Reducing numerical precision preserves classification accuracy in Mondrian Forests

Marc Vicuna[1], Martin Khannouz[1], Gregory Kiar[2], Yohan Chatelain[1], Tristan Glatard[1]

[1]Department of Computer Science and Software Engineering, Concordia University, Montreal, Canada
[2]Center for the Developing Brain, Child Mind Institute, New York, NY, USA

Abstract—**Mondrian Forests are a powerful data stream classification method, but their large memory footprint makes them ill-suited for low-resource platforms such as connected objects. We explored using reduced-precision floating-point representations to lower memory consumption and evaluated its effect on classification performance. We applied the Mondrian Forest implementation provided by OrpailleCC, a C++ collection of data stream algorithms, to two canonical datasets in human activity recognition: Recofit and Banos *et al*. Results show that the precision of floating-point values used by tree nodes can be reduced from 64 bits to 8 bits with no significant difference in F1 score. In some cases, reduced precision was shown to improve classification performance, presumably due to its regularization effect. We conclude that numerical precision is a relevant hyperparameter in the Mondrian Forest, and that commonly-used double precision values may not be necessary for optimal performance. Future work will evaluate the generalizability of these findings to other data stream classifiers.**

Index Terms—**numerical precision, memory footprint, Mondrian Forests, human activity recognition, data streams, supervised classification, floating-point representation**

I. INTRODUCTION

Mondrian Forests [1], an online variant of Random Forests, are powerful data stream classifiers that have been used in various applications. However, their considerable memory footprint limits their applicability to low-memory devices such as connected objects, which are common in online learning settings. We aim to decrease this footprint by reducing the numerical precision of floating-point data used throughout their training and application. Beginning from the OrpailleCC [2] implementation of Mondrian Forests, we applied this investigation to the popular use-case of human activity recognition using the Recofit [3] and Banos *et al* [4] datasets.

The floating-point data standard (IEEE-754) [5] is universally supported across numerical computing. The popularity of this standard is due to its flexibility to store both very large and very small numbers consistently. The size of these data structures, γ, is given by:

$$\gamma = p + e + 1, \qquad (1)$$

where p-bits are allocated to the mantissa (defining the value itself), e-bits are allocated to the exponent (its order of magnitude), and its sign fits in the remaining bit. Alongside the blueprint for constructing floating-point data, the IEEE-754 [5] norm specifies floating-point formats commonly used in CPUs, such as binary32 ($e = 8$, $p = 23$) and binary64 ($e = 11$, $p = 52$), also known as single and double precision

or "floats" and "doubles". Recently, deep learning has led to the emergence of reduced formats such as Microsoft's ms-fp8 Minifloats (8 bits), IEEE-754 binary16 (16 bits), bfloat16 (16 bits) [6], or posits (8, 16, 32, and 64 bits) [7]. However, reduced formats implicitly store less information compared to double precision and should therefore be used carefully.

Mondrian classifiers are built upon like-named Mondrian processes [8], which are multidimensional generalizations of Poisson processes. These can be interpreted as a recursive generative process that randomly makes univariate cuts partitioning the data space hierarchically into k-dimensional trees [9]. A key feature of these processes is that they are self consistent: a Mondrian process operating on the partition of some domain (e.g. a subdomain) is equivalent to a Mondrian process operating on the subdomain directly, allowing simultaneous subprocesses. While Mondrian processes can be infinite, Mondrian trees designate the classifiers formed by finite Mondrian processes, akin to decision trees. Each split is constructed by randomly selecting a feature and threshold, where the probability of selection is proportional to the normalized range of values for the feature. Finally, Mondrian Forests are constructed by ensembling a group of Mondrian trees, akin to Random Forest classifiers.

Data stream classifiers including Mondrian Forests, Micro-Cluster Nearest Neighbours [10], Hoeffding Trees [11], and simple Feedforward Neural Networks [12] are all commonly used in online contexts and were previously compared in [13] in the context of human activity recognition. Results showed superior performance of Mondrian Forests, and Hoeffding Trees. However, both methods remain too memory intensive for connected objects where memory is typically in the order of $\mathcal{O}(100\text{KB})$ [14]. Without compromising the breadth or depth of constructed forests, typically optimized hyperparameters which have a considerable impact on model performance, reducing the numerical precision used would potentially enable their application on memory-constrained platforms. However, given the inter-dependence of these parameters on model performance, numerical precision, memory consumption, and model parameterization must be studied together. This paper uniquely quantifies the impact of reducing numerical precision on the performance of Mondrian Forests classifiers for activity recognition and evaluates the impact these results have on their portability.

