Auto-SpMV: Automated Optimizing SpMV Kernels on GPU

Mina Ashoury
Department of Computer Science and Engineering, Shiraz University, Shiraz, Iran
minaashoury@gmail.com

Mohammad Loni
School of Innovation, Design and Engineering, Mälardalen University, Västerås, Sweden
mohammad.loni@mdu.se

Farshad Khunjush
Department of Computer Science and Engineering, Shiraz University, Shiraz, Iran
khunjush@shirazu.ac.ir

Masoud Daneshtalab
School of Innovation, Design and Engineering, Mälardalen University, Västerås, Sweden
masoud.daneshtalab@mdu.se

ABSTRACT

Sparse matrix-vector multiplication (SpMV) is an essential linear algebra operation that dominates the computing cost in many scientific applications. Due to providing massive parallelism and high memory bandwidth, GPUs are commonly used to accelerate SpMV kernels. Prior studies mainly focused on reducing the latency of SpMV kernels on GPU. However, few attempts have been made to improve the energy efficiency of SpMV kernels, resulting in GPUs being excluded from the range of low-power applications. Furthermore, prior work has primarily focused on optimizing the sparse format of SpMV kernels, the literature ignores evaluating the impact of tweaking compilation parameters. Lastly, Little attention has been paid to preparing a comprehensive training dataset of running SpMV kernels and fine-tuning the learning hyperparameters. To address these limitations, we present a novel framework, dubbed Auto-SpMV, that enables energy-efficient and low-latency SpMV kernels on GPU. To achieve the best run time performance, Auto-SpMV proposes two optimization modes: *compile-time* and *run-time*. In the *compile-time* mode, Auto-SpMV tweaks the compilation parameters, while in the *run-time* mode, Auto-SpMV selects the best sparse format for the input matrix. To achieve the best classification results, 1) we collect the largest dataset ever having 30 different sparse matrices running with more than 15K different configurations, and 2) we boost classification models by automatically fine-tuning the learning hyperparameters. Experimental results reveal that Auto-SpMV optimizes latency, energy consumption, average power, and energy efficiency in the *compile-time* mode by up to 51.9%, 52%, 33.2%, and 53%, respectively, compared to the default setting. Auto-SpMV optimizes average power and energy efficiency in the *run-time* mode by up to 34.6% and 99.7%, respectively, compared to the default setting.

Keywords Sparse Matrix-Vector Multiplication, GPU, Performance Modeling, Energy Efficiency

1 Introduction

Sparse matrix-vector multiplication (SpMV) is a key operation in a variety of scientific applications such as large simulation systems [35], graph learning [39], economic modeling [38], and more. It is extremely imperative to

optimize the SpMV kernel on modern hardware devices because SpMV dominates the computing cost in iterative problem-solving methods such as the calculation of Eigenvalues [36] and Krylov subspace [37]. GPUs are high-throughput devices that are widely used for accelerating a variety of scientific applications due to the ease of programming and the ability to provide large amounts of processing power and memory bandwidth.

Prior studies investigated how to improve the performance of SpMV kernels on GPU. [71,79,81] proposed automated tweaking of different parameters, such as the maximum number of registers per thread or thread block size according to the sparsity pattern of the input matrix. Another research direction was related to reducing the memory footprint of SpMV kernels by proposing different efficient sparse formats such as COO [88], CSR [88], ELL [88], BELL [90], and SELL [90]. However, as shown in many studies [23, 73, 78], using an inappropriate sparse format can result in significant performance degradation (up to 1.6× [41]). To address this limitation, recent research studies proposed selecting the optimal sparse format automatically to minimize the execution time and/or reduce the energy consumption of SpMV kernels on GPU [31, 32, 42, 78]. The proposed sparse format selection methods have been successful, but they ignore the importance of optimizing power consumption (MFLOPS/Watt), which prevents GPUs from processing SpMV kernels in many low-power applications [26]. Furthermore, GPU performance of the SpMV kernel has a high sensitivity to the compiler configuration parameters (thread block size, maximum number of registers per thread, and memory hierarchy configuration) [81]; nonetheless, the proper selection of these parameters has not been explored in the literature. Finally, it is not conclusive that the results of the prior works are reproducible, primarily since a comprehensive evaluation cannot be found in their report [33].

To tackle shortcomings of the state-of-the-art methods, we propose a machine learning framework that automatically tweaks both sparse format and compilation parameters, targeting different optimization objectives such as 1) latency (Second), 2) energy consumption (Joule), 3) average power consumption (Watt) and 4) energy efficiency (MFLOPS/Watt) of the SpMV kernel on GPU. We call this method automated tweaking of the SpMV kernel configuration parameters or Auto-SpMV. Auto-SpMV enables two modes of optimization: *compile-time* and *run-time*. In the *compile-time* mode, Auto-SpMV tweaks compiler parameters while in the *run-time* mode, Auto-SpMV automatically selects the most efficient sparse format and programming kernel according to the structural characteristics of the sparse input matrix.

In this paper, we consider CSR, ELL, BELL, and SELL sparse formats due to simple implementation and providing higher performance in many applications over other hybrid formats [73]. To characterize the performance of SpMV kernels for sparse matrices, we consider simple sparsity features of the sparse input matrix. In the learning process, a multi-class classifier is trained on a dataset of sparse matrices characterized by their sparsity features. Then, we use this model to predict the optimal sparse format, maximum number register per thread (`maxrregcount`), memory hierarchy configuration, and thread block size (TB size) for unseen inputs.

To achieve maximum learning performance, the classifier should be trained on a comprehensive dataset. In order to fulfill this requirement, we collect a dataset of 30 sparse matrices from the SuiteSparse matrix collection [61] that contains records of SpMV kernels running on GPU obtained from different compiler parameters and sparse formats. To the best of our knowledge, Auto-SpMV provides the largest training dataset containing 15520 records with ≈70 runs on two GPUs (Turing and Pascal architectures) compared to all existing studies on automated SpMV optimization. To improve the classification accuracy, we need to fine-tune the learning hyperparameters such as the learning rate of neural networks. In order to fulfill this requirement, we use an Automated Machine Learning (AutoML) tool to efficiently optimize the learning pipeline (Section 5.4). Then, we report the best classification results. In this paper, we also devise regression models to estimate various optimization objectives on a sparse input matrix.

Our main contributions are summarized as follows:

1. Enhancing the performance of SpMV kernels on GPU through the use of efficient machine learning classifiers. This will enable us to select the optimal compilation parameters and sparse format based on different optimization objectives.

2. As the general demand for green GPU processing grows [30], Auto-SpMV considers new optimization objectives: energy efficiency and average power consumption. Along with latency and energy consumption, evaluating these two optimization objectives is extremely tedious since an SpMV kernel must be run many times on GPU.

3. Using highly-optimized classifiers rather than relying on widely-used models trained with default learning hyperparameters.

4. Collecting the largest ever dataset containing the record of running 30 different sparse matrices with more ≈70M runs over 15k configuration settings on two different GPU architectures. Dataset, code, and configuration parameters will be available upon acceptance.

Results show that Auto-SpMV optimizes latency, energy consumption, average power, and energy efficiency in the *compile-time* mode by up to 51.9%, 52%, 33.2%, and 53%, respectively, over all selected sparse matrices compared to the best default configuration settings. Auto-SpMV optimizes average power and energy efficiency in the *run-time* mode by up to 34.6%, and 99.7%, respectively. We analyze the overhead introduced by the run time kernel selection and translation between different sparse formats. Finally, Our experimental results on two different NVIDIA® GPU devices, GTX 1650-mobile (Turing architecture) and GTX 1080 (Pascal architecture), show that Auto-SpMV prediction results are not sensitive to the underneath GPU architecture.

The rest of this paper is organized as follows: Section 2 introduces the sparse formats considered and how the SpMV kernel is executed on GPU. In Section 3, we introduce the problem of selecting sparse formats and compiler parameters. In Section 5, we present the Auto-SpMV optimization framework. Section 4 presents the configuration setup for evaluating our proposed method. In Section 7, we evaluate the Auto-SpMV framework and compare it to the other baselines. We thoroughly discuss the proposed results and analyze the processing overhead of Auto-SpMV in Section 8. In Section 9, we review the related studies. Finally, we conclude our paper in Section 10.

2 Background

2.1 GPU Architecture

GPUs are designed to maximize computing throughput by executing parallel workloads, even though sacrificing a single task's serial performance. GPUs offer a large number of Streaming Processors (SPs) equipped with fully-pipelined logic units. SPs are grouped in a set of Streaming Multiprocessors (SMs) to execute instructions in Single-Instruction Multiple-Threads (SIMT) mode. A kernel is a GPU-related part of an application and is executed by thousands of threads arranged in a grid of thread blocks defined by the programmer. Thread blocks are the groups of threads that are executed simultaneously in SMs. The SM runs threads in groups of 32 parallel threads called warps. To achieve the maximum throughput, every SM unit should have enough active warps to hide memory and instruction pipeline latencies.

GPUs feature a rich memory hierarchy configuration that includes L1 and L2 caching capabilities, texture memory, shared memory, and global memory. Memory hierarchy can be defined by the programmer, allowing L1 cache and shared memory to be configured manually. Thread blocks use shared memory that can be accessed and shared by all threads within a block. Programmers can identify the upper-bound number of registers used by a thread at the compile time by using the `maxrregcount` variable. According to the default settings, CUDA® compilers try to minimize register usage to maximize the number of thread blocks that can be active simultaneously within the SM, leading to maximizing occupancy within the SM. An overview of GPU architecture is illustrated in Figure 1.

2.2 SpMV Processing on GPU

The Sparse Matrix-Vector Multiplication (SpMV) is a non-trivial Level-2 BLAS operation that finds the dense vector product Y of a sparse matrix A and a dense vector X such that Y = A × X. Figure 2.a shows an example of SpMV operation for matrix A and dense vector X. Dense formats are inefficient for implementing SpMV kernels on GPU due to storing zeros and the zero multiplications result in inefficient memory usage and computation, respectively. Several sparse formats have been proposed to address this problem (Section 2.3).

2.3 Popular SpMV Formats

Several sparse formats have been proposed in order to improve the efficiency of running SpMV kernels on GPU by avoiding unnecessarily 1) storing zero elements, and 2) performing computations on them. Note that there is no best-for-all sparse format for all SpMV matrices [23, 73, 78]. In other words, if a sparse format is not carefully selected, significant performance degradation is likely to occur [78]. Due to the sparsity features of matrices, the sparse format can be used in flat, blocked, and composite types. Here is a brief explanation of the most popular sparse formats that we used in this paper.

CSR format [88]. According to the CSR format, non-zero elements and column indices are stored explicitly in *Data* and *Column Index* arrays, respectively. The boundaries of each row are saved in a third array called *Row Index*. Figure 2.b shows the CSR format of matrix A. This format requires coordination among threads within a warp to accumulate per-thread results together. Although this format requires no zero-padding, the access to the *Data* and *column Index* might be unaligned. Having random access to X, as well as the varied number of non-zero elements per row, this kernel may suffer from the load imbalance problem.
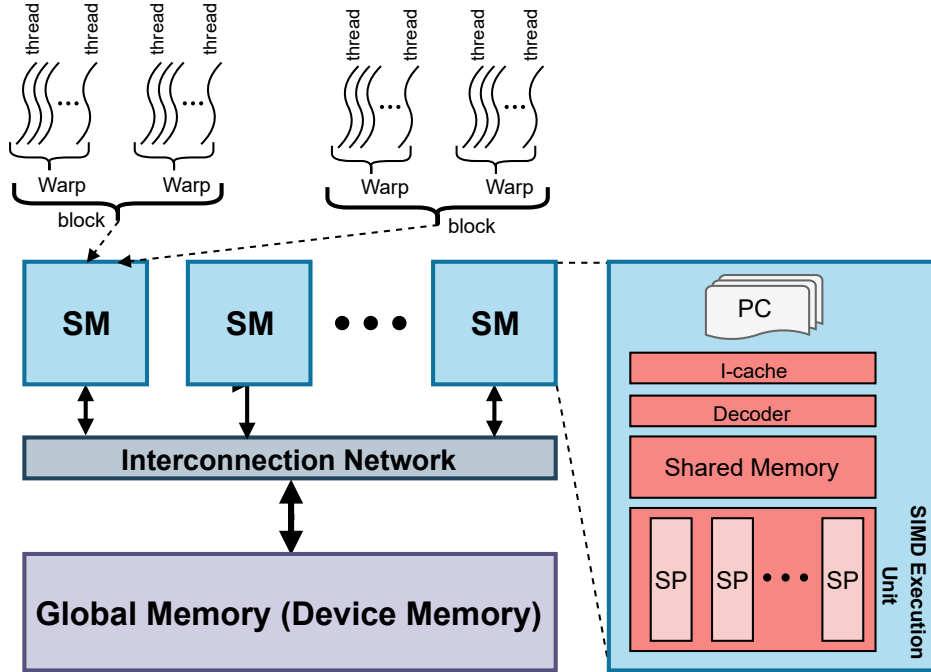
Figure 1: An overview of GPU architecture.

