

Utilizing Ensemble Learning for Performance and Power Modeling and Improvement of Parallel Cancer Deep Learning CANDLE Benchmarks

Xingfu Wu

*Argonne National Laboratory
The University of Chicago, USA
Email: xingfu.wu@anl.gov*

Valerie Taylor

*Argonne National Laboratory
The University of Chicago, USA
Email: vtaylor@anl.gov*

Abstract—Machine learning (ML) continues to grow in importance across nearly all domains and is a natural tool in modeling to learn from data. Often a tradeoff exists between a model’s ability to minimize bias and variance. In this paper, we utilize ensemble learning to combine linear, nonlinear, and tree-/rule-based ML methods to cope with the bias-variance tradeoff and result in more accurate models. Hardware performance counter values are correlated with properties of applications that impact performance and power on the underlying system. We use the datasets collected for two parallel cancer deep learning CANDLE benchmarks, NT3 (weak scaling) and P1B2 (strong scaling), to build performance and power models based on hardware performance counters using single-object and multiple-objects ensemble learning to identify the most important counters for improvement. Based on the insights from these models, we improve the performance and energy of P1B2 and NT3 by optimizing the deep learning environments TensorFlow, Keras, Horovod, and Python under the huge page size of 8 MB on the Cray XC40 Theta at Argonne National Laboratory. Experimental results show that ensemble learning not only produces more accurate models but also provides more robust performance counter ranking. We achieve up to 61.15% performance improvement and up to 62.58% energy saving for P1B2 and up to 55.81% performance improvement and up to 52.60% energy saving for NT3 on up to 24,576 cores.

1. Introduction

Energy-efficient scientific applications require insight into how high-performance computing (HPC) system components impact the applications’ power and performance. This insight can result from the development of performance and power models. Currently, HPC systems, especially petaflops supercomputers, consume a huge amount of power; and exascale HPC systems will be similarly constrained. Therefore, monitoring the power consumption of an HPC system is important for power management. Since direct online power measurement at high frequencies is impractical, hardware performance counters have been widely used as effective proxies to estimate power consumption [1], [42], [45]. Hardware performance counter values are correlated with properties of applications that impact

performance and power on the underlying system. In this paper, we use ensemble machine learning, which combines several machine learning models to create a new ensemble model [10], [31] for modeling performance and power based on performance counters and rank performance counters, with the aim of identifying the most important performance counters for application improvement.

Much of the previous work on power modeling and estimation is based on performance counters [7], [9], [17], [25], [45], [1], [5], [15], [28], [29], [35], [42], [46], [51], [53], [54], [57]. These counters are used to monitor system components such as CPU, GPU, memory, disk, and I/O. The values of these performance counters are then correlated with the power consumed by each system component to derive a power model for each system component. Many of these approaches use a small set of performance counters (13 counters or less) for power modeling. In our previous work [28], [54], we developed models of runtime and power based on 40 performance counters and found that the performance counters used for four different models (runtime, node power, CPU power, and memory power) were not the same. For instance, with six scientific applications we found that a total of 37 different performance counters were used for the models. Hence, in this paper we use ensemble machine learning to enhance our modeling methods.

Machine learning (ML) continues to grow in importance across nearly all domains and is a natural tool in modeling to learn from data. It is the process of developing a model in a way that we can understand and quantify the model’s prediction accuracy on future, yet-to-be-seen data [20]. A variety of ML models have been fit to different datasets in different languages. For instance, the machine learning caret package in R [21] provides 238 machine learning methods in R, and the scikit-learn [44] package provides many machine learning methods in Python. ML methods have their own way of learning the relationship between the predictors and the target object. Which method, then, should be selected for the final model?

Generally, prediction errors can be decomposed into two important subcomponents: error due to bias and error due to variance [21]. Error due to bias is the difference between the expected (or average) prediction of a model and the actual value. It measures how far off the model’s predictions are

from the actual value, which provides a sense of how well the model can conform to the underlying structure of the data. On the other hand, error due to variance is defined as the variability of a model prediction for a given data point. Generally, more complex models such as k-nearest neighbor, decision trees, and gradient boosting machines can have very high variance, which leads to overfitting. Simple models such as linear models, which are classical examples of high-bias models, tend not to overfit but to underfit because they are less flexible and rarely capture nonlinear relationships. Understanding how different sources of error lead to bias and variance helps us improve the data-fitting process, resulting in more accurate models. Often a tradeoff must be made between a model’s ability to minimize bias and variance.

Ensemble learning combines several ML models to create a new ensemble model [10], [31]. It is a powerful technique in machine learning, often outperforming other methods with respect to accuracy. However, this comes at the cost of increased algorithmic and model complexity. Should the ensemble model be selected for the final model? Should the ensemble learning provide the robust feature importance? When we consider multiple objects such as runtime, node power, CPU power, and memory power for improvement, how do we utilize multiple-object ensemble learning? In this paper, we address these issues and use ensemble learning to combine linear, nonlinear, and tree-/rule-based ML methods to cope with the bias-variance tradeoff and result in more accurate models.

The CANDLE project [3], [4], [55] focuses on building a single scalable deep neural network that can address three cancer challenge problems: (1) the RAS pathway problem of understanding the molecular basis of key protein interactions in the RAS/RAF pathway presented in 30% of cancers using unsupervised learning; (2) the drug response problem of developing predictive models for drug response to optimize preclinical drug screening and drive precision-medicine-based treatments for cancer patients using supervised learning; and (3) the treatment strategy problem of automating the analysis and extraction of information from millions of cancer patient records to determine optimal cancer treatment strategies using semi-supervised learning. The CANDLE benchmarks [4] implement deep learning architectures that are relevant to these three cancer problems. In our previous work [55], we discussed our parallel methodology and implemented and analyzed the Horovod CANDLE benchmarks (NT3, P1B1, P1B2, and P1B3), focusing on their scalability, performance, and power characteristics with different batch sizes, learning rates, and epochs on both Summit [49] with GPUs at Oak Ridge National Laboratory and Theta [50] with CPUs at Argonne National Laboratory. We identified the data-loading bottlenecks and then improved the performance and energy for better scalability by loading data in chunks without memory limitation. We achieved up to 78.25% in performance improvement and up to 78% in energy saving under strong scaling on up to 384 GPUs and up to 79.5% in performance improvement and up to 77.11% in energy saving under weak scaling on up to 3,072 GPUs on Summit.

We also achieved up to 45.22% performance improvement and up to 41.78% in energy saving under strong scaling on up to 384 nodes on Theta. However, these benchmarks still ran much slower than what we expected on Theta. This result motivated us to identify opportunities to further improve their performance and energy on Theta.

In this paper, we use the datasets collected for the parallel cancer deep learning CANDLE benchmarks NT3 (weak scaling) and P1B2 (strong scaling) to address these issues in performance and power modeling and to achieve improvement using single-object and multiple-object ensemble learning [30], [31], which is a combination of several machine learning methods from the R package caret. We focus on three types of machine learning: linear, nonlinear, and tree-/rule-based regressions. We analyze how a single ML method performs, then investigate how ensemble learning for 15 ML methods performs and compare the results and evaluate models with built-in feature selection. Further, we use unsupervised feature selection methods to confirm that ensemble learning provides more robust model and feature ranking. Then, based on the insights from these models, we improve the performance and energy of P1B2 and NT3 on Theta.

The remainder of this paper is organized as follows. Section 2 presents the modeling and improvement framework using ensemble learning. Section 3 briefly describes the parallel cancer deep learning CANDLE benchmarks NT3 and P1B2. Section 4 discusses performance and power modeling using single-object ensemble learning and investigates the performance counter ranking under different ML methods. Section 5 discusses performance and power modeling using multiple-objects ensemble learning. Section 6 presents performance and energy improvement and the experimental results. Section 7 summarizes our conclusions and briefly discusses possible future work.

2. Modeling and Improvement Framework Using Ensemble Learning

In this section, we propose a modeling and improvement framework using ensemble learning. We discuss how to use ensemble machine learning methods to model performance and power, and we identify the most important performance counters that affect the application performance and power.