II. MATERIALS AND METHODS

Using Verificarlo [15], the numerical precision of Mondrian Forest classifiers as implemented in OrpailleCC was controlled and modified. The reduced-precision models were tested with multiple human activity recognition datasets, Recofit and Banos *et al.*, and performance was evaluated with F1 score.

A. Mondrian Forest Implementation

OrpailleCC's implementation of Mondrian Forests pre-allocates a fixed amount of memory shared by all the trees in the forest, leading to a constant memory footprint. When the amount of allocated memory is reached, tree growth is stopped in the forest.

Mondrian trees are tuned using three parameters: base count, discount factor, and budget. The base count is used to initialize the prediction in the root of the tree, the discount factor influences how reliant nodes are upon their parent prediction, and the budget controls the tree depth. In our experiments, we use the same sets of hyperparameters as in [13] (Table I).

TABLE I: Hyperparameters used in the Mondrian Forests

Number of trees	Base count	Discount	Budget
1	0.0	1.0	1.0
5	0.0	1.0	0.4
10	0.0	1.0	0.4
50	0.0	1.0	0.2

Under memory constraints, increasing the number of trees does not necessarily imply better performance as it reduces tree depth. We tested Mondrian Forests with 1, 5, 10 and 50 trees, and in configurations with 0.6 MB, 1.2 MB or 3 MB of memory allocated.

B. Simulating Reduced Precision

The experiments presented here rely on Verificarlo [15], a tool that allows for the perturbation and manipulation of floating-point data. Verificarlo is an LLVM-based compiler that instruments floating-point instructions through different backends. The VPREC [16] backend (for "Virtual PRE-Cision") was used here to simulate reduced floating-point formats. VPREC computes each instrumented floating-point operation using the original format (double precision in our case) and rounds the result to the desired precision, thereby ensuring correct rounding. VPREC also checks that the modified exponent is in the requested dynamic range $[e_{min}, e_{max}]$, where $e_{min} = 2 - 2^{e-1}$ and $e_{max} = 2^{e-1} - 1$.

We compared two instrumentation approaches to reduce the precision of Mondrian Forests in OrpailleCC:

a) Node instrumentation (NI): reduces floating-point precision only in tree nodes bounds, since most of the memory used by the forest is occupied by nodes. For the node instrumentation, new node bounds are rounded only when they are stored. This approach reduces memory consumption at most by half, as in OrpailleCC the memory occupied by nodes is equally divided between integers and floating-point numbers.

b) Whole instrumentation (WI): reduces precision in all floating-point values, i.e., node bounds, hyperparameters, and random splits. While whole instrumentation only marginally reduces memory consumption compared to node instrumentation, it implies that most floating-point calculations could be done in a globally reduced floating-point format, leading to faster computations and lower energy consumption.

For both instrumentations, we tested Mondrian Forests on 52 levels of precision, corresponding to each precision ranging from the maximum for 64-bit floating-point data down to a single bit (p in Equation 1). We also tested various exponent lengths (e in Equation 1) ranging from 11 to 2.

C. Datasets

We tested reduced precision on the two publicly-available datasets in human activity recognition. The Banos *et al.* dataset [4], [17] (henceforth referred to as "Banos") is a human activity dataset with 17 participants and 9 sensors per participant. Each sensor samples 3D acceleration, gyroscope, and magnetic field, as well as orientation in a quaternion format, producing a total of 13 values. Sensors were sampled at 50 Hz, and each sample was associated with one of 33 possible activities. In addition to the 33 activities, an extra activity labelled 0 indicated no specific activity. We preprocessed the Banos dataset using non-overlapping windows of one second (50 samples) and 6 axes (acceleration and gyroscope of the right forearm sensor). We computed the average and the standard deviation over the window as features for each axis and assigned the most frequent label to the window. To ensure that our results were independent of the order of the data points in the dataset, 7 random orderings were tested.

The second dataset, Recofit [3], is a human activity dataset containing 94 participants. Similar to the Banos dataset, the activity labeled 0 indicated no specific activity and similar activities were merged (see Table 2 in [18]). We preprocessed the dataset similarly to the Banos, using non-overlapping windows of one second and 6 axes of data (acceleration and gyroscope) from one sensor. From these 6 axes, the average and the standard deviation over the window were used as features. Previously reported F1 scores were globally much lower in this dataset, which is hypothesized to be due to the higher number of participants.

III. RESULTS

We compared the classification performance obtained with the node (NI) and whole (WI) instrumentations to the original double precision Mondrian Forests implementation in OrpailleCC. F1 scores were evaluated using the prequential error, meaning that for each new data point the model was first tested and then trained. We computed the resulting F1 score every 50 elements.

