

Memory Efficient Weightless Neural Network using Bloom Filter

Leandro Santiago¹, Leticia Verona¹, Fabio Rangel¹, Fabrício Firmino¹,
Daniel S. Menasche¹, Wouter Caarls², Mauricio Breternitz Jr³,
Sandip Kundu⁴, Priscila M.V. Lima¹ and Felipe M. G. França¹ *

1- Federal University of Rio de Janeiro (UFRJ), Brazil, 2- PUC Rio, RJ, Brazil,
3- University of Lisbon, Lisbon, Portugal, 4- UMass, Amherst, USA

Abstract. Weightless Neural Networks (WNNs) are a kind of Artificial Neural Networks based on RAM memory broadly explored as solution for pattern recognition applications. Memory-oriented solutions for pattern recognition are typically very simple, and can be easily implemented in hardware and software. Nonetheless, the straightforward implementation of a WNN requires a large amount of memory resources making its adoption impracticable on memory constrained systems. In this paper, we propose a new model of WNN model which utilizes Bloom filters to implement RAM nodes. Bloom filters reduce memory requirements, and allow false positives when determining if a given pattern was already seen in data. We experimentally found that for pattern recognition purposes such false positives can build robustness into the system. The experimental results show that our model using Bloom filters achieves competitive accuracy, training time and testing time, consuming up to 6 orders of magnitude less memory resources in comparison with the standard Weightless Neural Network model.

1 Introduction

Weightless Neural Networks (WNNs) [1] are neuron models based on Random Access Memory (RAM) where each neuron is defined as RAM node. These models have been shown as attractive solutions to solve pattern recognition and artificial consciousness applications achieving competitive performance. WiSARD (Wilkie, Stoneham and Aleksander's Recognition Device) is the pioneering WNN distributed commercially [2] which provides simple and efficient implementation enabling to deploy learning capabilities into real-time and embedded systems.

The straightforward WiSARD implementation needs a considerable amount of memory resources to obtain good learning features. To address this problem, we propose a new WiSARD model that replaces RAM nodes with Bloom filters. Bloom filters [3, 4] are probabilistic data structures which represent a set as small bit array allowing the occurrences of false positives, that is, an element can be considered a member of set even it is not. Although false positives detract certain applications, we experimentally discovered that for pattern recognition purposes they can build robustness into the system (as dropout does to deep neural networks) – Bloom WiSARD presents similar accuracy when contrasted against WiSARD, but uses significantly less resources and, in this sense,

*The authors thank CAPES, CNPq and FAPERJ for the financial support for this work.

Table 1: Comparison of classifiers: simpler models such as Bloom WiSARD typically favor generalization. WiSARD and Dict WiSARD are logically equivalent whereas Bloom WiSARD has more degrees of freedom.

	Space per discriminator	Use of hashes for	Accuracy
WiSARD	$N2^M$ bits	no hashes	reference
Dict WiSARD	significantly less than WiSARD (worst case equal)	exact set membership (with collision checking)	equal to WiSARD (given mapping from input to RAM)
Bloom WiSARD	typically similar to Dict WiSARD (tunable by design)	approximate set membership (no collision checking)	potentially greater than WiSARD (hashes are tunable)

is more robust. Our experiments analyze accuracy, training time, testing time and memory consumption of our model compared against standard WiSARD and WiSARD implemented with hash tables (see Table 1).

2 Background: WiSARD and Bloom filters

WiSARD (Wilkie, Stoneham and Aleksander’s Recognition Device) is a multi-discriminator WNN model proposed in the early 80’s [2] that recognizes patterns from binary data. Each class is represented by a discriminator which contains a set of RAMs. A binary input with $N \times M$ bits is split into N tuples of M bits. Each tuple n , $n = 1, \dots, N$, is a memory address to an entry of the n -th RAM. Each RAM contains 2^M locations.

A *pseudo-random mapping* is a deterministic function that maps each binary input matrix to a set of N tuples of M bits each. The function is typically a pseudo-random shuffling of the binary input matrix, hence the name pseudo-random mapping. Each discriminator may be associated to a different pseudo-random mapping, that must remain the same across training and classification phases.