Because our problem is a regression problem, we focus on three types of machine learning with a total of 15 methods: linear, nonlinear, and tree-/rule-based regressions with built-in measurements of variable importance from the caret package in R [21]. For instance, multivariate adaptive regression spline and many tree-based models monitor the increase in performance that occurs when adding each variable to the model [20], and linear regression and logistic regression use quantifications based on the model coefficients or statistical measurements (such as t-statistics).

For linear regression, we use five methods: Lasso and elastic-net regularized generalized linear models (glmnet) [14], partial least squares (pls) [20], ridge regression [58],

principal component regression (pcr) [20], and elastic net regression (enet) [12]. For nonlinear regression, we use five methods: k-nearest neighbors (knn) [20], support vector machine with a linear kernel [18], multivariate adaptive regression spline [33], Gaussian process with radial basis function (gaussprRadial) [18], and bayesglm [20]. For tree/rule-based regression, we use five methods: random forests (rf) [24], cubist [22], conditional inference tree (ctree) [16], eXtreme Gradient Boosting (xgbTree) [6], and bagged (bootstrap aggregated) CART (Classification & Regression Trees) (treebag) [20].

To collect these ML models as an ensemble, we use caretEnsemble [31], which is a package for making ensembles of the caret models from R caret [21]. The caretEnsemble package has three primary functions: caretList(), caretEnsemble(), and caretStack(). caretList is a function for fitting many different caret models to the same dataset with the same resampling parameters. It returns a list of caret objects that can be passed to caretEnsemble or caretStack. It has almost the same arguments as train() from the R caret. caretEnsemble uses a generalized linear model (glm) to create a simple linear blend of models. It has two arguments that can be used to specify which models to fit: methodList and tuneList. methodList is a simple character vector of methods that will be fit with the default train parameters, while tuneList can be used to customize the call to each component model. varImp() [21] is used to extract the variable importance from each member of the ensemble, as well as the final ensemble model. caretStack uses a caret model to combine the outputs from several component caret models. It allows to move beyond simple blends of models to use metamodels to create ensemble collections of predictive models. However, it does not support varImp(). In this work, we use caretEnsemble to make an ensemble of the 15 ML methods to build the model and rank the performance counters.

Figure 1 shows the MuMMI modeling and improvement framework. This extends and enhances our previous MuMMI framework [54] by leveraging ensemble learning and feature selection. For an HPC application executed on a power-aware system, we collect runtime, power (node, CPU, and memory), and performance counter data. During the application execution we capture available underived performance counters using any counter measurement tool (PAPI [36], perf_events [37], perfmon2 [38], HPM [40], etc.). All performance counters are normalized by using the total cycles of the execution to create performance event rates for each counter. We then use ML methods to build the models for the metrics and rank the counters based on their variable importance. However, each ML method builds the model and provides the variable importance in distinct ways. Hence, we use ensemble learning to create the ensemble of these ML models, build a more accurate complex model, and provide the robust variable importance for ranking performance counters. Then we can use the counter ranking to identify the most important counters for potential application improvements and/or use the ensemble models to predict the execution time and power. In this work,

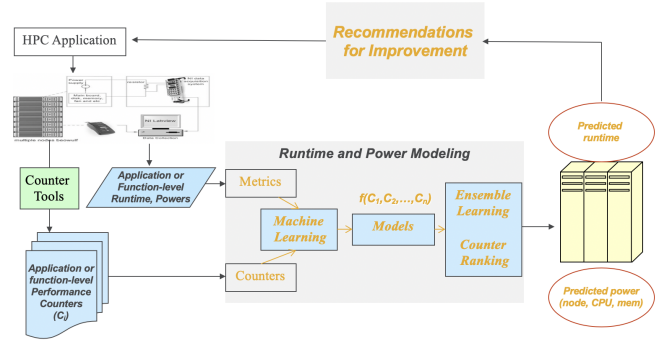


Figure 1. MuMMI Modeling and Improvement Framework

we focus on four metrics: runtime, node power, CPU power, and memory power.

3. Parallel Cancer Deep Learning CANDLER Benchmarks: P1B2 and NT3

In this section, we discuss two parallel cancer deep learning CANDLER benchmarks [4]: NT3 (weak scaling) and P1B2 (strong scaling). The NT3 benchmark is a 1D convolutional network for classifying RNA-seq gene expression profiles into normal or tumor tissue categories. This network follows the classic architecture of convolutional models with multiple 1D convolutional layers interleaved with pooling layers followed by final dense layers. The model is trained on the balanced 700 matched normal-tumor gene expression profile pairs available from the NCI Genomic Data Commons and acts as a quality control check for synthetically generated gene expression profiles. The full dataset of expression features contains 60,483 float columns transformed from RNA-seq FPKM-UQ values [4] that map to a column that contains the integer 0—1. The training data size for this benchmark is 597 MB, and the test data size is 150 MB. The number of epochs is one per node for the weak-scaling study. The batch size is 20 (default); and the total training samples are 1,120, with 60,483 features per sample. The optimizer is sgd (stochastic gradient descent).

The P1B2 benchmark is an MLP (multiLayer perceptron) network with regularization and five layers. Given patient somatic SNP data, it builds a deep learning network that can classify the cancer type based on sparse input data and evaluate the information content and predictive value in a molecular assay with auxiliary learning tasks. The training data size is 162 MB, and the test data size is 55 MB. The number of epochs is 384 for the strong-scaling study; the batch size is 60 (default); and the total training samples are 2,700 with 28,204 features per sample. The optimizer is rmsprop (root mean square propagation).

In this work, we use the two benchmarks P1B2 and NT3 to collect the available 26 performance counters using PyPAPI [41] and performance and power data on the Cray XC40 Theta [50]. Each node of Theta has 64 cores. The dataset p1b2.csv for P1B2 has 144 configurations, which

include 37 variables per configuration (such as application name and system name), 12 numbers of nodes (6 to 384), number of cores (384 to 24,576), learning rates, and 12 different batch sizes, number of epochs (384). The dataset nt3.csv for NT3 has 105 configurations, which include the same 37 variables per configuration, 15 numbers of nodes (1 to 384), number of cores (64 to 24,576), learning rates, 7 different batch sizes, and number of epochs (one per node) Both datasets include 26 performance counters and four metrics (runtime, node power, CPU power, and memory power). The number of cores used for PIB2 and NT3 is up to 24,576 (384 nodes).

The 26 performance counters are TOT_CYC (total cycles), TOT_INS (total instructions completed), BR_CN (conditional branch instructions), BR_NTK (conditional branch instructions not taken), L1_TCM (L1 total cache misses), L1_LDM (L1 load misses), L1_DCM (L1 data cache misses), L1_ICA (L1 instruction cache accesses), L1_ICH (L1 instruction cache hits), L1_ICM (L1 instruction cache misses), L2_TCM (L2 total cache misses), L2_TCA (L2 total cache assesses), L2_TCH (L2 total cache hits), L2_LDM (L2 load misses), TLB_DM (data translation lookaside buffer misses), BR_MSP (conditional branch instructions mispredicted), RES_STL (cycles stalled at any resource), SR_INS (store instructions), LD_INS (load instructions), BR_TKN (conditional branch instructions taken), BR_INS (branch instructions), L1_DCA (L1 data cache accesses), LST_INS (Load/store instructions completed), REF_CYC (reference clock cycles), STL_ICY (cycles with no instructions issued), and BR_UCN (unconditional branch instructions). Then TOT_CYC is used to normalize all the performance counters. We used the datasets to analyze the pairwise correlations among 25 hardware performance counters and four target objects as follows.

For NT3, we used weak scaling to collect the dataset nt3.csv with 105 configurations. For the same batch size, we expect that the runtime for different configurations is similar. Figure 2 shows the counter correlation matrix for 25 counters. This indicates that most counters are not correlated each other. Figure 3 presents the object correlation matrix for the four metrics. It indicates that runtime is inversely correlated with power. However, node power is highly correlated with CPU power because node power includes the CPU power, and CPU power is poorly correlated with memory power.