A. Reduced precision minimally affects performance

Figure 1 shows the difference in F1 scores ($\Delta F1 = F1_{p=i} - F1_{p=52}$) between reduced-precision and double precision classifiers across instrumentations and datasets. All
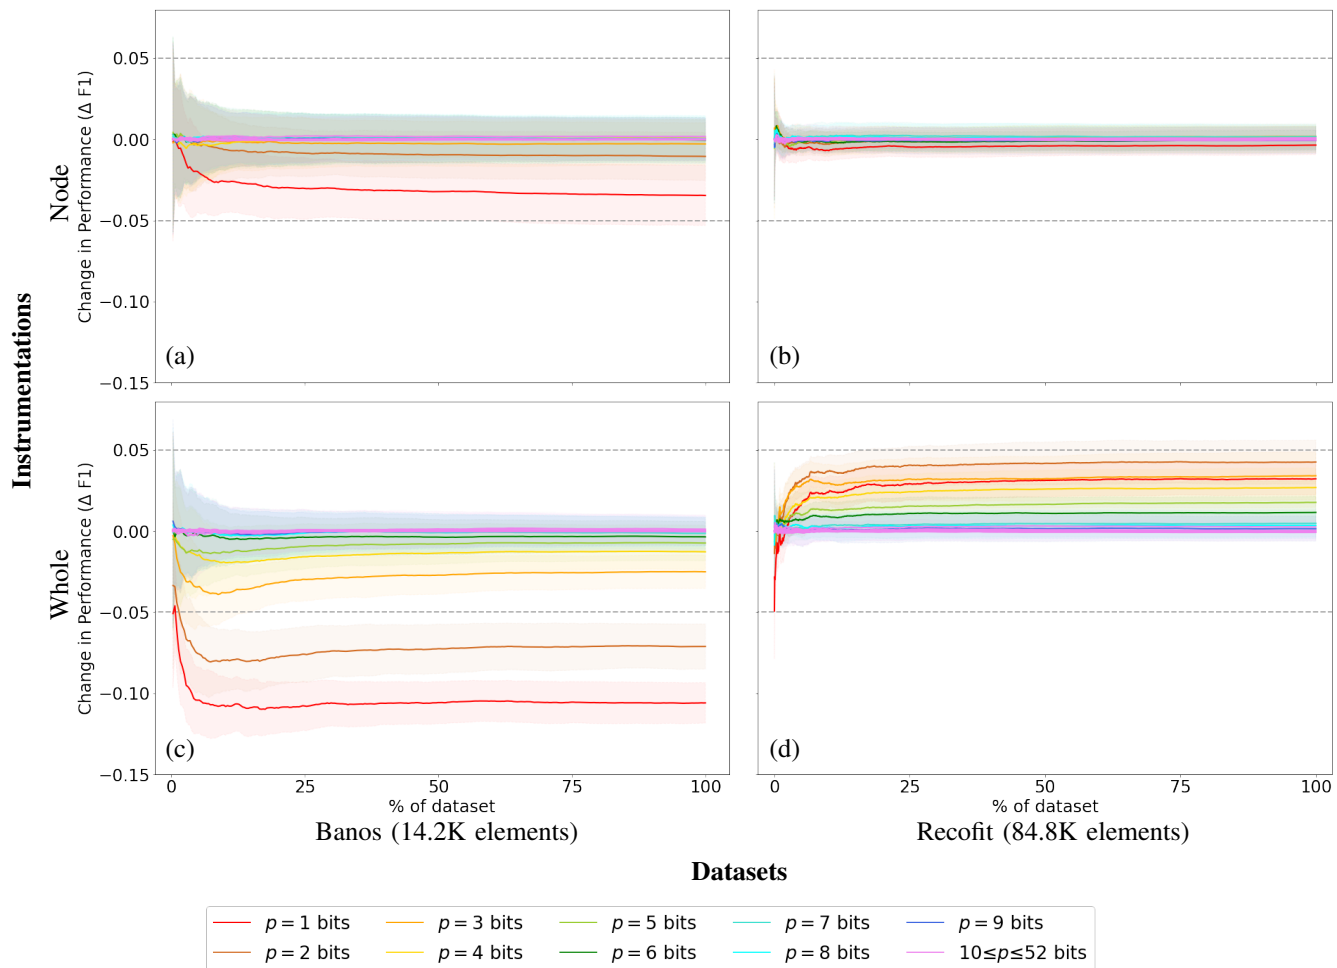
Fig. 1: Difference between the F1 scores obtained at reduced and double precision ($p = 52$ bits) with 3.0 MB of memory and 5 trees. Negative values indicate that the instrumented implementation performed worse than the original. Color shades represent the standard deviation in the corresponding Mondrian Forest. The grey lines at -0.05, 0.05 are for reference only. The 7 ordering of data points tested in Banos performed similarly.

datasets have been processed with an exponent length of 11 bits and with 3.0 MB of memory.

