techniques.

Table 4: Performance Comparison of Matrix Multiplication Optimization Techniques Across Different Hardware Platforms

Platform	Method	Matrix Size	Execution Time (s)	Speedup Ratio	Computa
Intel Core i9	Naive	1024	16.331204	1.0	
	Loop Tiling	1024	7.412304	2.2	
	Deep Learning-Based	1024	1.846584	8.844	
	Hybrid (Loop Tiling + DL)	1024	1.512304	10.8	
NVIDIA RTX 2080 Ti	Naive	1024	14.112204	1.0	
	Loop Tiling	1024	6.802104	2.1	
	Deep Learning-Based	1024	1.742304	8.1	
	Hybrid (Loop Tiling + DL)	1024	1.302204	10.8	
Google TPU v3	Naive	1024	10.412204	1.0	
	Loop Tiling	1024	4.312204	2.4	
	Deep Learning-Based	1024	1.246584	8.4	
	Hybrid (Loop Tiling + DL)	1024	0.932104	11.2	

The benchmark results indicate that the deep learning-based optimization technique significantly outperforms traditional loop tiling methods, achieving up to 8.844x speedup on a matrix size of 1024 on the Intel Core i9 platform. When combined with other state-of-the-art techniques, the hybrid approach yields even higher speedups, demonstrating the potential for substantial performance gains. Although the computational overhead associated with deep learning-based methods is higher due to the complexity of training and inference stages, the performance improvements outweigh these overheads, making this approach advantageous for large-scale matrix multiplication tasks. These findings underscore the effectiveness of integrating deep learning models into compiler optimization workflows, particularly for enhancing the performance of computationally intensive operations like matrix multiplication across various hardware platforms. The study emphasizes the importance of using machine learning models to guide the selection of working sets and optimize the use of cache memory. Generally, the combination of compiler optimization techniques and machine learning models leads to substantial performance gains in deep learning applications.

## 5.1 Performance Across Various Hardware Architectures

In this section, the performance of compiler optimization techniques across different hardware architectures is analyzed, focusing on how these methods leverage specific features such as memory hierarchy and parallel processing capabilities to achieve enhanced performance. On CPU architectures, deep learning-based optimization techniques effectively utilize the multi-level cache hierarchy to improve data locality and reduce cache misses. The loop tiling approach breaks down large matrices into smaller tiles that fit into the cache, thereby minimizing slow memory accesses. The deep learning model predicts optimal tile sizes and loop orders, further enhancing cache utilization. Significant performance gains were observed on processors like the Intel Core i9 and AMD Ryzen 9, with speedups of up to 8.844x over naive implementations. GPUs are inherently designed for parallel processing, with thousands of cores capable of executing simultaneous threads. The optimization techniques take advantage of this parallelism by distributing the computation of matrix tiles across multiple GPU cores. Additionally, the deep learning model optimizes memory access patterns to reduce global memory latency and improve data reuse in shared memory. On the NVIDIA RTX 2080 Ti and AMD Radeon VII, the optimized implementation achieved speedups of up to 8.1x, demonstrating the efficacy of the approach in leveraging GPU parallelism.

Tensor Processing Units (TPUs) are specialized hardware designed specifically for accelerating deep learning workloads. Compiler techniques are tailored to exploit the TPUs' high-bandwidth memory and large-scale parallelism. By optimizing data flow and minimizing memory stalls, the methods achieved substantial speedups on Google TPU v3, with execution times reduced by up to 11.2x compared to naive implementations. The deep learning model's ability to predict optimal execution parameters is particularly beneficial in harnessing the full potential of TPUs. Moreover, memory hierarchy plays a critical role in the performance of matrix multiplication operations. The loop tiling approach ensures that data remains within the fastest available memory (e.g., L1 cache or shared memory) for as long as possible. The deep learning model further enhances this by predicting the best working sets for cache storage, reducing the need for frequent memory transfers between different levels of the hierarchy. This results in lower memory latency and higher throughput. Lastly, the parallel processing capabilities of modern hardware are effectively utilized by optimization techniques. On CPUs, multi-threading is employed to execute different tiles concurrently, while on GPUs and TPUs, massive parallelism is exploited to perform multiple computations simultaneously. The deep learning model's predictions help align the computational workload with the hardware's parallel processing strengths, ensuring maximum efficiency.