For PIB2, we used strong scaling to collect the dataset plb2.csv with 144 configurations. For the same batch size, we expect that the runtime for different configurations is distinct and that it decreases with increasing numbers of nodes because the number of epochs per node is decreased. Figure 4 shows the counter correlation matrix for 25 counters. This indicates that most counters have some correlation each other except REF_CYC. Figure 5 presents the metrics correlation matrix for four metrics. It indicates that runtime is correlated with power. As with NT3, however, node power is highly correlated with CPU power because node power includes the CPU power, and CPU power is poorly correlated with memory power.

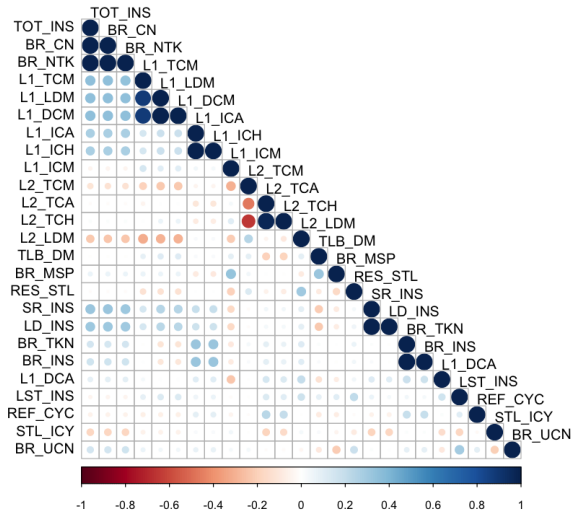


Figure 2. Counter correlation matrix for NT3

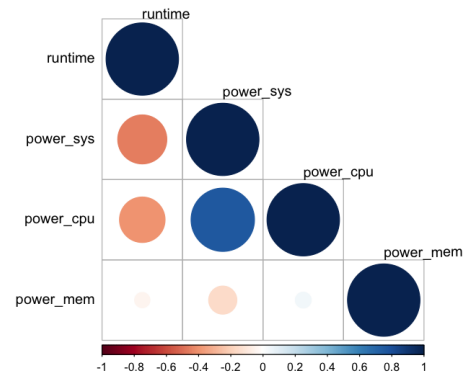


Figure 3. Object correlation matrix for NT3

4. Performance and Power Modeling Using Ensemble Learning

To make sure we use the same training and test sets, we use the `set.seed(3456)` in our R codes for all experiments so that creating the random objects can be reproduced. We define the relationship between the object metric (runtime, node power, CPU power, or memory power) and variables (25 performance counters) as follows: $\text{metric} \sim \text{TOT_INS} + \text{BR_CN} + \text{BR_NTK} + \text{L1_TCM} + \text{L1_LDM} + \text{L1_DCM} + \text{L1_ICA} + \text{L1_ICH} + \text{L1_ICM} + \text{L2_TCM} + \text{L2_TCA} +$

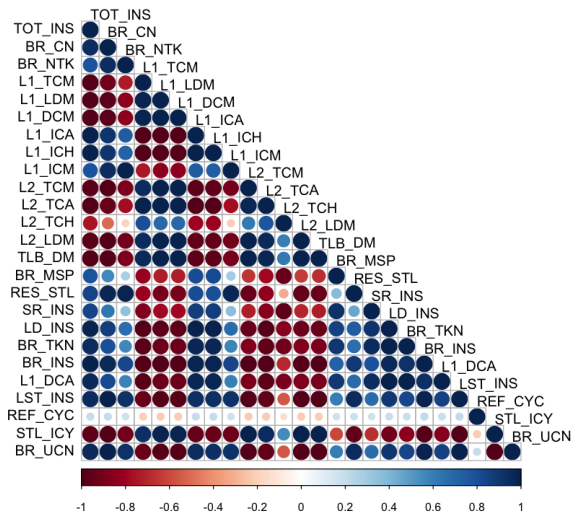


Figure 4. Counter correlation matrix for NT3

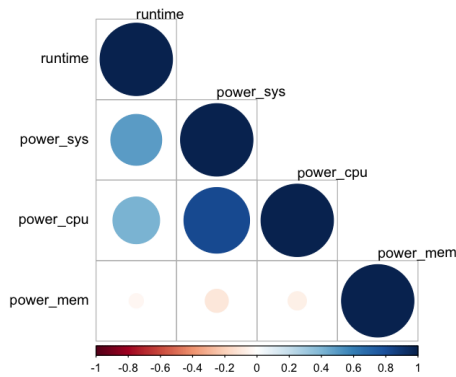


Figure 5. Object correlation matrix for NT3

L2_TCH + L2_LDM + TLB_DM + BR_MSP + RES_STL + SR_INS + LD_INS + BR_TKN + BR_INS + L1_DCA + LST_INS + REF_CYC + STL_ICY + BR_UCN.

When using models to predict a numeric outcome, some measure of accuracy is typically used to evaluate the effectiveness of the model. For instance, the most common method for characterizing a model’s predictive capabilities is to use the root mean square error (RMSE). RMSE is a function of the model residuals, which are computed as the observed values minus the predicted value. It is interpreted either as how far, on average, the residuals are from zero

or as the average distance between the observed values and the predicted values. Another common metric is R^2 , the coefficient of determination, which can be interpreted as the proportion of the information in the data that is explained by the model. It is the squared correlation coefficient between the observed and predicted values in the simplest version. It is a measure of correlation, not accuracy. Therefore, in this paper we use the metric RMSE to evaluate the effectiveness of the model using different ML methods and ensemble learning.

We choose 15 machine learning methods [21] from the R caret package: tree-/rule-based group – random forest (rf), cubist (cubist), eXtreme Gradient Boosting (xgbTree), conditional inference tree (ctree), and treebag; nonlinear group – k-nearest neighbors (knn), support vector machines with linear kernel (svmLinear), Gaussian process with radial basis function (gaussprRadial), multivariate adaptive regression spline (earth), and Bayesian generalized linear model (bayesglm); linear group – Lasso and elastic-net regularized generalized linear models (glmnet), partial least squares (pls), ridge regression (ridge), elastic net regression (enet), and principal component regression (pcr) because their RMSE is very close to each other in the same group. Then we use caretEnsemble to create the ensemble of the 15 ML methods and build the model and rank the performance counters.

4.1. Performance and Power Modeling

In this section, we use the 15 machine learning methods and caretEnsemble [31] to model performance (runtime) and node power of NT3 and P1B2 using the same training and test sets based on the 80/20% rule.

Table 1 shows the ensemble model for P1B2 that resulted from the ensemble of the following models: cubist, treebag, xgbTree, ctree, rf, glmnet, pls, ridge, pcr, enet, knn, svmLinear, earth, gaussprRadial, and bayesglm. The resulting RMSE for the ensemble model is 200.14 in performance and 3.70 in node power. They are the most accurate among these models.

Similarly, Table 2 shows the ensemble model for NT3 that resulted from the ensemble of the following models: cubist, treebag, xgbTree, ctree, rf, glmnet, pls, ridge, pcr, enet, knn, svmLinear, earth, gaussprRadial, and bayesglm. The resulting RMSE for the ensemble model is 154 in performance and 4.24 in node power. The ensemble model for runtime is the most accurate among these models; however, the power model using rf results in the smallest RMSE, 4.22. This is an exception. Overall, ensemble learning results in the most accurate performance and power models for P1B2 and NT3 in most cases.