Across all configurations, the F1 score difference stabilized within approximately the first 25% of the dataset and remained stable beyond that point, indicating an overall good tolerance of Mondrian Forests to low precision.

With node instrumentation, reductions in F1 scores remained within the standard deviation of performance for configurations with $p \geq 2$ (Figures 1a and 1b), indicating that using a low precision for the node bounds has almost no influence on the F1 score. This can likely be attributed to the randomness of Mondrian Forests: given the randomness of the imposed splits, performance does not depend highly on their precision. Generally, the error due to reduced precision remains within the existing randomness of node bounds.

In contrast, whole instrumentation (WI) shows significant deviations in F1 scores for configurations with $p \leq 6$ (Figures 1c and 1d), indicating the high impact of precision loss in model hyperparameters. Distinctly, reducing the precision below $p = 6$ decreased the F1 score in the Banos dataset but increased the F1 score in the more complex Recofit dataset.

Observed differences between NI and WI stem from low-precision representations of the split thresholds and budget, since all other hyperparameters can be represented precisely with low precision (see Table I). Low-precision split bounds and budget both lead to increased regularization. Low-precision split bounds reduce minimal distances between random splits, while a low-precision budget reduces tree depth. In complex datasets such as Recofit, which contain numerous activity classes and participants, regularization reduces

overfitting. In contrast, simpler datasets such as Banos may not benefit from this form of regularization since the risk for overfitting is reduced.

B. Reduced exponent size unimpactful down to 4 bits

Table II shows F1 score differences ($\Delta F1 = F1_{e=i} - F1_{e=11}$) between reduced-exponent and full-exponent Mondrian Forests. All classifiers used 52 bits of precision. The reported F1 score differences were calculated after having processed all the data points in the stream.

TABLE II: F1 score differences with reduced exponent length

Instru	Dataset	Memory	Trees	Exponent length e			
				2 bits	3 bits	4 bits	5 bits
NI	Banos	0.6 MB	1	-0.4321	-0.0446	0.0000	0.0000
			5	-0.5374	-0.0157	0.0000	0.0000
			10	-0.5344	-0.0173	0.0000	0.0000
			50	-0.3971	-0.0047	0.0000	0.0000
		1.2 MB	1	-0.4696	-0.0904	0.0000	0.0000
			5	-0.5937	-0.0275	0.0000	0.0000
			10	-0.5805	-0.0217	0.0000	0.0000
			50	-0.4771	-0.0045	0.0000	0.0000
		3.0 MB	1	-0.4628	-0.0878	0.0000	0.0000
			5	-0.6334	-0.0404	0.0000	0.0000
			10	-0.6396	-0.0331	0.0000	0.0000
			50	-0.5482	-0.0155	0.0000	0.0000
	Recofit	0.6 MB	1	-0.1763	-0.1291	-0.0031	0.0000
			5	-0.1491	-0.1070	-0.0009	0.0000
			10	-0.0996	-0.0700	0.0000	0.0000
			50	-0.0482	-0.0225	0.0000	0.0000
		1.2 MB	1	-0.1948	-0.1470	-0.0038	0.0000
			5	-0.1802	-0.1388	-0.0025	0.0000
			10	-0.1501	-0.1090	-0.0009	0.0000
			50	-0.0683	-0.0408	0.0000	0.0000
		3.0 MB	1	-0.2127	-0.1601	-0.0042	0.0000
			5	-0.2214	-0.1726	-0.0018	0.0000
			10	-0.1966	-0.1521	-0.0008	0.0000
			50	-0.1003	-0.0717	-0.0001	0.0000
WI	Banos	0.6 MB	1	N/A	-0.4628	-0.0006	0.0000
			5	N/A	-0.5590	-0.0137	0.0000
			10	N/A	-0.5404	-0.0309	0.0000
			50	N/A	-0.3984	-0.1729	0.0000
		1.2 MB	1	N/A	-0.4840	-0.0002	0.0000
			5	N/A	-0.6130	-0.0009	0.0000
			10	N/A	-0.5937	-0.0078	0.0000
			50	N/A	-0.4798	-0.1351	0.0000
		3.0 MB	1	N/A	-0.4923	-0.0001	0.0000
			5	N/A	-0.6571	-0.0031	0.0000
			10	N/A	-0.6586	-0.0091	0.0000
			50	N/A	-0.5590	-0.1388	0.0000
	Recofit	0.6 MB	1	N/A	-0.1856	-0.0121	0.0000
			5	N/A	-0.1552	-0.0389	0.0000
			10	N/A	-0.1069	-0.0260	0.0000
			50	N/A	-0.0548	-0.0260	0.0000
		1.2 MB	1	N/A	-0.2066	0.0088	0.0000
			5	N/A	-0.1874	-0.0017	0.0000
			10	N/A	-0.1577	-0.0235	0.0000
			50	N/A	-0.0753	-0.0220	0.0000
		3.0 MB	1	N/A	-0.2237	0.0159	0.0000
			5	N/A	-0.2321	0.0360	0.0000
			10	N/A	-0.2058	0.0212	0.0000
			50	N/A	-0.1073	-0.0107	0.0000