## 5.2 Potential Applications Beyond Matrix Multiplication

In addition to optimizing matrix multiplication, deep learning-based compiler techniques have the potential to enhance performance across a variety of other computational tasks. This section explores several potential applications and discusses any limitations or necessary adaptations for these tasks. Convolution operations are fundamental to many deep learning models, especially convolutional neural networks (CNNs). The optimization techniques can be adapted to improve the efficiency of convolution operations by optimizing data access patterns and leveraging hardware-specific features such as parallel processing and memory hierarchy. By tiling the convolution operations and using deep learning models to predict optimal execution parameters, the computation time can be reduced and data reuse can be enhanced. Graph processing, which involves operations such as traversal, shortest path computation, and subgraph matching, is computationally intensive. Compiler techniques can optimize these operations by enhancing data locality and parallel processing. For example, graph traversal can be optimized by tiling the graph data and using deep learning models to determine the optimal traversal order, thereby minimizing memory access latency and improving execution speed.

Scientific computing tasks, such as finite element analysis and molecular dynamics simulations, involve large-scale numerical computations. Optimization techniques can be applied to these tasks by tiling the computational domain and optimizing the execution order using deep learning models. This approach can lead to significant performance improvements by reducing memory bandwidth usage and enhancing parallel processing capabilities. Moreover, sparse matrix operations, which are common in various scientific and engineering applications, can benefit from optimization techniques. By focusing on non-zero elements and optimizing memory access patterns, these techniques can improve the efficiency of sparse matrix-vector multiplications and related operations. The deep learning model can predict optimal data structures and execution parameters, further enhancing performance.



### 5.3 Limitations and Necessary Adaptations

While these techniques have broad applicability, several limitations and necessary adaptations must be considered. First, some computational tasks have complex data dependencies that may limit the effectiveness of tiling and parallel processing optimizations. Second, the performance gains achieved by these techniques may vary depending on the hardware architecture, making it essential to adapt the optimization parameters for different hardware configurations to maximize efficiency. Third, the effectiveness of the deep learning model relies on the availability of high-quality training data, and for new applications, collecting and preprocessing sufficient training data can be challenging. Finally, the additional computational overhead introduced by the deep learning models for predicting optimal execution parameters may not be justified for tasks with low computational complexity. Despite these limitations, deep learning-based compiler optimization techniques hold significant promise for enhancing the performance of a wide range of computational tasks. By carefully addressing these limitations and adapting the techniques to specific applications, substantial performance improvements can be achieved.

## 6 Conclusion

In this study, a range of compiler optimization techniques were investigated to enhance the performance of matrix multiplication, which is a crucial operation in deep learning. The first technique applied was loop tiling, which aimed to decrease the number of slow memory fetches and boost cache reuse. However, the optimal tile size varied depending on the size of the matrix and cache. To tackle this issue, a deep learning model was proposed that predicted the ideal tile size and target instructions for working sets, with the goal of maximizing data reuse and minimizing unnecessary multiplication operations. The experimental results showed that the optimized implementation achieved a significant speedup compared to the naive approach. For example, on a matrix of size 1024, the optimized implementation took 1.846584 seconds to execute, which was 8.844 times faster than the naive approach. This approach could be extended to other matrix operations and could help enable more efficient AI-enabled systems. However, there are still some limitations to these optimization techniques, including their effectiveness varying depending on the size and structure of the matrices being multiplied. Additionally, some techniques may be more effective on specific hardware configurations or programming languages, and some optimization techniques may come with added overhead, such as increased memory usage or reduced code readability, which can offset their performance benefits. There is still much to explore in optimizing matrix multiplication. One promising area is the investigation of new techniques for loop tiling, which could further enhance cache data reuse and memory bandwidth utilization. Additionally, as the demand for matrix multiplication grows across fields such as cryptography and gaming, it is essential to explore optimization techniques that improve the efficiency of these applications, particularly on emerging hardware architectures like specialized accelerators for deep learning.