4.2. Performance Counter Ranking

In this section, we use P1B2 to explore performance counter ranking using model-based feature selection for 15 ML methods and ensemble learning and using unsupervised feature selection. We compare them to identify the most

TABLE 1. RMSE OF DIFFERENT ML MODELS FOR P1B2

Method		RMSE	
		Runtime	Power
Tree-based	cubist	278.28	3.97
	treebag	339.93	4.02
	xgbTree	319.15	3.86
	ctree	364.78	4.20
	rf	315.61	3.90
Linear	glmnet	261.76	3.90
	pls	289.72	3.88
	ridge	259.13	3.97
	pcr	349.05	3.87
	enet	257.29	3.90
Nonlinear	knn	359.69	4.11
	svmLinear	294.53	3.98
	earth	224.95	4.17
	gaussprRadial	351.78	3.98
	bayesglm	266.41	4.17
Ensemble		200.14	3.70

TABLE 2. RMSE OF DIFFERENT ML MODELS FOR NT3

Method		RMSE	
		Runtime	Power
Tree-based	cubist	193.41	4.50
	treebag	175.92	4.36
	xgbTree	183.68	4.41
	ctree	167.53	4.48
	rf	170.00	4.22
Linear	glmnet	187.91	4.34
	pls	199.55	4.61
	ridge	210.19	5.11
	pcr	176.32	4.38
	enet	170.42	4.27
Nonlinear	knn	177.30	4.27
	svmLinear	188.54	5.65
	earth	186.63	4.43
	gaussprRadial	177.22	4.24
	bayesglm	213.16	5.47
Ensemble		154.00	4.24

important performance counters that impact the performance and node power.

4.2.1. Model-Based (Supervised) Feature Selection. In this section, we use 15 machine learning methods and caretEnsemble to explore performance counter ranking. varImp() is used to extract the variable importance from each member of the ensemble, as well as the final ensemble model. Then we sum the variable importance values for counters to calculate the percentage for each counter to rank them.

Table 3 presents the top 6 counters in each performance model using the different ML methods. A total of 15 different counters are listed in the table. We observe that 7 of the 15 ML methods listed (ridge, pcr, enet, knn, svmLinear, gaussprRadial, and bayesglm) have the same

TABLE 3. PERFORMANCE COUNTER RANKING FOR PERFORMANCE MODELS

Method	Top 6 performance counters (from the highest to the lowest)						
	BR_CN	L1_ICM	LST_INS	L1_LDM	STL_ICY	TOT_INS	
Tree-based	cubist	BR_CN	L1_ICM	LST_INS	L1_LDM	STL_ICY	TOT_INS
	treebag	BR_CN	LST_INS	RES_STL	BR_UCN	BR_NTK	STL_ICY
	xgbTree	LST_INS	BR_CN	L2_TCH	BR_NTK	TOT_INS	LST_INS
	ctree	LST_INS	RES_STL	BR_CN	BR_UCN	STL_ICY	BR_NTK
	rf	RES_STL	BR_CN	LIS_INS	BR_UCN	BR_NTK	BR_INS
Linear	glmnet	L1_ICM	BR_NTK	L2_LDM	L1_TCM	BR_CN	L1_LDM
	pls	L1_ICM	RES_STL	LST_INS	BR_NTK	BR_MSP	BR_CN
	ridge	LST_INS	RES_STL	BR_CN	BR_UCN	STL_ICY	BR_NTK
	pcr	LST_INS	RES_STL	BR_CN	BR_UCN	STL_ICY	BR_NTK
	enet	LST_INS	RES_STL	BR_CN	BR_UCN	STL_ICY	BR_NTK
Nonlinear	knn	LST_INS	RES_STL	BR_CN	BR_UCN	STL_ICY	BR_NTK
	svmLinear	LST_INS	RES_STL	BR_CN	BR_UCN	STL_ICY	BR_NTK
	earth	BR_CN	L2_LDM	BR_TKN	STL_ICY	BR_NTK	BR_MSP
	gaussprRadial	LST_INS	RES_STL	BR_CN	BR_UCN	STL_ICY	BR_NTK
	bayesglm	LST_INS	RES_STL	BR_CN	BR_UCN	STL_ICY	BR_NTK
Ensemble (counter percentage)	BR_CN (11.90%)	RES_STL (7.09%)	LST_INS (6.39%)	L2_LDM (5.93%)	STL_ICY (5.28%)	L1_ICM (4.73%)	

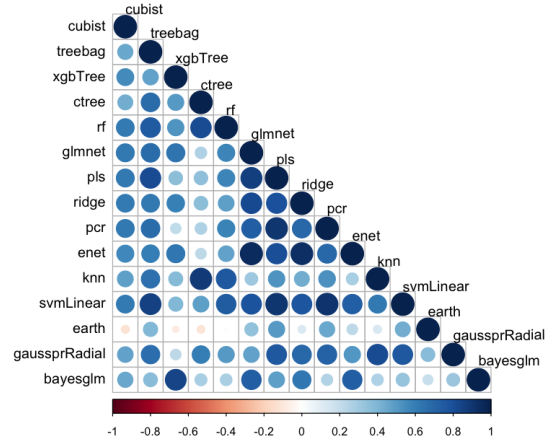


Figure 6. ML performance model correlation matrix for P1B2

performance counter ranking order. However, they are not highly correlated each other, as shown in Figure 6. The others have different performance counter rankings. We note that ensemble learning provides the overall counter ranking, which is different from all 15 ML methods and includes the top counter from each ML model. The counter percentage is the counter importance values divided by the summation of all counter importance values. We note that BR_CN, RES_STL, and LST_INS are the top 3 counters in performance modeling provided by ensemble learning.

Table 4 presents the top 6 counters in each node power model using different ML methods, which are different from that in Table 3. A total of 21 different counters are listed in the table. We observe that the same 7 ML methods (ridge, pcr, enet, knn, svmLinear, gaussprRadial, and bayesglm) have the same performance counter ranking order and that ridge and gaussprRadial have a high correlation with the others, as shown in Figure 7. We note that L1_ICM, LST_INS, and L2_LDM are the top 3 counters in power modeling

TABLE 4. PERFORMANCE COUNTER RANKING FOR POWER MODELS

Method		Top 6 performance counters (from the highest to the lowest)						
Tree-based	cubist	L2_LDM	L2_TCA	L1_ICA	LST_INS	L2_TCA	L1_ICH	
	treebag	L1_ICM	TLB_DM	L2_LDM	RES_STL	BR_NTK	LST_INS	
	xbgTree	L1_DCA	RES_STL	BR_NTK	BR_UCN	L2_TCM	RFE_CYC	
	ctree	L1_ICA	L1_ICH	TOT_INS	BR_INS	L1_LDM	L1_DCM	
	rf	TLB_DM	L1_ICM	L1_DCA	RES_STL	L2_TCH	BR_INS	
Linear	glmnet	L1_ICH	LST_INS	LD_INS				
	pls	L1_ICM	L1_ICA	TOT_INS	L1_TCM	LD_INS	BR_INS	
	ridge	L1_ICA	L1_ICH	TOT_INS	BR_INS	L1_DCM	L1_LDM	
	pcr	L1_ICA	L1_ICH	TOT_INS	BR_INS	L1_DCM	L1_LDM	
	enet	L1_ICA	L1_ICH	TOT_INS	BR_INS	L1_DCM	L1_LDM	
Nonlinear	knn	L1_ICA	L1_ICH	TOT_INS	BR_INS	L1_DCM	L1_LDM	
	svmLinear	L1_ICA	L1_ICH	TOT_INS	BR_INS	L1_DCM	L1_LDM	
	earth	L2_LDM						
	gaussprRadial	L1_ICA	L1_ICH	TOT_INS	BR_INS	L1_DCM	L1_LDM	
	bayesglm	L1_ICA	L1_ICH	TOT_INS	BR_INS	L1_DCM	L1_LDM	
Ensemble (counter percentage)		L1_ICH (9.78%)	LST_INS (5.17%)	L2_LDM (4.95%)	LD_INS (4.33%)	L1_ICA (4.32%)	TLB_DM (4.09%)	

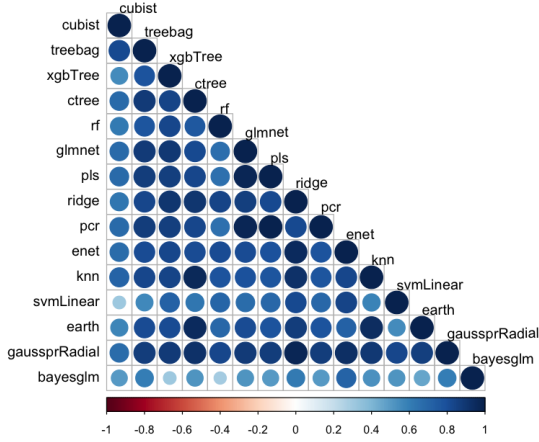


Figure 7. ML power model correlation matrix for P1B2

provided by ensemble learning.