For both datasets and instrumentations, all tree numbers, and all memory values, the F1 score did not significantly change until the exponent length was reduced to ≤ 3 bits, with some exceptions in WI. For lower exponents lengths, forests with larger numbers of trees tended to perform better. The whole instrumentation at $e = 2$ was not evaluated since it causes segmentation faults due to the creation of infinite values. For the tested datasets, an exponent size of 4 bits may be used without loss of classification performance. For other datasets, this limit could likely be estimated given awareness of the dynamic range of the acquired signals.

C. Larger forests are more robust to reductions in precision

Table III shows F1 score differences between reduced and double precision for each set of tested hyperparameters. All classifiers used an exponent length of 11 bits. As before, the reported F1 score differences were calculated after having processed all the complete stream of data.

TABLE III: F1 score differences with reduced precision

Instru	Dataset	Memory	Trees	Precision p					
				1 bit	2 bits	3 bits	4 bits	5 bits	6 bits
NI	Banos	0.6 MB	1	-0.0514	-0.0132	0.0119	0.0146	0.0135	0.0134
			5	-0.0032	0.0087	0.0098	-0.0065	0.0019	0.0007
			10	-0.0070	-0.0011	-0.0014	0.0012	0.0007	-0.0022
			50	-0.0077	-0.0026	-0.0040	-0.0029	-0.0006	-0.0015
		1.2 MB	1	-0.0561	-0.0274	-0.0082	-0.0114	-0.0022	-0.0019
			5	-0.0280	-0.0130	-0.0005	-0.0045	0.0013	-0.0025
			10	-0.0152	-0.0030	-0.0049	0.0021	0.0040	-0.0029
			50	-0.0036	0.0008	-0.0007	0.0008	0.0014	0.0027
		3.0 MB	1	-0.0692	-0.0323	-0.0058	-0.0070	0.0004	0.0006
			5	-0.0257	-0.0060	0.0037	-0.0022	0.0037	0.0017
			10	-0.0155	0.0004	0.0001	0.0042	0.0014	0.0029
			50	-0.0033	0.0022	-0.0006	-0.0010	0.0003	0.0012
	Recofit	0.6 MB	1	-0.0063	-0.0033	-0.0009	-0.0034	-0.0004	-0.0024
			5	-0.0024	0.0003	0.0000	-0.0013	-0.0013	-0.0019
			10	0.0006	0.0014	0.0005	0.0002	0.0012	0.0009
			50	0.0018	0.0006	0.0005	0.0003	0.0008	0.0001
		1.2 MB	1	-0.0154	-0.0051	-0.0012	-0.0014	-0.0025	-0.0012
			5	-0.0031	-0.0011	-0.0002	-0.0004	-0.0007	0.0001
			10	-0.0005	0.0025	0.0023	0.0010	0.0026	0.0010
			50	0.0018	-0.0006	0.0002	-0.0010	0.0001	-0.0005
		3.0 MB	1	-0.0140	-0.0072	-0.0007	-0.0005	-0.0005	-0.0003
			5	-0.0036	-0.0001	0.0018	-0.0001	0.0007	-0.0004
			10	-0.0026	0.0019	0.0013	-0.0007	0.0012	0.0000
			50	0.0007	0.0001	-0.0004	-0.0003	0.0004	-0.0011
WI	Banos	0.6 MB	1	-0.0755	-0.0570	-0.0330	-0.0239	-0.0072	-0.0064
			5	-0.0935	-0.0603	-0.0272	-0.0124	-0.0055	-0.0059
			10	-0.1150	-0.0510	-0.0172	-0.0090	-0.0038	-0.0041
			50	-0.2147	-0.1505	-0.0773	-0.0111	0.0020	0.0011
		1.2 MB	1	-0.0705	-0.0647	-0.0271	-0.0193	-0.0165	0.0016
			5	-0.1087	-0.0678	-0.0309	-0.0177	-0.0098	-0.0069
			10	-0.1203	-0.0624	-0.0217	-0.0107	-0.0011	-0.0030
			50	-0.2223	-0.1351	-0.0513	-0.0049	0.0025	0.0001
		3.0 MB	1	-0.0737	-0.0660	-0.0331	-0.0348	-0.0127	-0.0051
			5	-0.1038	-0.0729	-0.0294	-0.0149	-0.0106	-0.0068
			10	-0.1228	-0.0648	-0.0199	-0.0122	-0.0033	-0.0042
			50	-0.2035	-0.0994	-0.0405	-0.0079	0.0023	0.0002
	Recofit	0.6 MB	1	-0.0024	0.0121	0.0093	0.0058	0.0031	0.0014
			5	0.0159	0.0214	0.0088	0.0076	0.0066	0.0026
			10	0.0103	0.0206	0.0127	0.0123	0.0038	0.0024
			50	0.0112	0.0077	0.0018	0.0019	0.0018	0.0007
		1.2 MB	1	0.0084	0.0203	0.0099	0.0132	0.0046	0.0052
			5	0.0197	0.0295	0.0176	0.0131	0.0043	0.0057
			10	0.0091	0.0227	0.0135	0.0132	0.0052	0.0053
			50	0.0071	0.0058	0.0010	0.0020	0.0021	0.0011
		3.0 MB	1	0.0124	0.0282	0.0177	0.0148	0.0071	0.0105
			5	0.0322	0.0426	0.0341	0.0268	0.0177	0.0115
			10	0.0181	0.0356	0.0300	0.0264	0.0142	0.0060
			50	0.0097	0.0147	0.0036	0.0038	0.0038	0.0004