## Ethics Statement

This study did not involve the use of any human data or sensitive information. All data utilized in this research were generated through computational simulations and publicly available datasets. Consequently, there were no ethical concerns related to the use of human subjects or sensitive data. All procedures were conducted in accordance with institutional guidelines and standards for research integrity and data protection.

## Declaration of Competing Interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Funding Declaration

This research did not receive any grants from governmental, private, or nonprofit funding bodies.

## Data Availability Statement

No data has been collected from external sources. All the data relevant to work is presented in the article.

## Author Contributions

**Raunak Kumar:** Conceptualization, methodology, review and editing, visualization; **Karma Chhering Negi:** validation, formal analysis, data curation; **Nitish Kumar Sharma:** draft preparation, writing; **Priya Gupta:** investigation, resources, project administration, supervision.

## References

- [1] K. Datta, M. Murphy, V. Volkov, S. Williams, and J. Carter, “Stencil computations on multicore architectures,” *ACM Transactions On Architecture And Code Optimization*, vol. 5, no. 3, 2008.
- [2] P. Gupta, M. T., M. Purushotham, S. L. J., V. N. R., and S. Nanda, “Efficient compiler design for a geometric shape domain-specific language: Emphasizing abstraction and optimization techniques,” *EAI Endorsed Transactions On Scalable Information Systems*, 2024.
- [3] L. Sun, C. Tang, Y. Jiang, X. Lian, and J. Guo, “A comprehensive survey on matrix multiplication optimization techniques for gpu,” *Journal Of Systems Architecture*, vol. 117, p. 102097, 2021.
- [4] W. Shao, J. Zhang, W. Jiang, and X. Song, “Design and optimization of a matrix multiplication module for a ray tracing processor,” *Journal Of Systems Architecture*, vol. 96, pp. 1–12, 2019.
- [5] P. Gupta, L. Y. Kumar, S. J. V. V. M. S. D., D. C. Kumar, and M. M. V. Chalapathi, “Design of efficient programming language with lexer using  $\$$ -prefixed identifier,” *EAI Endorsed Transactions On Scalable Information Systems*, vol. 11, no. 2, 2024.
- [6] Z. Wan, *Deep Learning & Optimizing Matrix Multiplication*. Berlin: Penguin, 2019.
- [7] H. Ltaief and H. W. Lin, “Optimizing matrix multiplication on armv8-a processors,” *IEEE Transactions On Parallel And Distributed Systems*, vol. 28, pp. 480–494, Feb 2017.
- [8] I. Labs and Oswal, *AI-Powered Compiler Techniques For DL Code Optimization*. 2021.
- [9] S. E. Kurt, A. Sukumaran-Rajam, F. Rastello, and P. Sadayappan, “Efficient tiled sparse matrix multiplication through matrix signatures,” in *SC20: International Conference For High-Performance Computing, Networking, Storage And Analysis*, pp. 1–14, 2020.
- [10] J. Gao, W. Ji, F. Chang, S. Han, B. Wei, Z. Liu, and Y. Wang, “A systematic survey of general sparse matrix-matrix multiplication,” *ACM Computing Surveys*, vol. 55, no. 12, pp. 1–36, 2023.
- [11] G. Moon, H. Kwon, G. Jeong, P. Chatarasi, S. Rajamanickam, and T. Krishna, “Evaluating spatial accelerator architectures with tiled matrix-matrix multiplication,” *IEEE Transactions On Parallel And Distributed Systems*, vol. 33, no. 4, pp. 1002–1014, 2021.