ELL format [88]. In the ELL format, non-zero elements are grouped into a dense *Data* matrix, and the column indices of matrix A are stored in a *Column Index* matrix. Dimensions of the *Data* matrix will be $m \times max\_nnz$, in which max_nnz is the maximum number of non-zero elements in a row. Figure 2.c shows the ELL format of matrix A. Despite the efficiency of the ELL format, it does not necessarily guarantee contiguous access to the input vector X.

Blocked ELL (BELL) format [90]. This is a variation of the ELL format in which a block of non-zero elements is considered as an element of the ELL format. Figure 2.d shows the BELL format of matrix A with the default block size of 2×2. The BELL format is composed of two matrices: 1) *Data* matrix, which is used to store blocks of non-zero elements, and 2) *Column Index* matrix, which is used to store the block indices. In general, this format is suitable for matrices with a uniform distribution of non-zero blocks.

Sliced ELL (SELL) format [90]. Each slice of this format consists of a constant number of rows of matrix A. The length of each slice is the maximum number of non-zero elements per row in that slice. For compressing a matrix to SELL format, each slice of the matrix must be considered as a matrix packing in ELL format, and this is achieved by storing the pointers of each slice. SELL format is composed of three matrices including 1) *Slice Index* saves pointers to slices; 2) *Data* matrix stores slices of non-zero elements; and 3) *Column Index* matrix stores column indices for the *Data* elements. Figure 2.e shows the SELL format of matrix A with slice height set to 2. This format is suitable for sparse matrices with a variety of elements per row that are non-zero.
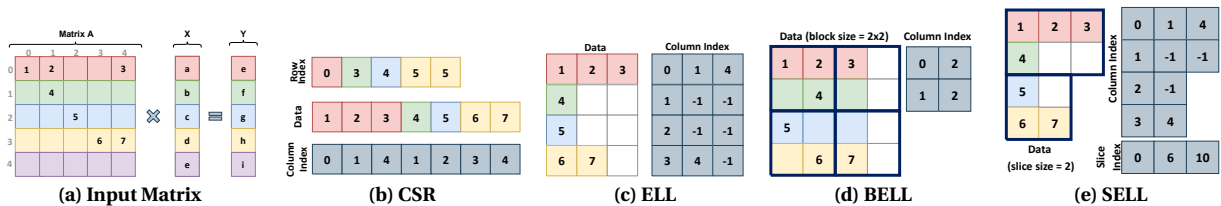

Figure 2: Illustration of (a) sparse matrix-vector multiplication, (b) CSR, (c) ELL, (d) BELL, and (e) SELL formats. Empty colorful cells are zero values of the sparse input matrix and white cells are zero padding.

4

3   Research Motivation

In this section, we first demonstrate the importance of tweaking compilation time parameters and the inefficiency of the default optimization scheme. Then, we discuss the limitations of prior learning models that attempted to model the latency and energy consumption of an SpMV kernel on GPU.

3.1   Inefficiency of the Default Optimization Parameters

Figure 3 compares the results of optimizing storage format and compiler parameters of the SpMV kernel by Auto-SpMV with the default CUDA® compiler configuration parameters for the *consph* sparse matrix given as a sample benchmark. Results are normalized to the Auto-SpMV (the higher is better). Because the CSR [88] format is popular, we have selected it as CUDA's default sparse format. The results demonstrate that Auto-SpMV provides at least 2.04× lower latency, 2.07× lower energy consumption, 1.08× lower average power consumption, and 2.086% better energy efficiency for the selected sparse matrix compared to the default CUDA® compilation parameters. Clearly, the default configuration parameters are not optimal, resulting in a significant degradation for all optimization objectives. Optimizing the compiler configuration parameters and the sparse format of SpMV kernels is therefore of paramount importance.
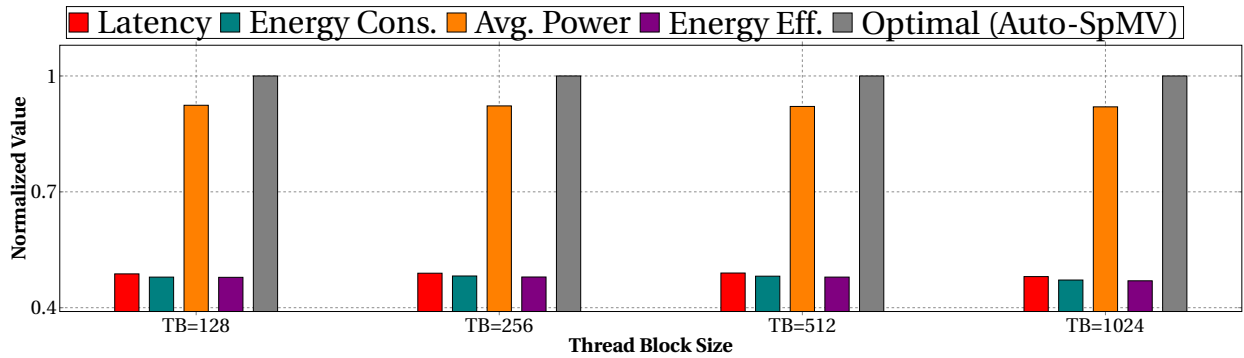

Figure 3: Comparing Auto-SpMV results with the default configuration parameters for the *consph* sparse matrix. Auto-SpMV is selected as the normalization baseline (the higher is better).

3.2   Limitations of Existing Learning Models

Due to the complexity of GPU architectures and the wide range of different sparsity features, it is very hard to describe the effect of each one of these sparsity features on the performance of SpMV kernels on GPU [78]. The lack of such thorough understanding is behind the motivation of using a machine learning approach. Efficient execution of the SpMV kernels on GPU requires optimizing several software and hardware parameters, resulting in a large optimization space for selecting the optimal values. Therefore, utilizing an exhaustive optimization algorithm to find the optimal solution is not feasible in practice since the procedure is extremely time-consuming [76]. Several analytical models have been proposed to estimate the performance of SpMV kernels and select the best sparse format [24]. However, analytical models cannot cope with the complexity of modern applications due to hardware complexities and requiring assumptions to simplify the problem [23].

To tackle the existing challenges, prior studies [21–23, 31, 32, 67] proposed to devise machine learning models to 1) estimate the performance and energy consumption of SpMV kernels on GPU, and 2) predict the near-optimal configuration parameters. The use of machine learning models offers a number of benefits, including that they are independent of hardware specifications, such as the number of cores per SM. Further, assumptions that simplify the problem are not necessary. Despite the success of prior studies in building machine learning models, they suffer from the inaccurate prediction (up to 16% accuracy loss [78]) due to 1) collecting a relatively small training dataset, and 2) evaluating a limited number of machine learning models. To address these challenges, Auto-SpMV devises highly accurate machine learning models trained on a large dataset consisting of 15520 samples. Approximately 70 million runs were performed on two different NVIDIA® GPU architectures, Pascal and Turing, in order to construct such a comprehensive training dataset.

5

4 Configuration Parameters

In this section, we provide an ablation study to individually examine the contribution of each compilation parameter and sparse format to the optimization objectives. Figure 4 shows the impact of optimizing each individual configuration setting of SpMV kernels by Auto-SpMV on various optimization objectives for the *eu-2005* sparse matrix given as a sample benchmark.
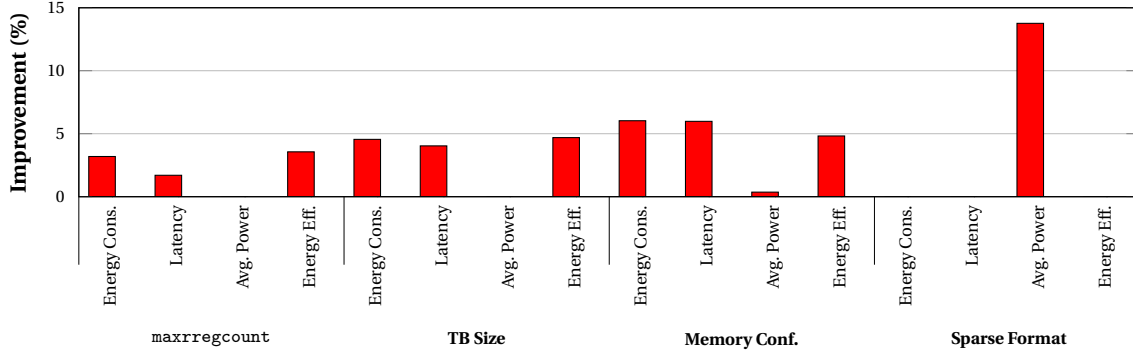

Figure 4: Individual improvement of each configuration parameter provided by Auto-SpMV over different optimization objectives. We select *eu-2005* as the sample benchmark.

The experimental results provided in Figure 4 have led to the following observations:

1. **Maximum number of registers per thread (`maxrregcount`).** Register memory is the thread-private memory that is divided among all resident threads on an SM. All local variables are placed in registers or local memory. Because local memory is placed in device DRAM (although it can be cached on-chip) and registers are on-chip, it is preferable for thread-private variables to be placed in registers. The maximum number of registers that can be used per thread in a kernel is controlled by the compiler. Nonetheless, the programmer can declare the number of registers used in every kernel in a compilation unit. Limiting the number of registers per thread can lead to increasing the number of blocks that can concurrently reside on an SM, which by itself can result in better latency hiding. Nevertheless, restricting the number of registers can cause register spilling which occurs when there are not enough registers available for a given task. Hence, registers can spill to global memory. Because of the opposing factors of register spilling and higher occupancy, some experiments are often needed to obtain the optimal configuration [58].

2. **Thread block size (TB Size).** The SpMV kernel is a memory-bound application, and a kernel with a higher occupancy rate has more latency-hiding ability. For improving kernel occupancy, we need to increase the size of the thread block. However, the higher value for thread block size might result in imbalanced resource usage in each SM. Determining the proper thread block size while sharing the workload equally across all threads results in fewer idle threads, thereby leading to improved performance. On the other hand, occupancy is an indicator of thread parallelism in a CUDA® program and can be increased by increasing the thread block size. However, maximizing the thread block size can lead to wasting achievable parallelism, e.g., in the case of suspended thread blocks, there are fewer other thread blocks to be active. Hence, the optimal number of threads per block is the result of the occupancy and performance trade-off and can be obtained after some experimentation.

3. **Memory hierarchy configuration.** For memory-bound kernels, such as SpMV, the optimization challenge is to reduce the memory latency time by leveraging the fast memory resources and caching capabilities of GPU. Exploiting the reconfigurability of the memory hierarchy and other hardware resources, such as registers, may lead to the energy optimization of GPU devices. The ideal case is that after each thread's first load, the next elements will reside in the L1 cache, but up to 2048 threads can run on an SM sharing L1 cache and likely replace each other's data. Although increasing the size of the L1 cache on GPU for data-intensive applications is beneficial, there might be some additional latency when using L1 due to the need to translate the global address you are accessing to a cache location. Hence, various input matrices with a variable number of non-zero elements in rows have different cache and shared memory needs.

4. **Sparse format.** The irregular computations involved in SpMV make its performance optimization challenging. Hence, enormous attempts have been devoted to devising data formats to store the sparse matrix with the aim of maximizing performance. Due to the sparsity features of matrices, sparse formats can be

6

used in flat, blocked, and composite types. Therefore, we select four widely-used sparse formats which have fewer specifications to sparsity features and can be used generally.