4.2.2. Unsupervised Feature Selection. Next, we use unsupervised feature selection and compare with the model-based feature selection. For unsupervised feature selection, we use two types of methods [20]: (1) wrapper methods – recursive feature elimination (RFE), genetic algorithm (GA), and simulated annealing (SA); and (2) filter methods – stepwise and selection by filter (SBF). Then we compare them with the results from the ensemble learning. We still use P1B2 as an example.

For the RFE method, it supports random forest (rfFuncs), linear regression (lmFuncs), and bagged tree (treebagFuncs). For the GA method, it supports random forest (rfGA) and bagged tree (treebagGA). For the SA method, it supports random forests (rfSA) and bagged tree (treebagSA). For the SBF method, it supports random forests (rfSBF), linear regression (lmSBF), and bagged tree (treebagSBF). Stepwise forward and backward selection use glm.

TABLE 5. PERFORMANCE COUNTER RANKING FOR PERFORMANCE MODELS

Method		Top 6 performance counters (from the highest to the lowest)							
Wrapper	RFE	rf	BR_CN	RES_STL	LST_INS	L1_ICM	BR_UCN	BR_NTK	
		glm	REF_CYC	BR_CN	L1_ICM	L1_ICA	L1_ICH	L2_LDM	
		treebag	BR_CN	LST_INS	RES_STL	BR_UCN	BR_NTK	TOT_INS	
	GA	rf	RES_STL	TOT_INS	L1_ICM	STL_ICY	BR_MSP	L2_LDM	
		treebag	BR_CN	RES_STL	TOT_INS	BR_UCN	L2_TCM	L2_LDM	
		SA	rf	TOT_INS	L1_ICM	TLB_DM	L2_TCM	BR_MSP	L1_DCM
Filter	SBF	treebag	RES_STL	BR_UCN	L2_TCM	TOT_INS	TLB_DM	L1_DCM	
		rf	BR_CN	RES_STL	L1_ICM	LST_INS	BR_MSP	TOT_INS	
		glm	BR_CN	L1_ICH	L1_ICA	L1_ICM	L2_LDM	TOT_INS	
	Step Wise	treebag	BR_CN	LST_INS	RES_STL	BR_UCN	TOT_INS	BR_NTK	
		glm	BR_CN	L1_ICM	TOT_INS	L1_TCM	L1_LDM	L2_LDM	

TABLE 6. PERFORMANCE COUNTER RANKING FOR POWER MODELS

Method		Top 6 performance counters (from the highest to the lowest)						
Wrapper	RFE	rf	TLB_DM	L1_ICM	L1_DCA	TOT_INS	RES_STL	L2_TCH
		glm	REF_CYC	L2_LDM	L1_LDM	L2_TCM	L1_ICH	L1_ICA
		treebag	L1_ICM	TLB_DM	BR_NTK	RES_STL	BR_CN	L1_LDM
	GA	rf	TLB_DM	SR_INS	TOT_INS	BR_NTK	L1_ICA	L1_DCA
		treebag	TLB_DM	BR_UCN	LST_INS	BR_TKN	BR_NTK	BR_MSP
		SA	rf	TLB_DM	TOT_INS	L1_ICM	L1_ICH	L1_TCM
Filter	SBF	treebag	L1_DCM	TLB_DM	L2_TCM	TOT_INS	L1_ICM	BR_MSP
		rf	TLB_DM	L1_ICM	L1_DCA	L2_TCH	SR_INS	RES_STL
		glm	L2_LDM	RES_STL	L1_ICH	L1_ICA	L2_TCM	LST_INS
	Step Wise	treebag	TLB_DM	STL_ICY	LST_INS	RES_STL	L1_ICM	BR_CN
		glm	L1_ICH					

Table 5 shows the performance counter ranking for performance models using unsupervised feature selection. It confirms that the top two counters are BR_CN and RES_STL provided by ensemble learning in Table 3. The other counters also occur in Table 5.

Table 6 shows the performance counter ranking for node power models using unsupervised feature selection. It shows that TLB_DM is the top counter in the top 6 counters in Table 4. Two of the top three counters, L1_ICH, LST_INS, and L2_LDM, provided by ensemble learning in Table 4 are the top counters in Table 6. Overall, compared with each ML method, we find that ensemble learning provides more robust performance count ranking.

4.3. Discussion

In summary, for P1B2, we used unsupervised feature selection methods to confirm that ensemble learning provides more robust performance count ranking than each individual ML method does. When we consider runtime and node power for application improvement for P1B2, we need to focus on the counters BR_CN, L1_ICH, and TLB_DM. Similarly, when we apply the same methodology to NT3, we find that the counters L1_ICM, L2_TCH, and TLB_DM are the focal counters for application improvement.

5. Multiple-Objects Ensemble Learning

In the preceding sections, we explored single-object (univariate) ensemble learning [31] to combine different ML methods and result in more accurate models and provide more robust performance counter ranking. These methods target only a single metric, either runtime or node power. In this section, we discuss multiple-objects (multivariate) ensemble learning using a multivariate tree boosting method mvtboost [30] to model performance and power and to rank the performance counters based on multiple objects.

TABLE 7. PERFORMANCE COUNTER RANKING FOR P1B2 USING ENSEMBLE LEARNING

Metric	Top 6 performance counters (from the highest to the lowest)					
runtime	BR_CN	RES_STL	TOT_INS	LD_INS	L1_ICM	BR_NTK
node power	TLB_DM	L1_ICM	REF_CYC	BR_INS	RES_STL	LST_INS
CPU power	L1_ICM	LST_INS	BR_INS	L2_LDM	L1_LDM	SR_INS
Memory power	L1_ICM	REF_CYC	BR_INS	L1_TCM	TLB_DM	BR_NTK

TABLE 8. PERFORMANCE COUNTER RANKING FOR NT3 USING ENSEMBLE LEARNING

Metric	Top 6 performance counters (from the highest to the lowest)					
runtime	L1_ICM	RES_STL	L1_LDM	L1_TCM	TOT_INS	L2_TCA
node power	L1_ICM	STL_ICY	BR_UCN	SR_INS	BR_INS	L2_TCA
CPU power	BR_MSP	L2_TCH	L1_TCM	L1_ICM	L2_TCM	REF_CYC
Memory power	STL_ICY	L1_ICM	REF_CYC	TOT_INS	RES_STL	TLB_DM

Boosted decision tree ensembles such as gradient boosting machine (gbm) [11] are powerful ensemble learning algorithms, allowing dependent variables to be nonlinear functions of predictors. mvtboost (multivariate tree boosting) extends gbm to multivariate, continuous object variables by fitting a separate univariate model of a common set of predictors to each object variable. This accounts for covariance in the object variables as in seemingly unrelated regression. This joint analysis of several object variables can be informative when we consider the four metrics runtime, node power, CPU power, and memory power for application improvement. In this section, we use a mvtboost method with 1,000 trees, a learning rate of 0.01, and a tree depth of 3 to model the performance and power of P1B2 and NT3 and analyze the impacts of different performance counters. We then compare the performance counter ranking with what we found in the preceding section.

Table 7 shows the performance counter ranking for P1B2 using multiple-objects ensemble learning. When we consider runtime, node power, CPU power, and memory power for the application improvement, we need to focus on the top counters BR_CN, TLB_DM, and L1_ICM, similar to the counters BR_CN, L1_ICM, and TLB_DM found in the preceding section. This further confirms that we need to focus on BR_CN, L1 cache, and TLB_DM to improve the application performance and power.

Table 8 shows the performance counter ranking for NT3 using multiple-objects ensemble learning. When we consider runtime, node power, CPU power, and memory power for the application improvement, we need to focus on the top counters L1_ICM, BR_MSP, and STL_ICY. In the preceding section, when we consider runtime and node power, we need to focus on L1_ICM, L2_TCH, and TLB_DM. Overall, we need to focus on L1 cache, BR_MSP, STL_ICY, and TLB_DM to improve the application performance and power.