For NI and WI, classifiers with a higher number of trees were consistently less affected by the precision reduction. This was likely because ensembling lowers the standard deviation in final predictions. The exception to the above observation was when using the WI for prediction of the Banos dataset with $p < 4$, where the loss of useful complexity in the node splits appeared amplified by the resulting lower number of nodes per tree.

We also observed the F1 score differences at $p > 2$ were non-significantly related to memory. The fact that the reduction in precision did not affect memory allocations differently suggests that these benefits could be translated to higher or lower memory allocations than the ones tested.

As mentioned in Figure 1 and observed here consistently, the NI performance was stable on every precision tested, while WI had different tendencies depending on the dataset. Computations based on Table III revealed that $p = 3$ was the lowest precision where all results are within 2 standard deviations of results at full precision.
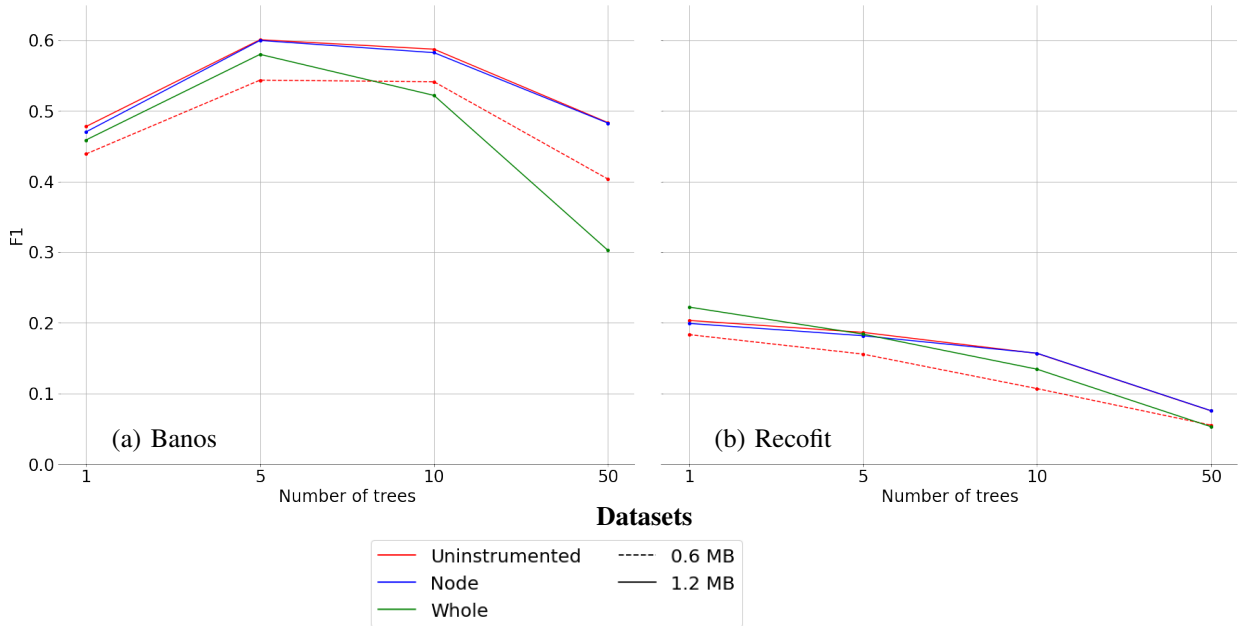
Fig. 2: F1 Score of the uninstrumented classifiers with 0.6 MB and 1.2 MB, NI and WI classifiers at 4 bit exponent and 3 bit precision (8 bits) with 1.2 MB.