At the training phase, initially all RAMs have their locations set to zero (0). Each training sample is treated by the corresponding discriminator which sets to one (1) all accessed RAM positions. At the classification phase, the input is sent to all discriminators generating responses per discriminators by summing all accessed RAM values. The discriminator with the highest response is chosen as representative class of the input.

Bloom filters [3] are space-efficient data structures for Approximate Membership Query (AMQ) which test whether an element belongs to a given set or not with a certain false positive probability. In other words, sometimes the membership query will respond that an element was stored while actually it was not inserted. A Bloom filter is composed of an m -bit array and k independent hash functions that map an element into k bit array positions.

The standard Bloom filter supports insertion and query operations. Initially, all bit array positions are zeroed. In the insertion operation, an element is

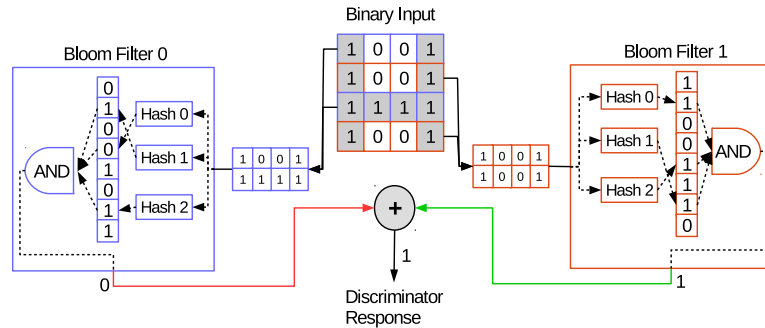


Fig. 1: Bloom WiSARD classifier: 16-bit input, 8-bit tuples and 2 Bloom filters.

mapped into k positions of the bit array using the k hash functions and the corresponding k bits are set to 1. The query operation looks up the k positions mapped from the input element, indicating it as either a member of the set, considering a false positive rate if all values are 1's, or a non-member when any value is 0. Note that a Bloom filter always reports a true negative whenever an element is not a member.

The probability of false positive p is affected by the parameters m , n and k , corresponding to bit array size, number of elements to store and number of hash functions, respectively [5]. Given the probability p and capacity n , the ideal parameters m and k are calculated through the following formulas: $m = -n \ln(p) / \ln(2)^2$ [6] and $k = m \ln(2) / n$ [5].

3 WiSARD based on Bloom Filters

The memory structures subsumed by a WiSARD WNN are typically sparse. We extend WiSARD by replacing RAMs with Bloom filters to reduce its memory resources by avoiding storage of irrelevant zero positions. The new model is termed Bloom WiSARD. The key idea is to store a set of tuples mapped to each Bloom filter and test if a given tuple belongs to its corresponding set.

Figure 1 presents the Bloom WiSARD design. On the training phase, the tuples are inserted into Bloom filters by updating the k bit array positions. On the classification phase, the tuples are queried into their associated Bloom filters returning whether each tuple is a member or not by ANDing all k bit values. Similar to WiSARD, the discriminator responses are calculated by summing the N Bloom filter membership results so that the highest response selects the appropriate discriminator to represent the input.

Our Bloom WiSARD implementation utilizes a double hashing technique [7] to generate k hash functions in the form: $h(i, k) = (h_1(k) + i \times h_2(k)) \pmod{n}$, where h_1 and h_2 are universal hash functions. We adopt MurmurHash as seed for h_1 and h_2 [8].

Table 2: Performance of classifiers in binary classification problems.

Dataset	WNN	Acc	Training (s)	Testing (s)	Memory (KB)
Adult	WiSARD	0.722	4.414	1.05	8978432
	Dict WiSARD	0.721	1.947	1.188	383.535
	Bloom WiSARD	0.718	1.932	1.166	80.173
Australian	WiSARD	0.843	0.002	0.001	4096
	Dict WiSARD	0.841	0.002	0.001	11.299
	Bloom WiSARD	0.834	0.002	0.001	8.613
Banana	WiSARD	0.87	0.052	0.028	13312
	Dict WiSARD	0.871	0.054	0.033	23.428
	Bloom WiSARD	0.864	0.058	0.036	3.047
Diabetes	WiSARD	0.698	0.001	0.0007	2048
	Dict WiSARD	0.689	0.001	0.0008	6.553
	Bloom WiSARD	0.69	0.001	0.0008	4.793
Liver	WiSARD	0.593	0.001	0.0007	5120
	Dict WiSARD	0.587	0.001	0.0008	6.387
	Bloom WiSARD	0.591	0.001	0.0009	2.344