6. Performance and Energy improvement

The CANDLE benchmarks P1B2 and NT3 are implemented in Python by using Keras, TensorFlow, and Horovod. These Python codes, like other scripting languages, do not have compiler optimization support and instead rely on the library, resource, and environment settings for better

performance. We can utilize what we learned from ensemble learning in preceding sections to identify better resource and environment settings for performance improvement. In this section, based on the insights from important performance counters we identified, we improve the performance and energy of these benchmarks by fine-tuning the applications, underlying library, software, and/or system.

For P1B2 and NT3 executed on Theta, the dominant phase is model training [55], which is the function call from TensorFlow. From our analysis in the preceding sections, the most important counters indicate that TensorFlow is the optimization target with huge page sizes. The Cray XC40 Theta [50] supports the huge page sizes of 2 MB, 4 MB, 8 MB, 16 MB, 32 MB, 64 MB, 128 MB, 256 MB, 512 MB, 1 GB, and 2 GB. The baseline environment used in the work includes Python3.7.5, TensorFlow 1.15, Keras 2.3, and Horovod 0.18 built from the packages by using pip. To optimize TensorFlow with the huge page sizes, we have to rebuild from the source codes for TensorFlow 1.15, Keras 2.3, and Horovod 0.19 based on Intel Python 3.6.8 and MKL-DNN with page sizes of 2 MB, 8 MB, 32 MB, 128 MB, and 1 GB. Then we use the benchmarks under these different environments to evaluate performance and energy. After we identify which page size results in the best performance, we further improve the application performance and energy.

6.1. Impacts of Different Huge Page Sizes

In this section, we use P1B2 and NT3 to evaluate the impacts of different huge page sizes of 2 MB, 8 MB, 32 MB, 128 MB, and 1 GB. For P1B2, we use a strong-scaling study (384 epochs in total) with the batch size of 60 to analyze the improvement. For NT3, we use a weak-scaling study (one epoch per node) with a batch size of 20 to analyze the improvement.

Table 9 shows performance and energy improvement for NT3 under different huge page sizes, where "2MB" stands for the huge page size of 2 MB. We rebuilt the deep learning environment from the source codes for TensorFlow 1.15, Keras 2.3, and Horovod 0.19 based on Intel Python 3.6.8 and MKL-DNN. Table 10 shows performance and energy improvement for P1B2 under different huge page sizes. For both benchmarks, as we showed in [55], the modeling training and data loading are the two dominant phases. Compared with the baseline, we find that building TensorFlow from source codes based on the latest MKL-DNN improves the model training performance significantly, and the huge page size of 8 MB outperforms the others in most cases. However, the data-loading time under the baseline with Python 3.7.5 (around 115 s for NT3, around 105 s for P1B2) is much smaller than that (around 160 s for NT3, around 125 s for P1B2) under the different huge page sizes with Python 3.6.8. We used the improved data loading methods from [55] for both cases. Basically, the data loading calls the Pandas function read_csv() to load the data. For all cases, the installed Pandas versions are the same (version 0.25.2). The difference is the Python version. In the

TABLE 9. PERFORMANCE AND ENERGY IMPROVEMENT FOR NT3

#Nodes	Method	Runtime (s)	Node Power (W)	Energy Saving
1	Baseline	792.77	147.97	---
	2MB	395.05	159.33	-46.34%
	8MB	392.79	152.49	-48.94%
	32MB	394.45	153.68	-48.32%
	128MB	390.23	160.72	-46.53%
	1GB	395.35	155.13	-47.72%
2	Baseline	813.07	147.52	---
	2MB	429.40	156.26	-44.06%
	8MB	426.89	159.68	-43.17%
	32MB	428.50	152.63	-45.47%
	128MB	434.99	154.57	-43.94%
	1GB	429.50	154.25	-44.77%
4	Baseline	833.28	147.18	---
	2MB	527.73	148.02	-36.31%
	8MB	501.61	156.88	-35.84%
	32MB	518.16	147.79	-37.56%
	128MB	500.53	154.05	-37.13%
	1GB	516.81	156.82	-33.92%
8	Baseline	841.18	143.63	---
	2MB	593.24	145.08	-28.76%
	8MB	559.27	146.56	-32.16%
	32MB	591.03	149.24	-26.99%
	128MG	583.01	148.26	-28.46%
	1GB	555.73	147.37	-32.21%

next section, we investigate this issue to further improve the performance and energy.

XLA (Accelerated Linear Algebra) [56] is a domain-specific compiler for linear algebra that can accelerate TensorFlow models with potentially no source code changes. It compiles subgraphs to reduce the execution time of short-lived operations to eliminate overhead from the TensorFlow runtime, fuses pipelined operations to reduce memory overhead, and specializes to known tensor shapes to allow for more aggressive constant propagation. It analyzes and schedules memory usage, in principle eliminating many intermediate storage buffers. As described in [56], XLA may improve the execution speed and memory usage on GPUs. However, it does not support CPUs well, as we experienced on Theta. Table 11 shows the XLA impacts using 8 nodes on Theta. For NT3 and PIB2, we enable XLA to rebuild TensorFlow to run them with/without enabling autoclustering by setting the TF_XLA_FLAGS environment variable as “-tf_xla_auto_jit=2 -tf_xla_cpu_global_jit” under the huge page size of 8 MB. Enabling XLA caused runtime to increase significantly; however, it results in the average node power decrease because of the XLA overhead. This is the case for all other experiments as well.

6.2. Further Improvement

For further improvement under the huge page size of 8 MB, we also rebuilt Python3.7.5, TensorFlow 1.15, Keras 2.3, and Horovod 0.19 from the source codes based on the latest Intel MKL-DNN. We then evaluated the performance

TABLE 10. PERFORMANCE AND ENERGY IMPROVEMENT FOR PIB2

#Nodes	Method	Runtime (s)	Node Power (W)	Energy Saving
1	Baseline	792.77	147.97	---
	2MB	395.05	159.33	-46.34%
	8MB	392.79	152.49	-48.94%
	32MB	394.45	153.68	-48.32%
	128MB	390.23	160.72	-46.53%
	1GB	395.35	155.13	-47.72%
2	Baseline	813.07	147.52	---
	2MB	429.40	156.26	-44.06%
	8MB	426.89	159.68	-43.17%
	32MB	428.50	152.63	-45.47%
	128MB	434.99	154.57	-43.94%
	1GB	429.50	154.25	-44.77%
4	Baseline	833.28	147.18	---
	2MB	527.73	148.02	-36.31%
	8MB	501.61	156.88	-35.84%
	32MB	518.16	147.79	-37.56%
	128MB	500.53	154.05	-37.13%
	1GB	516.81	156.82	-33.92%
8	Baseline	841.18	143.63	---
	2MB	593.24	145.08	-28.76%
	8MB	559.27	146.56	-32.16%
	32MB	591.03	149.24	-26.99%
	128MG	583.01	148.26	-28.46%
	1GB	555.73	147.37	-32.21%

TABLE 11. XLA IMPACTS USING 8 NODES UNDER THE HUGE PAGE SIZE OF 8 MB ON THETA

Application	Method	Runtime (s)	Node Power (W)
NT3	No XLA	559.27	146.56
	XLA	679.22	138.05
PIB2	No XLA	899.63	138.99
	XLA	979.13	135.44

and energy and compared them with “Baseline” and “2MB”, as described in the preceding section.

Table 12 shows the performance and energy improvement for NT3 on up to 384 nodes (24,576 cores) under three different deep learning environments. We note that for the weak-scaling case study of NT3, the workload per node is the same (1 epoch per node), and we focus on improving the TensorFlow training performance. With an increase in the number of nodes, the overall energy saving percentage decreases because of the increase of Horovod communication overhead. This overhead also results in a decrease in the average node power. Overall, we use the further improvement environment to achieve up to 55.81% performance improvement and up to 52.60% energy saving on up to 24,576 cores. Compared with the “2MB” and the baseline, we further improve not only the model training time but also the data-loading time. Thus, we achieve much better energy saving.