D. Expected memory footprint reduction

Under constant memory constraints, reducing memory consumption could allow trees to grow deeper, thereby potentially improving classification performance. To quantify this performance improvement, Figure 2 shows performance using double precision with 0.6 MB of memory and reduced precision (8 bits) with 1.2 MB of memory. As the execution platform does not natively support 8-bit minifloats, we achieve similar memory scaling by doubling available memory. In reality, taking into account integer values in the implementation, the amount of memory consumed by the uninstrumented implementation is only $1.8\times$ larger than for NI or WI. Indeed, the memory consumed by the uninstrumented implementation is $I + F$, where I is the memory consumed by integer values and F is the memory consumed by floating-point values ($I = F$ in our implementation), while the memory consumed by NI or WI with 8-bit minifloats is $I + F/8$. The limit of 0.6 MB was selected as a realistic memory allocation bound for connected objects [14].

The F1 score obtained with NI at 1.2 MB showed a 16.62% average increase (10.79% on Banos, 22.45% on Recofit) compared to the uninstrumented classifier at 0.6 MB, which quantifies the classification performance improvement expected from reduced precision under memory constraints on platform-supporting minifloats. Besides, at 1.2 MB, NI has a very similar behavior to the uninstrumented 1.2 MB classifier (superposed blue and red curves), confirming that 8-bit bounds for nodes are practically possible and do not affect performance.

As observed previously, WI was more variable, leading to inconsistent F1 score improvements or decreases depending on the number of trees.

IV. RELATED WORK

Mixed and reduced precision have been studied in various contexts with popular classifiers in recent years. Rojek [19] dynamically reduced the energy consumption of Random Forest training up to 36% by using mixed precision data. The objective of reducing energy consumption and memory on tree-based classifiers is common to our goal, however, Mondrian forests were not previously included in such a test. It is also not known how tight the relationship is between energy savings obtained in the context of supercomputers versus those in connected devices. Since Random Forests and Mondrian Forests are very similar, we hypothesize efficient implementations of our improvements on Mondrian Forests would lead to lower energy consumption.

Yuval *et al.* [20] demonstrated the use of online and offline Random Forests to learn a parameterization from the coarse-grained output of a three-dimensional high-resolution idealized atmospheric model. In their experiment, numerical precision is low (16 bits) and the classifiers are trained across resolutions. We note from their work that online random forests have more stable performance across resolutions than offline ones. This could indicate our observed performance stability across numerical precisions is correlated characteristics of online algorithms.

Wang *et al.* [21] successfully reduced a wide spectrum of DNN data and computations for training to 8-bit floating-point

numbers with a potential of 2-4x improved throughput. Recent research [22] in Deep Neural Networks (DNNs) demonstrates that DNNs can have more than 99% of their training reduced to half precision with similar accuracy to corresponding unreduced training. The comparison of unreduced, dynamically reduced, and mostly reduced trained performances is similar to our approach, but deals with a very different machine learning model, and very low-precision formats are not explored.

V. CONCLUSION

Node bounds in Mondrian Forests can be implemented using an 8-bit floating-point format with negligible impact on the F1 score, reducing the entire memory footprint by approximately $1.8\times$ compared to double precision implementations on 0.6 MB. For a given amount of memory, decreasing node size allows Mondrian Forests to grow deeper trees or have more trees, potentially resulting in higher performance. In our configurations, performance improvements resulting from memory reduction at 0.6 MB improved the F1 score by 16.62% on average. Improvements are hypothesized to be proportionally greater on lower memory at an optimal tree number.

Reducing the precision of random cuts and the budget hyperparameter in Mondrian trees may improve or worsen classification performance depending on the dataset. While the observed performance improvements are encouraging, future research is required to fully understand when such improvements can be expected. In the meantime, hyperparameterizing numerical precision using variable-length floating-points formats is a practical way to test this strategy.

Overall, reduced precision creates important opportunities for data stream classification in memory-constrained platforms such as connected objects. While we focused on reducing memory consumption, reduced precision would also decrease computation time substantially. Fully leveraging this result requires reduced-precision floating-point formats to be implemented in the target systems. Future research is required to test this technique with other data stream classifiers.

VI. CODE & DATA AVAILABILITY

The data, experimental results, plots, and code involved in our experiments are publicly available at https://github.com/big-data-lab-team/benchmark-har-data-stream/tree/veripaille.