Although the main method for improving the efficiency of SpMV kernels on GPU is choosing the best sparse format [29]; the kernel optimization does not end at this point. As demonstrated in this section, optimizing the compiler parameters has a higher impact on improving the efficiency of SpMV kernels on GPU. Therefore, Auto-SpMV enables optimizing both *compile-time* (thread block size, `maxrregcount`, and memory hierarchy configuration) and *run-time* parameters (sparse format) by devising highly accurate machine learning models. Furthermore, Auto-SpMV is capable of finding optimal configuration settings for a wider range of optimization objectives, including latency, energy consumption, energy efficiency, and average power consumption.

5   Auto-SpMV Optimization Method

5.1   Method Overview

As shown in Figure 5, Auto-SpMV optimizes SpMV kernels in the *compile-time* and *run-time* modes for various optimization objectives. In the *compile-time* mode (Section 5.2), the compiler configuration parameters are optimized. In the *run-time* mode (Section 5.3), Auto-SpMV optimizes the sparse format of the sparse input matrix. Both optimization modes are conducted on the CPU. In the rest of this section, we present the details of optimization modes, the sparsity features extracted from input matrices, and the Auto-SpMV performance prediction pipeline.
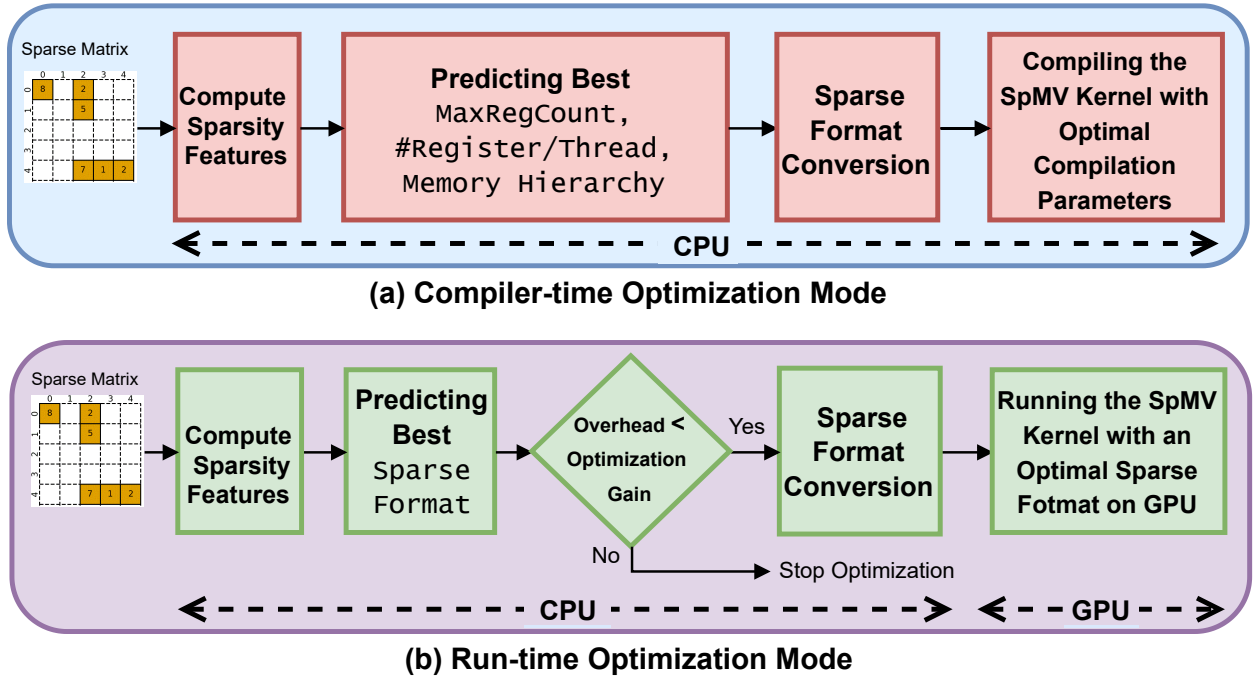

(a) Compiler-time Optimization Mode

(b) Run-time Optimization Mode

Figure 5: The overview of the proposed Auto-SpMV optimization framework. (a) *compile-time* Optimization Mode. (b) *run-time* Optimization Mode.

5.2   *Compile-time* Optimization Mode

The aim of *compile-time* optimization mode is to achieve the maximum GPU efficiency by tweaking the CUDA® compilation parameters. As shown in Figure 5, the *compile-time* optimization mode includes the following steps:

1. Compute the sparsity features.
2. Predict the optimal GPU compilation parameters.
3. Convert the sparse input matrix to CSR storage format (as default format) and compile CSR SpMV kernel with the optimal CUDA® compilation parameters.

We do not report the optimization overhead for the *compile-time* optimization mode since the whole procedure is performed at the compilation time.

5.3 *Run-time* Optimization Mode

The aim of the *run-time* optimization mode is to find the best-performing SpMV kernel and its sparse format for a sparse input matrix at run time. Given a sparse matrix in a default sparse format with the default compilation parameters, the following steps are performed to predict its best sparse format:

1. Compute the sparsity features.
2. Predict the optimal sparse format for the given sparse matrix.
3. Estimate the optimization overhead.
4. Convert the sparse input matrix to the predicted best sparse format if the optimization benefit is higher than the estimated optimization overhead.

With the aim of preventing unnecessary conversions for applications that are not composed of iterative solvers, Auto-SpMV leverages machine learning-based models for estimating the overhead of feature extraction and conversion to determine whether it is worthwhile to perform optimizations and tolerate the introduced overheads by subtracting estimated overhead (Section 7.5) from the predicted gain. Figure 6 shows the performance of estimating feature extraction and conversion overheads during the *run-time* optimization mode. Results show that Auto-SpMV accurately estimates the overhead of feature extraction and conversion at run time.

In our work, we used the implementations of the SpMV kernel under CSR, ELL, BELL, and SELL formats. The dataset used includes 30 sparse matrices from the SuiteSparse matrix collection [61] (Section 6.1) Computing sparsity features, overhead/sparse format predictions, and conversion to the predicted best sparse format are done at run time of the application on the CPU. The overhead of online sparse format optimization is discussed in Section 7.5.
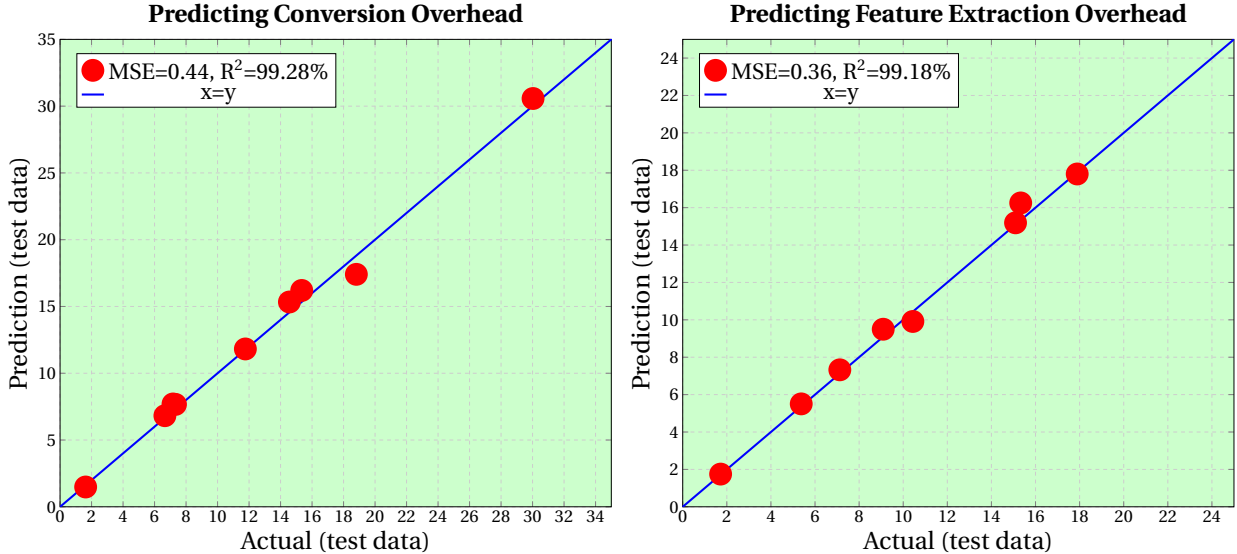

Figure 6: The performance of estimating feature extraction and conversion overheads during the online optimization mode.

5.4 Predicting Optimal Configuration Parameters

In general, we have two alternatives to learn the best configuration parameters:

1. Multi-class classification: train a classifier that learns the optimal configuration parameters based on extracted sparsity features. Note that different classifiers should be trained for different optimization objectives.
2. Regression: first, train different regression models for estimating different optimization objectives of the SpMV kernel. Then, select the optimal configuration parameters with the best estimations.

8

Auto-SpMV relies on the multi-class classification approach. However, we also provide the results of regression models that estimate different optimization objectives of the SpMV kernel (Section 7.4). Using a machine learning algorithm, we can predict the best configuration parameters and sparse format of SpMV kernels by the following steps:

1. **Step 1.** Extracting sparsity features of the sparse matrices that can represent the characteristics of running an SpMV kernel on GPU.

2. **Step 2.** Constructing a learning dataset of diverse sparse matrices.

3. **Step 3.** Fine-tuning machine learning algorithms to provide the most accurate predictions.

In this paper, we utilize supervised learning algorithms for both classification and regression tasks. Further, we employed Optuna, a state-of-the-art hyperparameter optimization library for AutoML [75]. Using Bayesian optimization [63] as the search method, Optuna automatically selects appropriate hyperparameters for a given dataset.

The main shortcomings of prior studies are 1) relying on one single learning model (e.g., BestSF only uses an SVM model [78]), and 2) constructing a relatively small-scale optimization space (e.g., [80] only contains two optimization candidates). On the other hand, Auto-SpMV constructs a large-scale dataset containing four optimization objectives and fine-tunes six different learning models (nearest centroid, decision tree, non-linear SVM, gradient boosting, random forest, and multi-layer perceptron) using Optuna to provide the most accurate results. Table 1 provides the hyperparameters and the corresponding range used to fine-tune the learning models. Section 6 presents the specification of the learning models after fine-tuning.

Table 1: The range of hyperparameters of each learning model selected for fine-tuning.

Machine Learning Model	hyperparameters
MLP	hidden layer size:{20, 50, 100, 150, 200}, #layer size:{1, 2, 3, 4, 5, 10}, Activation Function:{"Identity", "Logistic", "Tanh", "ReLU"}
Decision Tree	criterion:{"gini", "entropy", "log_loss"}, splitter:{"best", "random"}
Nearest Centroid Classifier	metric:{"manhattan", "euclidean", "minkowski"}
Non-linear SVM	kernel: {"linear", "poly", "rbf", "sigmoid", "precomputed"}
Gradient Boosting	#Estimators: 50, 100, 150, 200, learning rate: 0.1, 0.01, 0.001
Random Forest	criterion: {"gini", "entropy", "log_loss"}

5.5 Sparsity Features

In order to determine the optimal configuration parameters in both optimization modes, Auto-SpMV first needs to extract eight different sparsity features from the sparse input matrix. These sparsity features are briefly described in Table 2.

Table 2: Sparsity Features Selected by Auto-SpMV.

Feature	Definition
n	The number of rows
nnz	The number of non-zero elements
Avg_nnz	The average number of non-zero elements in rows
Var_nnz	The variance of non-zero elements in rows
ELL_ratio	The ratio of non-zero elements to the size of the matrix in the ELL format
Median	The median of non-zero elements in rows
Mode	The mode of non-zero elements in rows
Std_nnz	The standard deviation of non-zero elements in rows

We selected features based on 1) having the minimum computation overhead at run-time and 2) the performance impact reported by related papers [78]. The complexity of GPUs and the wide range of sparsity patterns make it impossible to explain exactly how each sparsity feature impacts the performance of the SpMV kernel on GPUs. Nevertheless, we list the following general observations regarding the performance of different sparse formats:

9

1. The number of rows (n) and the number of non-zero elements (**nnz**) characterize the volume of the SpMV kernel computation running on GPU.