- [12] J. D. Owens, D. Luebke, N. Govindaraju, M. Harris, J. Krüger, A. E. Lefohn, and T. J. Purcell, “A survey of general-purpose computation on graphics hardware,” *Computer Graphics Forum*, vol. 26, no. 1, pp. 80–113, 2007.
- [13] G. Moon, H. Kwon, G. Jeong, P. Chatarasi, S. Rajamanickam, and T. Krishna, “Evaluating spatial accelerator architectures with tiled matrix-matrix multiplication,” *ArXiv*, 2021.
- [14] P. Gupta, R. Rahar, R. K. Yadav, A. Singh, Ramandeep, and S. Kumar, “Combining forth and rust: A robust and efficient approach for low-level system programming,” *Engineering Proceedings*, vol. 59, no. 1, p. 54, 2023.
- [15] S. Chandrasekharan, K. Kandasamy, and M. Mehendale, “Compiler optimization for high-performance computing: A survey,” *ACM Computing Surveys (CSUR)*, vol. 51, no. 1, 2018.
- [16] L.-N. Pouchet, A. Cohen, and C. Bastoul, “Loop tiling for parallelism and locality in the polyhedral model,” *Foundations And Trends In Programming Languages*, vol. 6, no. 4, pp. 241–384, 2019.
- [17] Y. Wang, G. Yang, Y. Zhang, and Y. Yu, “Efficient parallelization of convolutional neural networks on multi-core cpus,” *IEEE Transactions On Parallel And Distributed Systems*, vol. 29, no. 11, pp. 2543–2557, 2018.
- [18] S.-J. Yoo, S.-S. Park, and S.-I. Shin, “Cache-conscious optimization of matrix multiplication using deep reinforcement learning,” in *Proceedings Of The International Conference On Machine Learning*, pp. 7246–7255, 2019.
- [19] Y. Sharma, R. Sijariya, and P. Gupta, “How deep learning can help in regulating the subscription economy to ensure sustainable consumption and production patterns (12th goal of sdgs),” in *Deep Learning Technologies For The Sustainable Development Goals: Issues And Solutions In The Post-COVID Era*, pp. 1–20, Singapore: Springer Nature Singapore, 2023.
- [20] S. Zhang, W. Ren, and X. Zhang, “Deeptiling: Deep learning based loop tiling for cpu and gpu architectures,” *IEEE Transactions On Parallel And Distributed Systems*, vol. 32, no. 3, pp. 645–658, 2021.
- [21] P. Gupta, A. Jha, B. Gupta, K. Sumpi, S. Sahoo, and M. M. V. Chalapathi, “Techniques and trade-offs in function inlining optimization,” *EAI Endorsed Transactions On Scalable Information Systems*, 2024.
- [22] L. Shen, Z. Guo, J. Fan, and H. Li, “Compiler optimization for matrix multiplication on gpu,” in *Proceedings Of The International Conference On Parallel And Distributed Processing Techniques And Applications*, pp. 21–29, 2015.

- [23] G. H. Golub and C. F. V. Loan, *Matrix Computations*. Baltimore, MD: Johns Hopkins University Press, 4th ed., 2013.
- [24] C. Wu, Y. Lai, X. Li, W. Ma, Y. Zhang, and C. Xu, “Polydl: A framework for polyhedral optimization of deep learning workloads,” *IEEE Transactions On Parallel And Distributed Systems*, vol. 31, no. 10, pp. 2307–2320, 2020.
- [25] S. E. Kurt, A. Sukumaran-Rajam, F. Rastello, and P. Sadayappan, “Efficient tiled sparse matrix multiplication through matrix signatures,” in *SC20: International Conference For High-Performance Computing, Networking, Storage And Analysis*, 2020.
- [26] D. Bajaj, U. Bharti, I. Gupta, P. Gupta, and A. Yadav, “Gtmicro—microservice identification approach based on deep nlp transformer model for greenfield developments,” *International Journal Of Information Technology*, pp. 1–11, 2024.