REFERENCES

[1] B. Lakshminarayanan, D. M. Roy, and Y. W. Teh, "Mondrian forests: Efficient online random forests," *arXiv preprint arXiv:1406.2673*, 2014.

[2] M. Khannouz, B. Li, and T. Glatard, "OrpailleCC: a library for data stream analysis on embedded systems," in *Journal of Open Source Software*, 2019.

[3] D. Morris, T. S. Saponas, A. Guillory, and I. Kelner, "Recofit: using a wearable sensor to find, recognize, and count repetitive exercises," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 3225–3234, 2014.

[4] O. Baños, M. Damas, H. Pomares, I. Rojas, M. A. Tóth, and O. Amft, "A benchmark dataset to evaluate sensor displacement in activity recognition," in *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, pp. 1026–1035, 2012.

[5] M. S. Committee *et al.*, "754-2019-ieee standard for floating-point arithmetic," 2019.

[6] D. Kalamkar, D. Mudigere, N. Mellempudi, D. Das, K. Banerjee, S. Avancha, D. T. Vooturi, N. Jammalamadaka, J. Huang, H. Yuen, *et al.*, "A study of bfloat16 for deep learning training," *arXiv preprint arXiv:1905.12322*, 2019.

[7] J. L. Gustafson and I. T. Yonemoto, "Beating floating point at its own game: Posit arithmetic," *Supercomputing Frontiers and Innovations*, vol. 4, no. 2, pp. 71–86, 2017.

[8] D. M. Roy and Y. Teh, "The Mondrian process," in *Advances in Neural Information Processing Systems* (D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, eds.), vol. 21, Curran Associates, Inc., 2009.

[9] J. L. Bentley, "Multidimensional binary search trees used for associative searching," *Commun. ACM*, vol. 18, p. 509–517, Sept. 1975.

[10] M. Tennant, F. Stahl, O. Rana, and J. B. Gomes, "Scalable real-time classification of data streams with concept drift," *Future Generation Computer Systems*, vol. 75, pp. 187–199, 2017.

[11] P. Domingos and G. Hulten, "Mining high-speed data streams," in *Proceedings of the Sixth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '00, (New York, NY, USA), p. 71–80, Association for Computing Machinery, 2000.

[12] T. L. Fine, *Feedforward neural network methodology*. Statistics for engineering and information science., Springer, 2021.

[13] M. Khannouz and T. Glatard, "A benchmark of data stream classification for human activity recognition on connected objects," *Sensors*, vol. 20, no. 22, p. 6486, 2020.

[14] M. Research, "Neblina datasheet: Neblina v2 module bluetooth® smart 9 axis motion tracking," 2017.

[15] C. Denis, P. D. O. Castro, and E. Petit, "Verificarlo: Checking floating point accuracy through monte carlo arithmetic," in *2016 IEEE 23nd Symposium on Computer Arithmetic (ARITH)*, pp. 55–62, IEEE, 2016.

[16] Y. Chatelain, E. Petit, P. de Oliveira Castro, G. Lartigue, and D. Defour, "Automatic exploration of reduced floating-point representations in iterative methods," in *European Conference on Parallel Processing*, pp. 481–494, Springer, 2019.

[17] O. Banos, M. A. Toth, M. Damas, H. Pomares, and I. Rojas, "Dealing with the effects of sensor displacement in wearable activity recognition," *Sensors*, vol. 14, no. 6, pp. 9995–10023, 2014.

[18] A. Dehghani, O. Sarbishei, T. Glatard, and E. Shihab, "A Quantitative Comparison of Overlapping and Non-Overlapping Sliding Windows for Human Activity Recognition Using Inertial Sensors," *Sensors*, vol. 19, no. 22, p. 5026, 2019.

[19] K. Rojek, "Machine learning method for energy reduction by utilizing dynamic mixed precision on gpu-based supercomputers," *Concurrency and Computation: Practice and Experience*, vol. 31, no. 6, p. e4644, 2019. e4644 cpe.4644.

[20] J. Yuval and P. A. O'Gorman, "Stable machine-learning parameterization of subgrid processes for climate modeling at a range of resolutions," *Nature Communications*, vol. 11, no. 1, p. 3295, 2020.

[21] N. Wang, J. Choi, D. Brand, C. Chen, and K. Gopalakrishnan, "Training deep neural networks with 8-bit floating point numbers," *CoRR*, vol. abs/1812.08011, 2018.

[22] J. Osorio, A. Armejach, E. Petit, and M. Casas, "A dynamic approach to accelerate deep learning training," *ICLR 2020 Conference Blind Submission*, 2020.