2. The performance of the SpMV on GPU, depending on the used sparse format, is sensitive to the dispersion level of the non-zero elements on the rows of the sparse matrix. To characterize this sensitivity, we included `Avg_nz`, `Var_nnz`, `Std_nnz` sparsity features.

3. CSR SpMV kernel gives good performance for matrices with regular distribution of the non-zero elements, characterized by large values of `Avg_nz` and relatively small values of `Std_nnz`, for which these kernels do not suffer from unbalanced distribution among the threads.

4. Small values of `ELL_ratio` is a good indicator for the existence of many blocks of non-zero elements in the matrix which favors the ELL format because, in this case, ELL will less suffer from the problem of excessive zero padding that introduces more computation than needed.

6 Experimental Setup

6.1 Benchmark

SuiteSparse [61] provides an extensive collection of sparse matrices. SuiteSparse is derived from real-world applications, such as numerical linear algebra, computational fluid dynamics, computer graphics, and financial modeling. We select 30 sparse matrices from the SuiteSparse library on the basis of three criteria: 1) matrices with a wide range of dimension size (14,340<n<1,489,752), 2) matrices with a wide range of non-zero elements (800,800<(nnz)<19,235,140), and 3) matrices with the minimum similarity between their sparsity features to cover more diverse benchmarks. Figure 7 presents the distribution of sparsity features for the selected matrices sorted by the ascending order of *nnz*. Figure 8 shows the Pearson correlation of sparsity features. The results show that there exists a low degree of correlation among the sparsity features of selected matrices.

6.2 Hardware Specification

In this paper, we utilize two different NVIDIA® GPU architectures, Pascal and Turing, to evaluate the sensitivity of learning results to the variation of hardware devices. Table 3 presents the specification of utilized hardware devices. We use NVIDIA® CUDA® v. 11 to compile SpMV kernels on GPU.

6.3 GPU Performance Monitoring

In this paper, we consider four optimization objectives, including 1) latency, 2) energy consumption, 3) average power consumption, and 4) energy efficiency (MFLOPS/Watt). The execution time and energy consumption are measured using the multi-threading technique by a separate thread running on the host CPU besides another thread which is responsible for allocating memory and launching the kernel on GPU. For computing, the execution time, function `QueryPerformanceCounter()` of the NVIDIA® Management Library (NVML) is utilized. This function retrieves the current value of the performance counter, with a high-resolution timestamp ($< 1us$). Depending on the kernel complexity, We run each kernel 500-200000 times on GPU due to the required time intervals for GPU on-chip sensors in returning the measured instant power consumption [85]. The average of these runs is then reported as the kernel's execution time.

We employ GPU power sensors to obtain energy and power consumption. The thread running on the CPU queries the internal sensors via the NVML interface, which returns the power readings in milliwatts. Energy and average power consumption are computed without taking into account idle power samples. The arithmetic mean of power samples is considered as the average power consumption, and the integral of power samples over time samples is regarded as the energy consumption of an SpMV kernel. Finally, energy efficiency (MFLOPS/Watt) is calculated by dividing the achieved mega floating-point operations per second (MFLOPS) by the average power consumption.

6.4 Learning Algorithms

Table 4 presents the specification of the best learning models after fine-tuning. The rest of the specification of learning models is set to the default hyperparameters of Scikit-learn machine learning library [62]. We used 80% of the dataset for training and 20% for validating the prediction results.
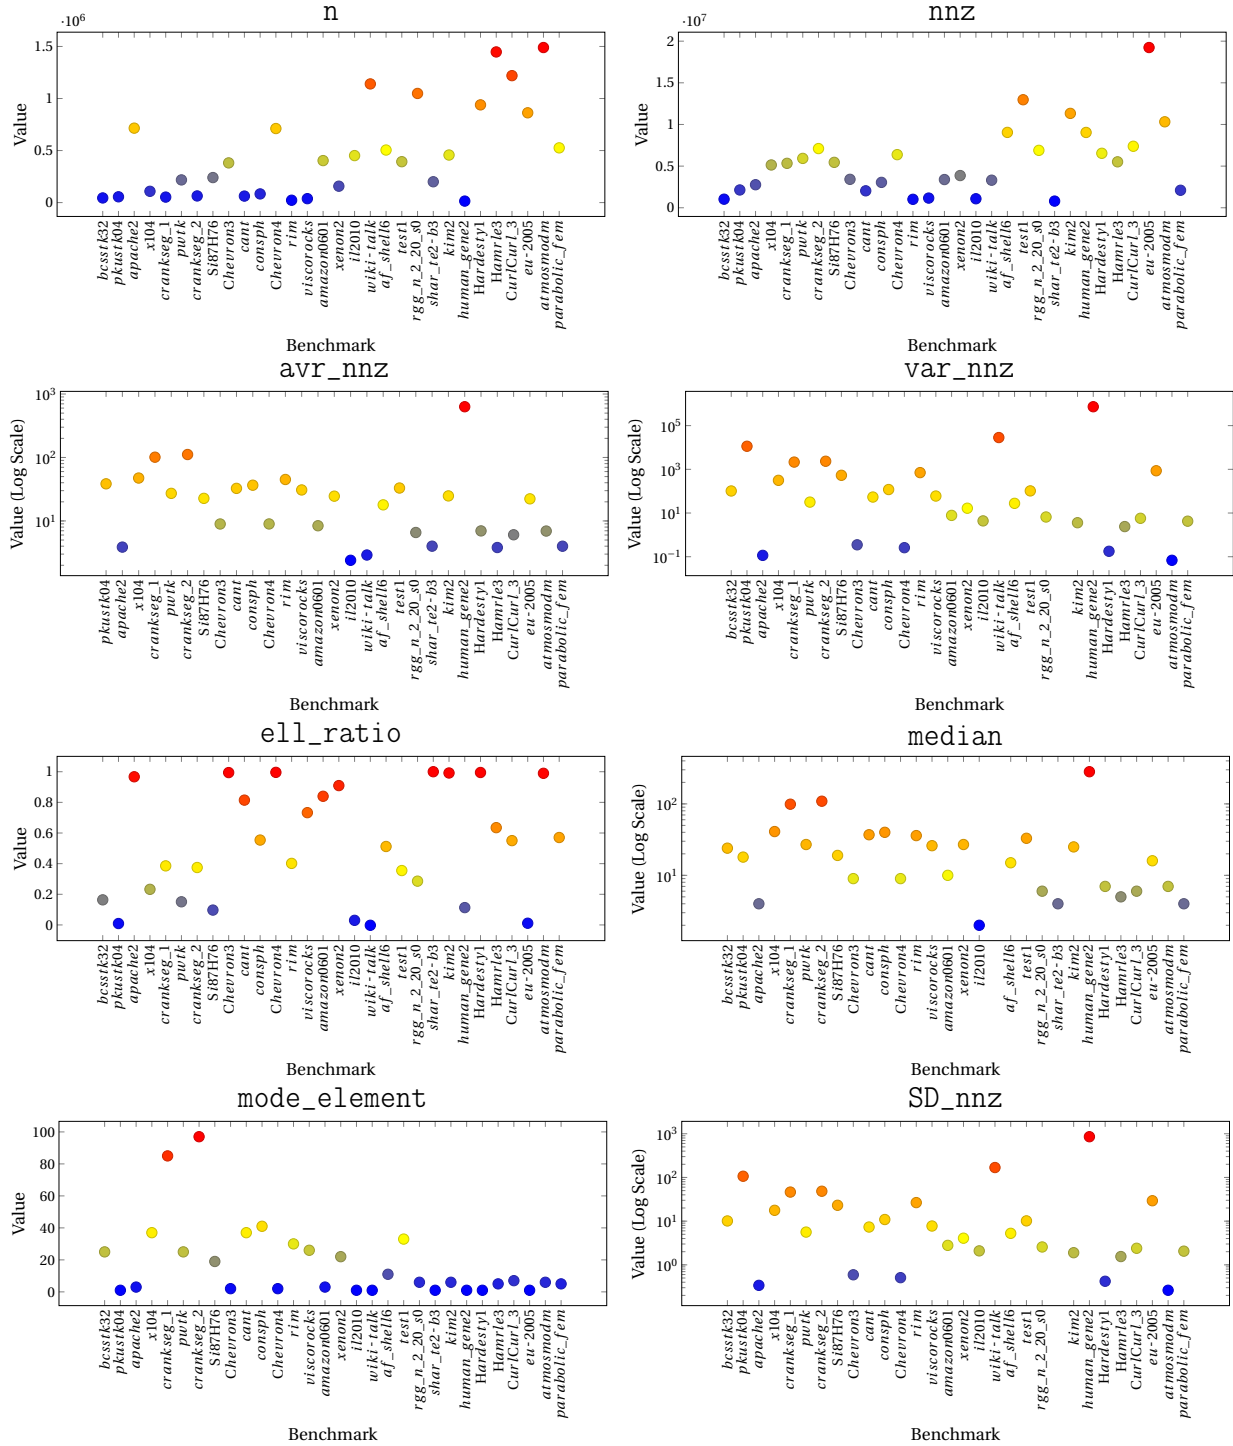
Figure 7: The distribution of the structural features over the studied matrices sorted by increasing values of *nnz*. We see that the selected matrices cover a wide range of sparsity features.
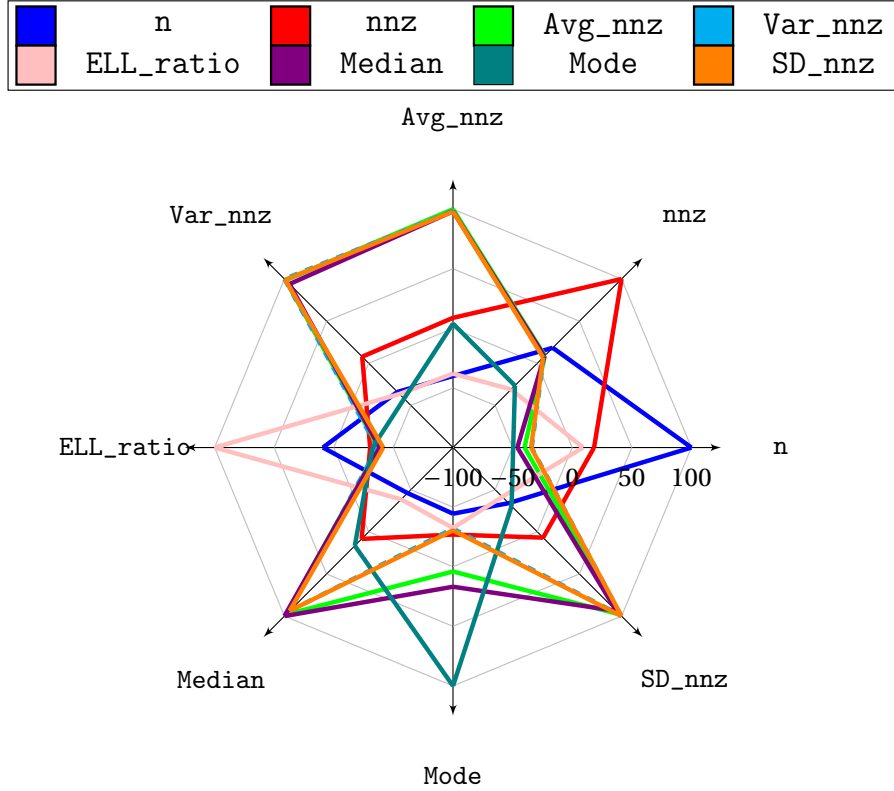
Figure 8: The Pearson correlation (%) of sparsity features of selected benchmarks. We see that there is no correlation between them.

Table 3: Hardware Specification.

GPU Device	Specification	Value
GTX 1650-mobile	Clock Frequency	1.6 GHz
	On-chip Memory	4GB GDDR5X
	# CUDA® Core	896
GTX 1080	Base Clock Frequency	1.6 GHz
	On-chip Memory	8GB GDDR5X
	# CUDA® Core	2560
CPU Device	**Specification**	**Value**
AMD® Ryzen 5	Clock Frequency	3.6 GHz
	Memory	16 GB

12

Table 4: Summarizing the specification of learning models after fine-tuning utilized for predicting the best hyperparameters (classification) and estimating various objective functions (regression) of SpMV kernels on GPU.