Table 13 shows the performance and energy improvement for PIB2 on up to 192 nodes (12,288 cores) under three different deep learning environments. We note that for the strong-scaling case study of PIB2, the workload per node decreases with increasing numbers of nodes. We focus

TABLE 12. PERFORMANCE AND ENERGY IMPROVEMENT FOR NT3 ON UP TO 24,576 CORES

#Nodes	Method	Runtime (s)	Node Power (W)	Energy Saving
1	Baseline	792.77	147.97	---
	2MB	395.05	159.33	-46.34%
	Further	350.31	158.73	-52.60%
2	Baseline	813.07	147.52	---
	2MB	429.40	156.26	-44.06%
	Further	380.24	153.62	-51.30%
4	Baseline	833.28	147.18	---
	2MB	527.73	148.02	-36.31%
	Further	440.92	153.04	-44.98%
8	Baseline	841.18	143.63	---
	2MB	593.24	145.08	-28.76%
	Further	539.57	145.62	-34.97%
32	Baseline	862.35	142.16	---
	2MB	659.6	149.02	-19.82%
	Further	573.49	145.68	-31.85%
64	Baseline	863.49	140.49	---
	2MB	663.56	148.79	-18.61%
	Further	620.98	145.25	-25.65%
128	Baseline	886.42	142.12	---
	2MB	728.54	147.15	-14.90%
	Further	618.87	140.63	-30.92%
192	Baseline	939.45	140.13	---
	2MB	787.25	142.48	-14.80%
	Further	747.02	137.18	-22.16%
384	Baseline	964.8	143.81	---
	2MB	877.43	142.95	-9.60%
	Further	794.57	140.63	-19.47%

TABLE 13. PERFORMANCE AND ENERGY IMPROVEMENT FOR PIB2

#Nodes	Method	Runtime (s)	Node Power (W)	Energy Saving
6	Baseline	3410.52	136.68	---
	2MB	1502.66	139.03	-55.18%
	Further	1369.7	135.78	-60.10%
8	Baseline	2329.38	139.65	---
	2MB	987.86	140.47	-57.34%
	Further	905.08	134.51	-62.58%
16	Baseline	1269.04	130.04	---
	2MB	721.95	129.59	-43.31%
	Further	640.91	130.18	-49.44%
32	Baseline	858.32	129.37	---
	2MB	444.58	130.86	-47.61%
	Further	371.8	132.09	-55.77%
64	Baseline	543.17	124.45	---
	2MB	416.61	131.38	-19.03%
	Further	353.54	130.31	-31.85%
128	Baseline	369.93	128.29	---
	2MB	345.45	127.31	-7.33%
	Further	299.05	125.89	-20.67%
192	Baseline	328.27	128.70	---
	2MB	317.41	126.96	-4.62%
	Further	264.45	126.38	-20.89%

on improving the TensorFlow training performance. With the increase in the number of nodes, the overall energy saving percentage decreases because of the decrease in the workload per node and the increase in the Horovod communication overhead. This also results in a decrease in the average node power. Overall, we use the further improvement environment to achieve up to 61.15% performance improvement and up to 62.58% energy saving on up to 12,288 cores.

7. Conclusions

We used the datasets collected for the two benchmarks NT3 (weak scaling) and PIB2 (strong scaling) to build performance and power models based on hardware performance counters using single-object and multiple-object ensemble learning. We utilized ensemble learning to combine linear, nonlinear, and tree-/rule-based ML methods to cope with the bias-variance tradeoff and result in more accurate models; and we ranked the performance counters to identify the most important counters for improvement hints. Based on the insights from these models, we improved the performance and energy of PIB2 and NT3 by optimizing the deep learning environments TensorFlow, Keras, Horovod, and Python under the huge page size of 8 MB on Theta. Experimental results show that ensemble learning not only leads to more accurate models but also provides more robust performance counter ranking. We achieved up to 61.15% performance improvement and up to 62.58% energy saving for PIB2, and up to 55.81% performance improvement and up to 52.60% energy saving for NT3 on up to 24,576 cores on Cray XC40 Theta.

Overall, ensemble learning provides a broad, robust view of learning from the data. Applying ensemble learning to science should help better identify key factors that impact science or decision-making [PR06]. We learned from this work that all HPC systems support the number of huge page sizes to accelerate scientific applications with recompiling requirement; however, there is no huge page size support for the ML and deep learning applications. For the deep learning applications and others using script languages such as Python, the huge page support should accelerate the time-consuming training process as well. Which huge page size will result in the best performance, however, depends on the application characteristics and the underlying systems.

Acknowledgments

This work was supported in part by Laboratory Directed Research and Development (LDRD) funding from Argonne National Laboratory, provided by the Director, Office of Science, of the U.S. Department of Energy under contract DE-AC02-06CH11357; in part by the U.S. Department of Energy, Office of Science, under contract DE-AC02-06CH11357; and in part by NSF grant CCF-1801856. We acknowledge the Argonne Leadership Computing Facility (ALCF) for use of Theta under the DOE INCITE project

CANDLE and ALCF project EE-ECP. We acknowledge Huihuo Zheng from ALCF for the TensorFlow and Horovod environments on Theta.