	Learning Algorithm	hyperparameter & Values
Classification	Nearest Centroid	metric=manhattan, #neighbors=5
	Decision Tree	Depth = 13
	Non-linear SVM	kernel=rbf, C=1.0, degree=3, gamma=scale
	Gradient boosting	loss=hinge, alpha=0.0001, #epochs=1000
	Random Forest	#estimators=100 The maximum Depth=15

	Multi-layer Perceptron	#layers=5, #nodes/layer=100
		activation=ReLU, #epochs= 200
		Optimizer=Adam and learning rate=0.001

Regression	Bayesian Ridge	#iter=300, tol=0.001
	Lasso	alpha=1.0 , #epochs=1000
	Random Forest	#estimators=100, Depth=None
	Decision Tree	Depth=None
	Lars	#non-zero coefs=500, eps=2.22045e-16,

	Multi-layer Perceptron	#layers=5, #nodes/layer=200
		activation=ReLU, #epochs= 200
		Optimizer=Adam and learning rate=0.0001

7 Experimental Results

7.1 Results of the *Compile-time* Optimization Mode

Figure 9 compares the performance improvement of Auto-SpMV in the *compile-time* optimization mode over the default parameters with the CSR sparsity format for different optimization objectives. The upper and lower bounds of each bar show the best and the worse results when tweaking the thread block size as the programmer is responsible for selecting this parameter. Results show that Auto-SpMV achieves up to 51.9%, 52%, 33.2%, and 53% improvement for all selected sparse matrices if we optimize the compiler parameters in terms of latency, energy consumption, average power, and energy efficiency, respectively.

7.2 Results of the *run-time* Optimization Mode

Figure 10 compares performance improvement of Auto-SpMV in the *run-time* optimization mode over the default sparsity format, CSR, for different optimization objectives. To ensure a fair comparison, we set the compilation parameters to the optimal values. By optimizing the sparse format at run time, Auto-SpMV provides up to 34.6% improvement in average power, as well as 99.7% improvement in energy efficiency for all selected sparse matrices. As CSR is the best sparse format for the latency and energy consumption optimization objectives, Auto-SpMV's performance is identical to that of the default setting.

7.3 Classification Results

Table 5 shows the best classification results achieved by Auto-SpMV for predicting the best configuration settings. We report accuracy and F1-Score as the evaluation metrics. Results show that Auto-SpMV successfully devises highly accurate learning models with 100% accuracy for predicting the optimal settings for TB size, `maxrregcount`, and memory hierarchy configuration. Note that the best result is provided by tuning the hyperparameters of the decision tree classifier [77].

Table 6 compares the classification results of Auto-SpMV with state-of-the-art. At least 10% more accurate classification is achieved by Auto-SpMV compared to the state-of-the-art. By evaluating different learning models and constructing a large-scale optimization space, Auto-SpMV is able to learn efficiently.
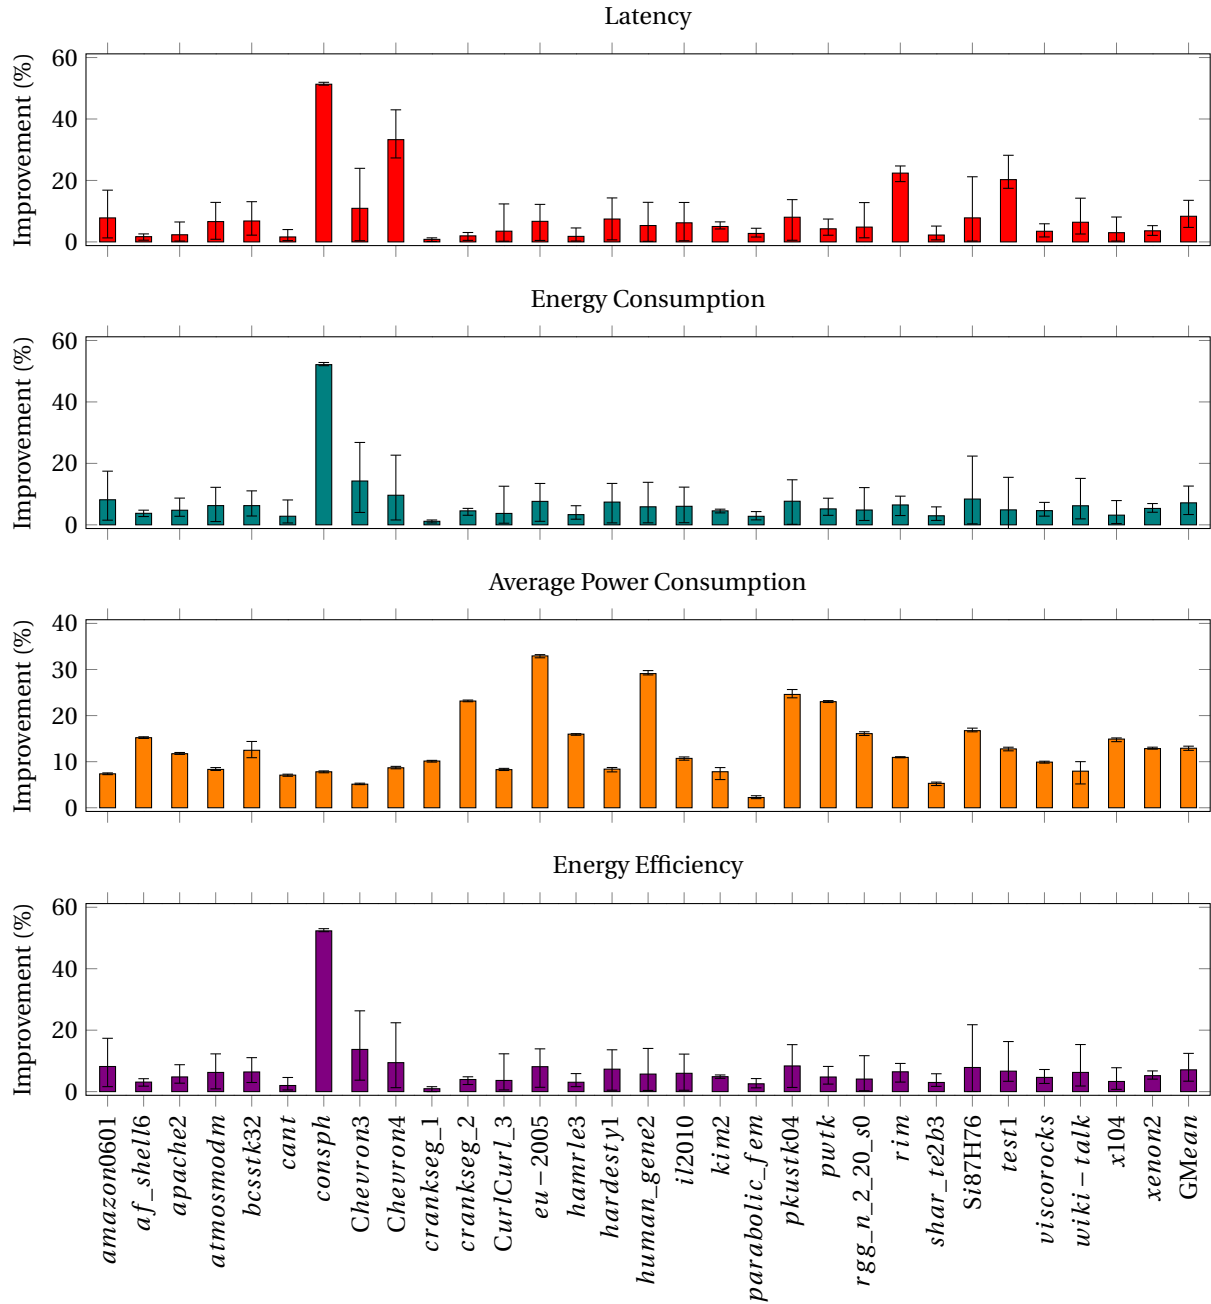
Figure 9: Comparing Auto-SpMV with the default configuration settings for different Optimization objectives.
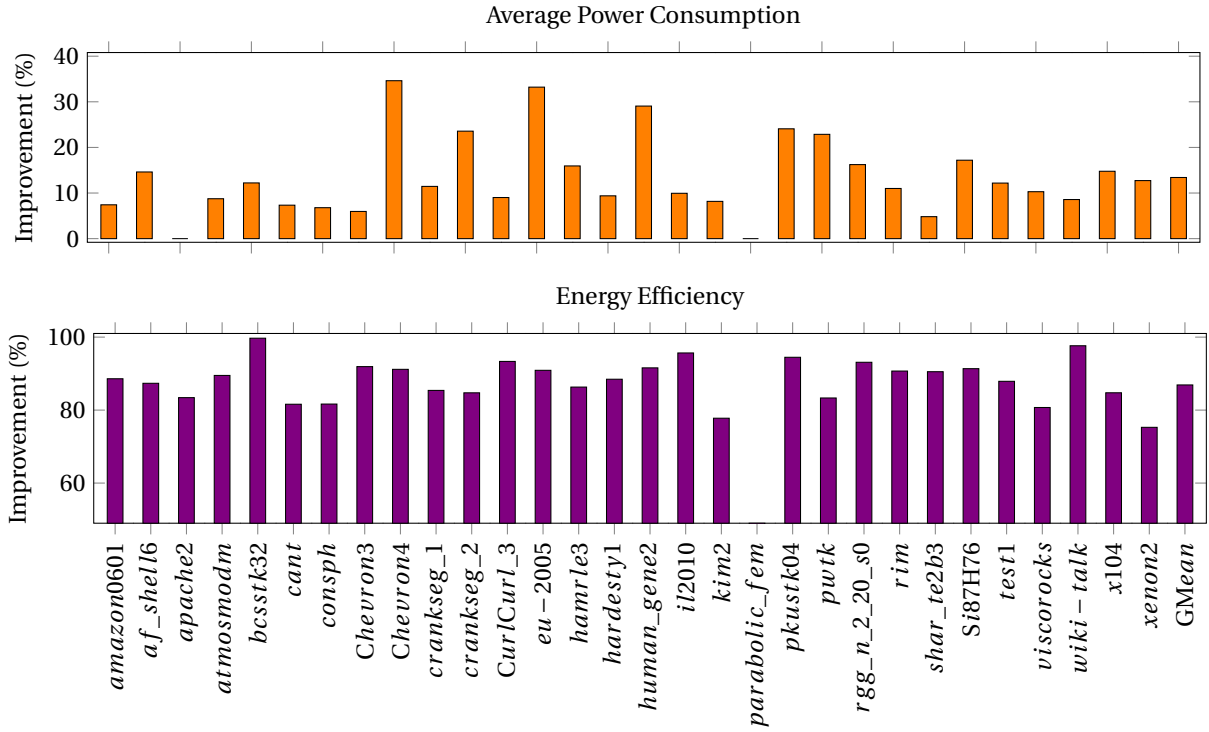
Figure 10: Auto-SpMV run time optimization over the default sparse format for different Optimization objectives.

Table 5: The classification accuracy of the optimal learning models (Decision Tree) used for predicting the best configuration settings for running SpMV kernels on GPU.

Classification	Latency		Energy		Average Power		Energy Efficiency	
(%)	Acc.	F1-Score	Acc.	F1-Score	Acc.	F1-Score	Acc.	F1-Score
TB Size	100	100	100	70	100	87.5	100	100
maxregcount	100	100	100	88.9	100	92.5	100	92.9
Memory	100	91.7	100	81.8	100	50	100	91.7

Table 6: Comparison of classification results of Auto-SpMV with state-of-the-art.

	Learning Model	Hardware Architecture	Accuracy (Execution Time)	Accuracy (Energy)
BestSF [78]	SVM	Maxwell	82%	-
[74]	Bagged Trees Classifier	Pascal	89%	84%
[32]	CNN	Pascal	90%	-
Auto-SpMV (ours)	Decision Tree	Turing	**100%**	**100%**

15

7.4 Regression Results

Figure 11 shows the regression results of estimating different optimization objectives. We report R^2 correlation, and MSE estimation metrics to identify the optimal learning model for each Optimization objective. Results show that random forest is the best learning model for estimating the energy consumption and energy efficiency of an SpMV kernel by providing 99.11% and 99.94% R^2 correlation, and 0.0086 and 0.00063 MSE, respectively. The decision tree is the best learning model for estimating average power by providing $R^2 = 99.99\%$ and MSE $= 2.82 \times 10^{-5}$. MLP is the best learning model for estimating latency by providing MSE $= 1.9 \times 10^{-2}$. Our success in using ML algorithms for estimating different objective functions is mainly due to 1) using a large training dataset that covers a broad range of sparse matrices with different sizes and sparsity features, and 2) tweaking the hyperparameters of learning models.
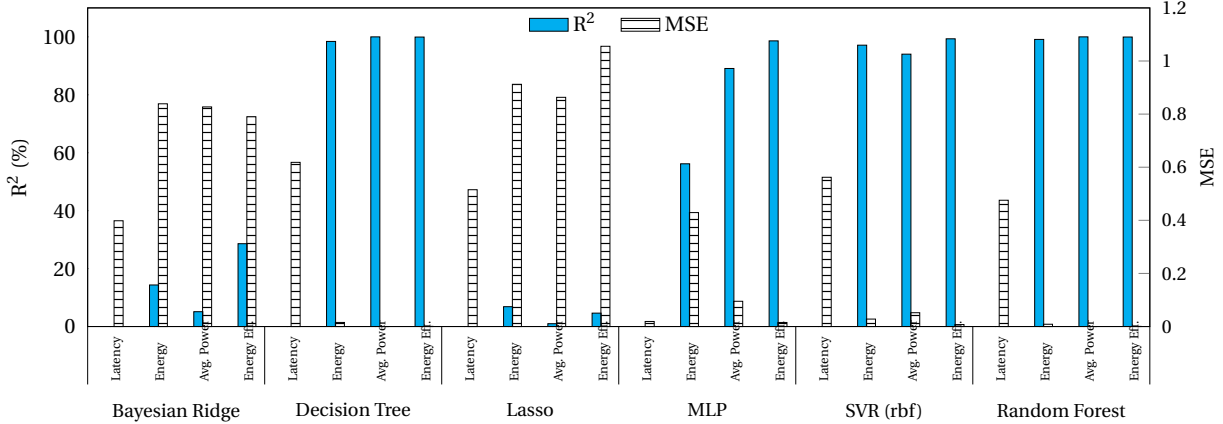

Figure 11: Evaluating different regression models for the estimation of different optimization objectives.

7.5 Run Time Prediction Overhead

In this section, we analyze the overhead of the optimization modes. Due to leveraging an overhead prediction mechanism (Section 5.3), Auto-SpMV does not perform unnecessary conversions when the conversion overhead is higher than the optimization benefit. The overhead of the optimization modes is computed based on the f_latency + o_latency + p_latency + c_latency. f_latency is the elapsed time to extract features, o_latency is the time needed to predict the conversion overhead, p_latency is the time needed to predict the optimal sparse format, and c_latency is the latency of converting to the predicted sparse format. The input matrix is often saved in a default sparse format, which is the COO sparse format in SuiteSparse [61]. Note that all the steps of the run time decision-making phase are executed on the CPU side using parallel implementations with Numpy Python arrays and matrices programming library.

Our measurements show that overall run time overhead is dominated by the sum of f_latency and c_latency since the o_latency and p_latency are constant values between $\approx$20 milliseconds while two other latencies are introduced according to the characteristics of the given input matrix. Table 7 shows the overhead of selected matrices if we consider lower latency as the optimization objective. The average f_latency+c_latency of our dataset is 24.2 seconds for the dataset while the maximum value is 87.8 seconds for a sparse matrix with nearly 1.9×10^7 non-zero elements (nnz). As also shown in [23, 78], this overhead is negligible for applications running large-scale sparse matrices such as Preconditioned Conjugate Gradient method which is an iterative solver. Thus, the total prediction overhead is insignificant for applications where an SpMV kernel is executed several times with the same input matrix in iterative methods such as eigenvalue problems and linear systems [69, 78].

7.6 GPU Sensitivity of Learning Models

We must verify that our classification results are independent of the underneath GPU architecture. To this end, we repeat a subset of evaluations for $amazon0601$, $crankseg\_2$, $bcsstk32$, $x104$, $il2010$, $Chevron3$ sparse matrices on the NVIDIA® 1080 GPU device with the Pascal architecture. Then, we compare the performance of predicted configuration settings provided by the Auto-SpMV classifier with the actual measurements on the NVIDIA® 1080 GPU device. Note that the prediction results of Auto-SpMV are based on running SpMV kernels on the NVIDIA® GTX 1650-mobile device with the Turing architecture. Figure 12 compares the optimal measurements of running SpMV

16

Table 7: Optimization Overhead of the Auto-SpMV Optimization Modes in seconds. Sorted based on the ascending order of nnz.

Matrix Name	nnz	f_latency	c_latency	f_latency+c_latency
shar_te2-b3	800800	1.71875	1.625	3.34375
rim	1014951	1.578125	2.046875	3.625
bcsstk32	1029655	1.71	2.125	3.835
il2010	1082232	3.625	2.5	6.125
viscorocks	1162244	1.90625	2.4375	4.34375
cant	2034917	3.4531	4.59	8.0431
parabolic_fem	2100225	5.46875	4.984375	10.45313
pkustk04	2137125	3.78	4.53125	8.31125
apache2	2766523	7.84	6.06	13.9
consph	3046907	5.375	6.65625	12.03125
wiki-talk-temporal	3309592	10.4375	7.3281	17.7656
amazon0601	3387388	7.125	7.171875	14.29688
Chevron3	3413113	7.0625	7.328125	14.39063
xenon2	3866688	6.75	9.375	16.125
x104	5138004	9.093	11.76563	20.85863
crankseg_1	5333507	9.78025	11.75	21.53025
Si87H76	5451000	9.828125	11.90625	21.73438
Hamrle3	5514242	15.09375	12.89063	27.98438
pwtk	5926171	11.39	13.8593	25.2493
Chevron4	6376412	13.76563	14.71875	28.48438
Hardesty1	6539157	15.09375	14.5625	29.65625
rgg_n_2_20_s0	6891620	15.32813	15.34375	30.67188
crankseg_2	7106348	13.03125	15.25	28.28125
CurlCurl_3	7382096	17.89063	18.8125	36.70313
human_gene2	9041364	17.01563	21.70313	38.71875
af_shell6	9046865	18.78125	21.4687	40.24995
atmosmodm	10319760	24.14063	23.90625	48.04688
kim2	11330020	22.70313	27.10938	49.8125
test1	12968200	23.78125	30.03125	53.8125
eu-2005	19235140	39.82813	47.98438	87.8125

17

kernels on an NVIDIA[®] GPU with the results predicted by Auto-SpMV. Results are normalized to the actual measurements. Results show that Auto-SpMV can predict the optimal configuration settings (TB Size and `maxrregcount`) with a marginal different (up to 2% performance loss). Accordingly, Auto-SpMV predicts the optimal configuration settings without relying on the underneath GPU architecture.
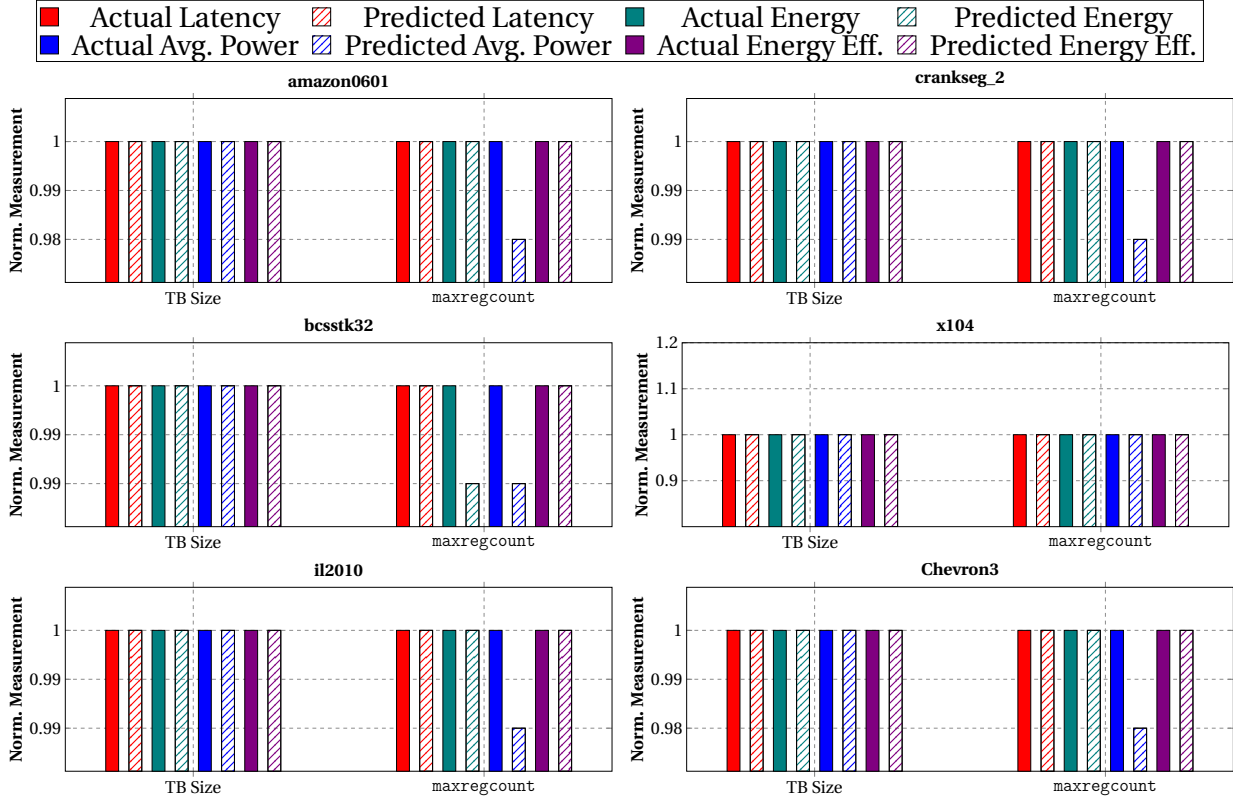

Figure 12: Evaluating efficiency of Auto-SpMV for predicting optimal configuration settings of new GPU devices.

8 Discussion

Our findings are summarized as follows:

1. There are no optimal configuration settings for all the sparse input matrices even on the same sparse format.

2. Carelessly selecting compiler parameters and storage format results in a significant performance loss of 13.3%, 5.4%, 2.0%, and 13.9% on average for latency, energy consumption, average power, and energy efficiency optimization objective, respectively (Figure 9).

3. Considering the CSR as the default sparse format, a careless choice of sparse format (*compile-time* mode) leads to a significant performance loss of 13.4% and 86.89% on average for the average power and energy efficiency optimization objectives, respectively (Figure 10).

4. Programmers cannot achieve the maximum improvement even when they try to select the optimal thread block sizes since `maxrregcount` and memory hierarchy configuration are selected by the CUDA[®] compiler.

5. Results demonstrate that CSR is not the best format for average power consumption and energy efficiency optimization objectives.

6. Increasing `maxrregcount` can reduce memory latency. In contrast, reducing `maxrregcount` can increase occupancy by allowing more threads to be active and thereby more warps to run. Therefore, an inefficient `maxrregcount` leads to 1) register spilling in global memory, and 2) decreased occupancy rate, resulting in increased latency and significant energy consumption [81].

7. Increasing the thread block size can result in more threads per SM, and therefore, a higher thread occupancy. On the other hand, when thread blocks are suspended, fewer thread blocks are active, resulting in less parallelism. Therefore, programmers should carefully consider the trade-off between increasing the occupancy of SMs and decreasing the number of thread blocks in order to achieve a higher level of parallelism, and as a result, be more efficient in terms of energy consumption and performance.

8. In order to fully utilize GPU computing power, memory access efficiency is essential. By determining the optimal size of the L1 cache and shared memory based on the sparse input matrix features, we can ensure better memory access efficiency and thus improved performance and energy efficiency.

9. The selection of a sparse format has a profound impact on data locality, cache performance, and memory bandwidth utilization. Results show that a significant amount of processing time and energy can be saved by automating the selection of the sparse format of SpMV kernels on GPU.