References

- [1] W. Bircher, and L. John, Complete System Power Estimation Using Processor Performance Events. *IEEE Tran. on Comp.*, 61(4), 2012.
- [2] R. Bertran, Y. Sugawara, H.M. Jacobson, A. Buyuktosunoglu, and P. Bose, Application-Level Power and Performance Characterization and Optimization on IBM Blue Gene/Q systems, *IBM J. Res. and Dev.*, 57(1/2), 2013.
- [3] CANDLE: Cancer Distributed Learning Environment, <http://candle.cels.anl.gov>.
- [4] CANDLE Benchmarks: <https://github.com/ECP-CANDLE/Benchmarks> (Accessed in Sep. 2019).
- [5] X. Chen, C. Xu, R. Dick, and Z. Mao, Performance and Power Modeling in a Multi-Programmed Multi-Core Env., *DAC2010*, 2010.
- [6] T. Chen, T. He, M. Benesty, et al., Extreme Gradient Boosting, Package “xgboost,” August 1, 2019.
- [7] M. Curtis-Maury, J. Dzierwa, C. Antonopoulos, and D. S. Nikolopoulos, Online Power-Performance Adaptation of Multithreaded Programs Using Hardware Event-Based Prediction, *ICS’06*, 2006.
- [8] M. Curtis-Maury, A. Shah, F. Blagojevic, D. S. Nikolopoulos, B. R. de Supinski, and M. Schulz, Prediction Models for Multidimensional Power-Performance Optimization on Many Cores, *PACT’08*.
- [9] G. Contreras and M. Martonosi, Power Prediction for Intel XScale Processors Using Performance Monitoring Unit Events, *ISLPED’05*.
- [10] P. Flach, *Machine Learning: The Art and Science of Algorithms that Make Sense of Data*, Cambridge University Press, 2012.
- [11] J. H. Friedman, Greedy function approximation: a gradient boosting machine, *Annals of Statistics*, 1189–1232, 2001.
- [12] J. Friedman, T. Hastie, and R. Tibshirani, Regularization Paths for Generalized Linear Models via Coordinate Descent, *Journal of Statistical Software* 33(1), 2010.
- [13] R. Ge, X. Feng, S. Song, et al., PowerPack: Energy Profiling and Analysis of High-Performance Systems and Applications, *IEEE Trans. on Para. and Dis. Sys.* 21(5), 2010.
- [14] glmnet: Lasso and Elastic-Net Regularized Generalized Linear Models, https://cran.r-project.org/web/packages/glmnet/vignettes/glmnet_beta.pdf.
- [15] J. L. Greathouse and G. H. Loh, Machine Learning for Performance and Power Modeling of Heterogeneous Systems (Invited Paper), in *IEEE/ACM International Conference on Computer-Aided Design (ICCAD’18)*, San Diego, CA, November 5–8, 2018.
- [16] T. Hothorn, K. Hornik, C. Strobl, and A. Zeileis, A Laboratory for Recursive Partytioning, Package “Party,” March 5, 2020.
- [17] C. Isci and M. Martonosi, Runtime Power Monitoring in High-End Processors: Methodology and Empirical Data, *36th IEEE/ACM Intern. Sym. on Microarchitecture*, 2003.
- [18] A. Karatzoglou, A. Smola, and K. Hornik, kernlab – An S4 Package for Kernel Methods in R, *Journal of Statistical Software*, 11(9), 1–20, 2004.
- [19] N. Kappiah, V. Freeh, and D. Lowenthal, Just In Time Dynamic Voltage Scaling: Exploiting Inter-Node Slack to Save Energy in MPI Programs, *SC05*, 2005.
- [20] M. Kuhn and K. Johnson, *Applied Predictive Modeling*, Springer, 2013.
- [21] M. Kuhn, The caret Package, <https://topepo.github.io/caret/index.html>, March 27, 2019 (<https://cran.r-project.org/web/packages/caret/>).
- [22] M. Kuhn, S. Weston, C. Keefer, N. Coulter, and R. Quinlan, Rule- and Instance-Based regression Modeling, Package “Cubist,” January 10, 2020. <https://topepo.github.io/Cubist>.
- [23] D. Li, B. de Supinski, M. Schulz, K. Cameron, and D. Nikolopoulos, Hybrid MPI/OpenMP Power-Aware Computing, *IEEE Intern. Conf. on Para. & Dist. Proc. Symp. (IPDPS2010)*, 2010.
- [24] A. Liaw, and M. Wiener, Breiman and Cutler’s random forests for classification and regression, Package “randomForest,” March 25, 2018.
- [25] Lim M, Porterfield A, Fowler R (2010) SoftPower: Fine-Grain Power Estimations Using Performance Counters, *19th Intl. High Perf. Dis. Comp. (HPDC ’10)*, 2010.
- [26] J. H. Laros III, P. Pokorny, and D. DeBonis, PowerInsight – A Commodity Power Measurement Capability, *International Green Computing Conference*, 2013.
- [27] C. Lively, X. Wu, V. Taylor, S. Moore, H. Chang, and K. Cameron, Energy and Performance Characteristics of Different Parallel Implementations of Scientific Applications on Multicore Systems, *Intern. J. High Perf. Comp. and App.*, 25(3), 2011.
- [28] C. Lively, V. Taylor, X. Wu, H. Chang, C. Su, K. Cameron, S. Moore, and D. Terpstra, E-AMOM: An Energy-Aware Modeling and Optimization Methodology for Scientific Applications on Multicore Systems, *Comp. Sci. – Res. and Dev.*, 29(3), 2014.
- [29] C. Lively, X. Wu, V. Taylor, S. Moore, H. Chang, C. Su, and K. Cameron, Power-Aware Predictive Models of Hybrid (MPI/OpenMP) Scientific Applications on Multicore System, *Comp. Sci. – Res. and Dev.*, 27(4), 2012.
- [30] P. Miller, mvtboost example, Dec. 5, 2016, https://cran.r-project.org/web/packages/mvtboost/vignettes/mvtboost_vignette.html
- [31] Z. Mayer, A Brief Introduction to caretEnsemble, Jan. 31, 2016.
- [32] A. Miyoshi, C. Lefurgy, E. V. Hensbergen, R. Rajamony, and R. Rajkumar, Critical Power Slope: Understanding the Runtime Effects of Frequency Scaling, *ICS’02*, June 2002.
- [33] S. Milborrow, Multivariate Adaptive Regression Splines, package “earth,” November 9, 2019.
- [34] MuMMI project, <http://www.mummi.org/info>.
- [35] H. Nagasaka, N. Maruyama, A. Nukada, T. Endo, and S. Matsuoka, Statistical Power Modeling of GPU Kernels Using Performance Counters, *Intern. Green Comp. Conf.*, 2010.
- [36] PAPI (Performance API), <http://icl.cs.utk.edu/papi/>
- [37] perf_event, https://perf.wiki.kernel.org/index.php/Main_Page.
- [38] PerfMon2, <http://perfmon2.sourceforge.net>.
- [39] R. Polikar, Ensemble Based Systems in Decision Making, *IEEE Circuits and Systems Magazine*, 2006. violin package: <https://cran.r-project.org/web/packages/violinplot/index.html>. <https://www.r-graph-gallery.com/violin-plot/>
- [40] C. Pospiech, Hardware Performance Monitor (HPM) Toolkit Users Guide, *Adv. Comp. Tech. Center, IBM Research*, June 2008.
- [41] PyPAPI, <https://flozz.github.io/ppypapi>
- [42] R. Rodrigues, A. Annamalai, I. Koren, and S. Kundu, A Study on the Use of Performance Counters to Estimate Power in Microprocessors, *IEEE Trans. on Cir. and Sys.*, 60(12), 2013.
- [43] E. Rotem, A. Naveh, D. Rajwan, A. Ananthkrishnan, and E. Weissmann, Power-management Architecture of the Intel Micro-architecture Code-named Sandy Bridge, *IEEE Micro*, 32(2), 2012.
- [44] scikit-learn, *Machine Learning in Python*, <https://scikit-learn.org/stable/index.html>
- [45] K. Singh, M. Bhadhauria, and S. A. McKee, Real Time Power Estimation and Thread Scheduling via Performance Counters, *Workshop on Design, Arch., and Sim. of Chip Multi-Processors*, 2008.

- [46] S. Song, C. Su, B. Rountree and K. Cameron, A Simplified and Accurate Model of Power-Performance Efficiency on Emergent GPU Architectures, IPDPS2013, 2013.
- [47] R. Springer, D. Lowenthal, B. Rountree, and V. Freech, Minimizing Execution Time in MPI Programs on an Energy-Constrained, Power-Scalable Cluster, PPOPP'06, 2006.
- [48] M. Suleman, M. Qureshi, and Y. Patt, Feedback-Driven Threading: Power-Efficient and High-Performance Execution of Multi-Threaded Workloads on CMPs, ASPLOS'08, 2008.
- [49] Summit, <https://www.olcf.ornl.gov/olcf-resources/compute-systems/summit/>
- [50] Theta, Cray XC40 system, <https://www.alcf.anl.gov/theta>.
- [51] G.L. Tsafack Chetsa, L. Lefevre, J.M. Pierson, P. Stolf, and G. Da Costa, Exploiting Performance Counters to Predict and Improve Energy Performance of HPC Systems, Future Generation Computer Systems, Vol. 36, July 2014.
- [52] X. Wu, V. Taylor, C. Lively, H. Chang, B. Li, K. Cameron, D. Terpstra and S. Moore, MuMMI: Multiple Metrics Modeling Infrastructure (Book Chapter), Tools for High Performance Computing 2013, Springer, 2014.
- [53] X. Wu and V. Taylor, Utilizing Hardware Performance Counters to Model and Optimize the Energy and Performance of Large Scaled Scientific Applications on Power-Aware Supercomputers, in IPDPS2016 Workshop on High Performance, Power-Aware Computing, Chicago, IL, May 23–27, 2016.
- [54] X. Wu, V. Taylor, J. Cook, and P. Mucci, Using Performance-Power Modeling to Improve Energy Efficiency of HPC Applications, IEEE Computer 49(10), 20–29, Oct. 2016.
- [55] X. Wu, V. Taylor, J. M. Wozniak, R. Stevens, T. Brettin, and F. Xia, Performance, Energy, and Scalability Analysis and Improvement of Parallel Cancer Deep Learning CANDLER Benchmarks, in 48th International Conference on Parallel Processing, Kyoto, Japan, August 5–8, 2019.
- [56] XLA architecture, <https://www.tensorflow.org/xla/architecture>
- [57] X. Zheng, L. K. John, and A. Gerstlauer, LACross: Learning-Based Analytical Cross-Platform Performance and Power Prediction, Int. J. Parallel Prog. 45:1488?–1514, 2017.
- [58] H. Zou and T. Hastie, Elastic-Net for Sparse Estimation and Sparse PCA, Package “elasticnet,” August 31, 2018.