9 Related Work

To the best of our knowledge, Auto-SpMV is the first automated framework that optimizes the performance and energy efficiency of SpMV kernels on GPU considering hardware and software optimization simultaneously. It is also the first study to demonstrate that compiler parameters play an integral role in GPU utilization, run time performance, and energy consumption. The efficient running of SpMV kernels on GPU has been extensively researched in the past. Prior studies are categorized as 1) proposing novel SpMV sparse formats, 2) software optimization techniques for SpMV kernels, and 3) performance or energy estimation of SpMV kernels on GPU. In this section, we briefly discuss these methods and compare them with this paper.

SpMV sparse formats. [2, 4] present the earliest studies on accelerating SpMV kernels on GPU. They provide a detailed analysis of the memory access pattern of CUDA® implementation of the classical sparse formats, including ELL, CSR, and COO. During the last decade, many attempts have been made to propose a range of SpMV sparse formats [3, 12, 13, 15–18]. In general, these formats significantly improve memory saving and can be applied to any sparse matrix. However, as we demonstrate in Section 3, carelessly selecting the sparse format and compiler parameters incurs significant performance degradation as a result of low GPU utilization. To maximize the performance of SpMV kernels on GPU, in this paper, we predict the best configuration settings using highly optimized machine learning models trained on a large-scale dataset.

Several SpMV sparse formats utilized specific sparse matrix characteristics in order to provide higher efficient storage usage and indexing [5, 6, 8]. For example, CSR-VI and CSR-DU use non-zero values similarity to achieve more efficient sparse matrix compression [8]. CSX first analyzes a sparse matrix to find patterns such as repeated values or dense diagonals. Then encodes the matrix using sparse formats appropriate for the founded patterns [5]. However, these approaches are limited in applicability to matrices containing specific sparsity features, besides, requiring costly pre-processing to identify sparsity features. Therefore, these specific formats are neglected in our work.

Software optimization techniques for SpMV kernels. Many existing works investigate the potential of different performance enhancement techniques such as auto-tuning [19, 20], partitioning the sparse matrix and combining sparse formats [43, 45], improving the reuse of cache and registers [47, 48] and hardware-software co-design optimizations [49, 51]. Despite the performance improvement of these techniques, balancing the trade-off between energy consumption and performance in running SpMV on energy-hungry GPUs has not been addressed in the literature.

Performance or energy estimation of SpMV kernels on GPU. Many studies focused on accelerating SpMV kernels by proposing a performance estimation model to select the best configuration settings [21–23, 31, 32, 34, 42, 78]. These techniques mainly leverage machine learning techniques to learn the features of the sparse matrix and provide highly accurate predictions. In contrast with the state-of-the-art [34], our work considers more sparse matrix features to improve prediction accuracy. Besides latency and energy consumption [53, 57, 67], we model the energy efficiency and average power consumption of SpMV kernels to propose the best configuration settings for low-power applications, e.g., running a geometric flow task on embedded devices [72]. We also optimize different learning models using an AutoML tool to find the best-performing prediction model for a given sparse matrix.

10 Conclusion

This paper proposes Auto-SpMV, an automated framework for optimizing SpMV kernels on GPU by predicting the best-performing sparse format and compiler parameters. Auto-SpMV enables optimizing different objectives

including latency, energy consumption, average power consumption, and energy efficiency. Our experimental results using real-world benchmarks reveal that Auto-SpMV achieves remarkable performance and energy efficiency compared to the best default compiler parameters and sparse format. We hope our new findings on the performance characteristics of SpMV kernels on GPU inspire new techniques for improving the performance of future performance modeling algorithms.

References

[1] Im, E. & Yelick, K. Optimizing Sparse Matrix-Vector Multiplication on SMP.. *PPSC.* (1999)

[2] Bell, N. & Garland, M. Implementing sparse matrix-vector multiplication on throughput-oriented processors. *Proceedings Of The Conference On High Performance Computing Networking, Storage And Analysis.* pp. 1-11 (2009)

[3] Langr, D. & Tvrdik, P. Evaluation criteria for sparse matrix storage formats. *IEEE Transactions On Parallel And Distributed Systems.* **27**, 428-440 (2015)

[4] Bell, N. & Garland, M. Efficient sparse matrix-vector multiplication on CUDA. (Nvidia Technical Report NVR-2008-004, Nvidia Corporation,2008)

[5] Kourtis, K., Karakasis, V., Goumas, G. & Koziris, N. CSX: an extended compression format for SpMV on shared memory systems. *ACM SIGPLAN Notices.* **46**, 247-256 (2011)

[6] Li, J., Tan, G., Chen, M. & Sun, N. SMAT: an input adaptive auto-tuner for sparse matrix-vector multiplication. *Proceedings Of The 34th ACM SIGPLAN Conference On Programming Language Design And Implementation.* pp. 117-126 (2013)

[7] Su, B. & Keutzer, K. clSpMV: A cross-platform OpenCL SpMV framework on GPUs. *Proceedings Of The 26th ACM International Conference On Supercomputing.* pp. 353-364 (2012)

[8] Kourtis, K., Goumas, G. & Koziris, N. Optimizing sparse matrix-vector multiplication using index and value compression. *Proceedings Of The 5th Conference On Computing Frontiers.* pp. 87-96 (2008)

[9] Buluç, A., Fineman, J., Frigo, M., Gilbert, J. & Leiserson, C. Parallel sparse matrix-vector and matrix-transpose-vector multiplication using compressed sparse blocks. *Proceedings Of The Twenty-first Annual Symposium On Parallelism In Algorithms And Architectures.* pp. 233-244 (2009)

[10] Willcock, J. & Lumsdaine, A. Accelerating sparse matrix computations via data compression. *Proceedings Of The 20th Annual International Conference On Supercomputing.* pp. 307-316 (2006)

[11] Maggioni, M. & Berger-Wolf, T. AdELL: An adaptive warp-balancing ELL format for efficient sparse matrix-vector multiplication on GPUs. *2013 42nd International Conference On Parallel Processing.* pp. 11-20 (2013)

[12] Maggioni, M. & Berger-Wolf, T. CoAdELL: Adaptivity and compression for improving sparse matrix-vector multiplication on GPUs. *2014 IEEE International Parallel & Distributed Processing Symposium Workshops.* pp. 933-940 (2014)

[13] Liu, W. & Vinter, B. CSR5: An efficient storage format for cross-platform sparse matrix-vector multiplication. *Proceedings Of The 29th ACM On International Conference On Supercomputing.* pp. 339-350 (2015)

[14] Kreutzer, M., Hager, G., Wellein, G., Fehske, H. & Bishop, A. A unified sparse matrix data format for efficient general sparse matrix-vector multiplication on modern processors with wide SIMD units. *SIAM Journal On Scientific Computing.* **36**, C401-C423 (2014)

[15] Guo, D., Gropp, W. & Olson, L. A hybrid format for better performance of sparse matrix-vector multiplication on a GPU. *The International Journal Of High Performance Computing Applications.* **30**, 103-120 (2016)

[16] Chen, X., Xie, P., Chi, L., Liu, J. & Gong, C. An efficient SIMD compression format for sparse matrix-vector multiplication. *Concurrency And Computation: Practice And Experience.* **30**, e4800 (2018)

[17] Mofrad, M., Melhem, R., Ahmad, Y. & Hammoud, M. Efficient distributed graph analytics using triply compressed sparse format. *2019 IEEE International Conference On Cluster Computing (CLUSTER).* pp. 1-11 (2019)

[18] Li, Y., Xie, P., Chen, X., Liu, J., Yang, B., Li, S., Gong, C., Gan, X. & Xu, H. VBSF: a new storage format for SIMD sparse matrix–vector multiplication on modern processors. *The Journal Of Supercomputing.* **76**, 2063-2081 (2020)

[19] Yan, S., Li, C., Zhang, Y. & Zhou, H. yaSpMV: yet another SpMV framework on GPUs. *Acm Sigplan Notices.* **49**, 107-118 (2014)

[20] Hou, K., Feng, W. & Che, S. Auto-tuning strategies for parallelizing sparse matrix-vector (spmv) multiplication on multi-and many-core processors. *2017 IEEE International Parallel And Distributed Processing Symposium Workshops (IPDPSW)*. pp. 713-722 (2017)

[21] Tan, G., Liu, J. & Li, J. Design and implementation of adaptive SpMV library for multicore and many-core architecture. *ACM Transactions On Mathematical Software (TOMS)*. **44**, 1-25 (2018)

[22] Guo, P. & Zhang, C. Sparse Matrix Selection for CSR-Based SpMV Using Deep Learning. *2019 IEEE 5th International Conference On Computer And Communications (ICCC)*. pp. 2097-2101 (2019)

[23] Benatia, A., Ji, W., Wang, Y. & Shi, F. Machine learning approach for the predicting performance of SpMV on GPU. *2016 IEEE 22nd International Conference On Parallel And Distributed Systems (ICPADS)*. pp. 894-901 (2016)

[24] Guo, P., Wang, L. & Chen, P. A performance modeling and optimization analysis tool for sparse matrix-vector multiplication on GPUs. *IEEE Transactions On Parallel And Distributed Systems*. **25**, 1112-1123 (2013)

[25] Li, K., Yang, W. & Li, K. Performance analysis and optimization for SpMV on GPU using probabilistic modeling. *IEEE Transactions On Parallel And Distributed Systems*. **26**, 196-205 (2014)

[26] Mach, P. & Becvar, Z. Mobile edge computing: A survey on architecture and computation offloading. *IEEE Communications Surveys & Tutorials*. **19**, 1628-1656 (2017)

[27] Cass, S. Nvidia makes it easy to embed AI: The Jetson nano packs a lot of machine-learning power into DIY projects-[Hands on]. *IEEE Spectrum*. **57**, 14-16 (2020)

[28] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V. & Others Scikit-learn: Machine learning in Python. *The Journal Of Machine Learning Research*. **12** pp. 2825-2830 (2011)

[29] Kanellopoulos, K., Vijaykumar, N., Giannoula, C., Azizi, R., Koppula, S., Ghiasi, N., Shahroodi, T., Luna, J. & Mutlu, O. Smash: Co-designing software compression and hardware-accelerated indexing for efficient sparse matrix operations. *Proceedings Of The 52nd Annual IEEE/ACM International Symposium On Microarchitecture*. pp. 600-614 (2019)

[30] Qasaimeh, M., Denolf, K., Lo, J., Vissers, K., Zambreno, J. & Jones, P. Comparing energy efficiency of CPU, GPU and FPGA implementations for vision kernels. *2019 IEEE International Conference On Embedded Software And Systems (ICESS)*. pp. 1-8 (2019)

[31] Nisa, I., Siegel, C., Rajam, A., Vishnu, A. & Sadayappan, P. Effective machine learning based format selection and performance modeling for SpMV on GPUs. *2018 IEEE International Parallel And Distributed Processing Symposium Workshops (IPDPSW)*. pp. 1056-1065 (2018)

[32] Zhao, Y., Li, J., Liao, C. & Shen, X. Bridging the gap between deep learning and sparse matrix format selection. *Proceedings Of The 23rd ACM SIGPLAN Symposium On Principles And Practice Of Parallel Programming*. pp. 94-108 (2018)

[33] Lindauer, M. & Hutter, F. Best practices for scientific research on neural architecture search. *Journal Of Machine Learning Research*. **21**, 1-18 (2020)

[34] Li, M., Ao, Y. & Yang, C. Adaptive SpMV/SpMSpV on GPUs for Input Vectors of Varied Sparsity. *ArXiv Preprint ArXiv:2006.16767*. (2020)

[35] Oyarzun, G., Peyrolon, D., Alvarez, C. & Martorell, X. An fpga cached sparse matrix vector product (spmv) for unstructured computational fluid dynamics simulations. *ArXiv Preprint ArXiv:2107.12371*. (2021)

[36] Sgherzi, F., Parravicini, A. & Santambrogio, M. A Mixed Precision, Multi-GPU Design for Large-scale Top-K Sparse Eigenproblems. *ArXiv Preprint ArXiv:2201.07498*. (2022)

[37] Ahamed, A. & Magoulès, F. Accelerated solution of Helmholtz equation with Iterative Krylov Methods on GPU. *2014 IEEE Intl Conf On High Performance Computing And Communications, 2014 IEEE 6th Intl Symp On Cyberspace Safety And Security, 2014 IEEE 11th Intl Conf On Embedded Software And Syst (HPCC, CSS, ICESS)*. pp. 54-61 (2014)

[38] Schiesser, W. Computational mathematics in engineering and applied science: ODEs, DAEs, and PDEs. (CRC press,2014)

[39] Qiu, S., Liang, Y. & Wang, Z. Optimizing Sparse Matrix Multiplications for Graph Neural Networks. *ArXiv Preprint ArXiv:2111.00352*. (2021)

[40] Krawezik, G., Kuntz, S. & Kogge, P. Implementing sparse linear algebra kernels on the Lucata Pathfinder-A computer. *2020 IEEE High Performance Extreme Computing Conference (HPEC)*. pp. 1-6 (2020)

[41] Lehnert, C., Berrendorf, R., Ecker, J. & Mannuss, F. Performance prediction and ranking of spmv kernels on gpu architectures. *European Conference On Parallel Processing*. pp. 90-102 (2016)

[42] Sedaghati, N., Mu, T., Pouchet, L., Parthasarathy, S. & Sadayappan, P. Automatic selection of sparse matrix representation on GPUs. *Proceedings Of The 29th ACM On International Conference On Supercomputing*. pp. 99-108 (2015)

[43] Yang, W., Li, K. & Li, K. A parallel computing method using blocked format with optimal partitioning for SpMV on GPU. *Journal Of Computer And System Sciences*. **92** pp. 152-170 (2018)

[44] Su, B. & Keutzer, K. clSpMV: A cross-platform OpenCL SpMV framework on GPUs. *Proceedings Of The 26th ACM International Conference On Supercomputing*. pp. 353-364 (2012)

[45] Muhammed, T., Mehmood, R., Albeshri, A. & Katib, I. Suraa: A novel method and tool for loadbalanced and coalesced SpMV computations on gpus. *Applied Sciences*. **9**, 947 (2019)

[46] Naumov, M., Chien, L., Vandermersch, P. & Kapasi, U. Cusparse library. *GPU Technology Conference*. (2010)

[47] Xu, W., Zhang, H., Jiao, S., Wang, D., Song, F. & Liu, Z. Optimizing sparse matrix vector multiplication using cache blocking method on Fermi GPU. *2012 13th ACIS International Conference On Software Engineering, Artificial Intelligence, Networking And Parallel/Distributed Computing*. pp. 231-235 (2012)

[48] Nagasaka, Y., Nukada, A. & Matsuoka, S. Cache-aware sparse matrix formats for Kepler GPU. *2014 20th IEEE International Conference On Parallel And Distributed Systems (ICPADS)*. pp. 281-288 (2014)

[49] Kanellopoulos, K., Vijaykumar, N., Giannoula, C., Azizi, R., Koppula, S., Ghiasi, N., Shahroodi, T., Luna, J. & Mutlu, O. Smash: Co-designing software compression and hardware-accelerated indexing for efficient sparse matrix operations. *Proceedings Of The 52nd Annual IEEE/ACM International Symposium On Microarchitecture*. pp. 600-614 (2019)

[50] Sadi, F., Fileggi, L. & Franchetti, F. Algorithm and hardware co-optimized solution for large SpMV problems. *2017 IEEE High Performance Extreme Computing Conference (HPEC)*. pp. 1-7 (2017)

[51] Animesh, S. Algorithm Architecture Co-Design for Dense and Sparse Matrix Computations. (Arizona State University,2018)

[52] Mullen, J., Wolf, M. & Klein, A. Pakck: Performance and power analysis of key computational kernels on cpus and gpus. *2013 IEEE High Performance Extreme Computing Conference (HPEC)*. pp. 1-6 (2013)

[53] Giefers, H., Staar, P., Bekas, C. & Hagleitner, C. Analyzing the energy-efficiency of sparse matrix multiplication on heterogeneous systems: A comparative study of GPU, Xeon Phi and FPGA. *2016 IEEE International Symposium On Performance Analysis Of Systems And Software (ISPASS)*. pp. 46-56 (2016)

[54] Subramaniam, B., Saunders, W., Scogland, T. & Feng, W. Trends in energy-efficient computing: A perspective from the Green500. *2013 International Green Computing Conference Proceedings*. pp. 1-8 (2013)

[55] Woo, D. & Lee, H. Extending Amdahl's law for energy-efficient computing in the many-core era. *Computer*. **41**, 24-31 (2008)

[56] Mukunoki, D. & Ogita, T. Performance and energy consumption of accurate and mixed-precision linear algebra kernels on GPUs. *Journal Of Computational And Applied Mathematics*. **372** pp. 112701 (2020)

[57] Anzt, H., Tomov, S. & Dongarra, J. Energy efficiency and performance frontiers for sparse computations on GPU supercomputers. *Proceedings Of The Sixth International Workshop On Programming Models And Applications For Multicores And Manycores*. pp. 1-10 (2015)

[58] Fatica, M. & Ruetsch, G. CUDA Fortran for scientists and engineers. (Elsevier,2014)

[59] Li, Z., Wang, Y., Zhi, T. & Chen, T. A survey of neural network accelerators. *Frontiers Of Computer Science*. **11**, 746-761 (2017)

[60] Moons, B., Bankman, D. & Verhelst, M. Embedded Deep Learning. (Springer,2019)

[61] Davis, T. & Hu, Y. The University of Florida sparse matrix collection. *ACM Transactions On Mathematical Software (TOMS)*. **38**, 1-25 (2011)

[62] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V. & Others Scikit-learn:Machine learning in Python. *The Journal Of Machine Learning Research*. **12** pp. 2825-2830 (2011)

[63] Shahriari, B., Swersky, K., Wang, Z., Adams, R. & De Freitas, N. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings Of The IEEE*. **104**, 148-175 (2015)

[64] Ryoo, S., Rodrigues, C., Baghsorkhi, S., Stone, S., Kirk, D. & Hwu, W. Optimization principles and application performance evaluation of a multithreaded GPU using CUDA. *Proceedings Of The 13th ACM SIGPLAN Symposium On Principles And Practice Of Parallel Programming*. pp. 73-82 (2008)

[65] Documentation, C. v10. 0, 2018.

[66] Barreda, M., Dolz, M. & Castano, M. Convolutional neural nets for estimating the run time and energy consumption of the sparse matrix-vector product. *The International Journal Of High Performance Computing Applications*. **35**, 268-281 (2021)

[67] Benatia, A., Ji, W., Wang, Y. & Shi, F. Energy evaluation of Sparse Matrix-Vector Multiplication on GPU. *2016 Seventh International Green And Sustainable Computing Conference (IGSC)*. pp. 1-6 (2016)

[68] Braun, L., Nikas, S., Song, C., Heuveline, V. & Fröning, H. A simple model for portable and fast prediction of execution time and power consumption of GPU kernels. *ACM Transactions On Architecture And Code Optimization (TACO)*. **18**, 1-25 (2020)

[69] Li, R., Xi, Y., Erlandson, L. & Saad, Y. The eigenvalues slicing library (EVSL): Algorithms, implementation, and software. *SIAM Journal On Scientific Computing*. **41**, C393-C415 (2019)

[70] Kurzak, J., Bader, D. & Dongarra, J. Scientific computing with multicore and accelerators. (CRC Press,2010)

[71] Guo, P., Huang, H., Chen, Q., Wang, L., Lee, E. & Chen, P. A model-driven partitioning and auto-tuning integrated framework for sparse matrix-vector multiplication on GPUs. *Proceedings Of The 2011 TeraGrid Conference: Extreme Digital Discovery*. pp. 1-8 (2011)

[72] Wang, Y., Feng, B. & Ding, Y. TC-GNN: Accelerating Sparse Graph Neural Network Computation Via Dense Tensor Core on GPUs. *ArXiv Preprint ArXiv:2112.02052*. (2021)

[73] Filippone, S., Cardellini, V., Barbieri, D. & Fanfarillo, A. Sparse matrix-vector multiplication on GPGPUs. *ACM Transactions On Mathematical Software (TOMS)*. **43**, 1-49 (2017)

[74] Dufrechou, E., Ezzatti, P. & Quintana-Orti, E. Selecting optimal SpMV realizations for GPUs via machine learning. *The International Journal Of High Performance Computing Applications*. **35**, 254-267 (2021)

[75] Akiba, T., Sano, S., Yanase, T., Ohta, T. & Koyama, M. Optuna: A next-generation hyperparameter optimization framework. *Proceedings Of The 25th ACM SIGKDD International Conference On Knowledge Discovery & Data Mining*. pp. 2623-2631 (2019)

[76] Loni, M., Sinaei, S., Zoljodi, A., Daneshtalab, M. & Sjödin, M. DeepMaker: A multi-objective optimization framework for deep neural networks in embedded systems. *Microprocessors And Microsystems*. **73** pp. 102989 (2020)

[77] Safavian, S. & Landgrebe, D. A survey of decision tree classifier methodology. *IEEE Transactions On Systems, Man, And Cybernetics*. **21**, 660-674 (1991)

[78] Benatia, A., Ji, W., Wang, Y. & Shi, F. BestSF: a sparse meta-format for optimizing SpMV on GPU. *ACM Transactions On Architecture And Code Optimization (TACO)*. **15**, 1-27 (2018)

[79] Abu-Sufah, W. & Karim, A. An effective approach for implementing sparse matrix-vector multiplication on graphics processing units. *2012 IEEE 14th International Conference On High Performance Computing And Communication & 2012 IEEE 9th International Conference On Embedded Software And Systems*. pp. 453-460 (2012)

[80] Tsai, Y., Cojean, T. & Anzt, H. Sparse linear algebra on AMD and NVIDIA GPUs–the race is on. *International Conference On High Performance Computing*. pp. 309-327 (2020)

[81] Zardoshti, P., Khunjush, F. & Sarbazi-Azad, H. Adaptive sparse matrix representation for efficient matrix–vector multiplication. *The Journal Of Supercomputing*. **72**, 3366-3386 (2016)

[82] CUSP The NVIDIA library of generic paralle algorithms for sparse linear algebra and graph computations on cuda architecture gpus.. (https://developer.nvidia.com/cusp), [Online]

[83] Davis, J. & Chung, E. SpMV: A memory-bound application on the GPU stuck between a rock and a hard place. *Microsoft Research Silicon Valley, Technical Report14 September*. **2012** (2012)

[84] Ashari, A., Sedaghati, N., Eisenlohr, J., Parthasarath, S. & Sadayappan, P. Fast sparse matrix-vector multiplication on GPUs for graph applications. *SC'14: Proceedings Of The International Conference For High Performance Computing, Networking, Storage And Analysis*. pp. 781-792 (2014)

[85] Burtscher, M., Zecena, I. & Zong, Z. Measuring GPU power with the K20 built-in sensor. *Proceedings Of Workshop On General Purpose Processing Using GPUs*. pp. 28-36 (2014)

[86] Bradley, T. GPU performance analysis and optimisation. *NVIDIA Corporation.* (2012)

[87] Steinberger, M., Derlery, A., Zayer, R. & Seidel, H. How naive is naive SpMV on the GPU?. *2016 IEEE High Performance Extreme Computing Conference (HPEC).* pp. 1-8 (2016)

[88] Bell, N. & Garland, M. Implementing sparse matrix-vector multiplication on throughput-oriented processors. *Proceedings Of The Conference On High Performance Computing Networking, Storage And Analysis.* pp. 1-11 (2009)

[89] Godwin, J., Holewinski, J. & Sadayappan, P. High-performance sparse matrix-vector multiplication on GPUs for structured grid computations. *Proceedings Of The 5th Annual Workshop On General Purpose Processing With Graphics Processing Units.* pp. 47-56 (2012)

[90] Choi, J., Singh, A. & Vuduc, R. Model-driven autotuning of sparse matrix-vector multiply on GPUs. *ACM Sigplan Notices.* **45**, 115-126 (2010)

[91] Osisanwo, F., Akinsola, J., Awodele, O., Hinmikaiye, J., Olakanmi, O. & Akinjobi, J. Supervised machine learning algorithms: classification and comparison. *International Journal Of Computer Trends And Technology (IJCTT).* **48**, 128-138 (2017)