4 Experiments and Results

Dataset and experimental setup To evaluate our proposed model, we compare Bloom WiSARD to two different WiSARD versions: standard WiSARD and dictionary WiSARD (see Table 1). The latter version is implemented with Hash Tables instead of RAMs to store the tuples values as key-value pair at each position, with the key representing the memory address and, the value, the tuple [9]. We select the MNIST database [10] and a subset of binary classification and multiclass classification datasets used in [11]. Most of the problems were taken from UCI public repository [12] and they have different characteristics in terms of number of samples, number of classes and number of features. Some datasets do not provide the training set and testing set in separated files. For these datasets, we adopt the same methodology applied in [11]: we randomly shuffle the data and partition it in 3 parts, such that 2/3 and 1/3 are used for training and testing sets, respectively.

The experiments were performed on an Intel Core i7-6700(3.40GHz) processor with 32GB of RAM running Ubuntu Linux 16.04. The core of all WiSARD experiments was implemented in a single-thread C++11 library accessed through a Python interface. To convert the input attributes to binary format, we concatenate all binary attributes using thermometer (resp., hot encoding) to transform the continuous (resp., categorical) attributes. The input size, number of RAMs and tuple size varied according to the dataset, but were kept across all considered WiSARD architectures. Bloom filters are setup with 10% of false positive probability. The capacities were empirically selected from each dataset and m and k were obtained through the formulas presented in Section 2.

Table 3: Performance of classifiers in multiclass classification problems.

Dataset	WNN	Acc	Training (s)	Testing (s)	Memory (KB)
Ecoli	WiSARD	0.793	0.0005	0.0005	7168
	Dict WiSARD	0.799	0.0005	0.0005	5.664
	Bloom WiSARD	0.799	0.0005	0.0007	3.281
Iris	WiSARD	0.985	0.0001	0.000009	1536
	Dict WiSARD	0.977	0.0001	0.000008	0.747
	Bloom WiSARD	0.976	0.0001	0.0001	0.703
Letter	WiSARD	0.845	1.483	0.16	10223616
	Dict WiSARD	0.846	0.0717	0.22	121.748
	Bloom WiSARD	0.848	0.07	0.208	91.292
MNIST	WiSARD	0.917	4.317	0.33	9175040
	Dict WiSARD	0.916	0.811	0.475	1368.457
	Bloom WiSARD	0.915	0.77	0.369	819.049
Satimage	WiSARD	0.851	0.048	0.034	27648
	Dict WiSARD	0.853	0.05	0.049	69.141
	Bloom WiSARD	0.851	0.053	0.05	12.656
Shuttle	WiSARD	0.87	0.119	0.064	8064
	Dict WiSARD	0.869	0.12	0.078	4.956
	Bloom WiSARD	0.868	0.132	0.103	3.691
Vehicle	WiSARD	0.67	0.003	0.0021	9216
	Dict WiSARD	0.672	0.003	0.0026	17.617
	Bloom WiSARD	0.662	0.003	0.0028	4.219
Vowel	WiSARD	0.876	0.0023	0.0025	14080
	Dict WiSARD	0.876	0.0023	0.0032	16.221
	Bloom WiSARD	0.876	0.0022	0.0036	6.445
Wine	WiSARD	0.932	0.0006	0.0003	4992
	Dict WiSARD	0.924	0.0005	0.0003	4.248
	Bloom WiSARD	0.926	0.0005	0.0004	2.285

Results and discussion All results are obtained through the mean of 20 runs with negligible standard deviation. Tables 2 and 3 show the results for binary classification and multiclass classification datasets, respectively. Note that the accuracy of Dict WiSARD and WiSARD slightly differ as we used different pseudo-random mappings at each training epoch (see Table 1). Overall, Bloom WiSARD achieved comparable accuracy, training time and testing time when compared against WiSARD and Dict WiSARD, while consuming a smaller amount of memory. Bloom WiSARD’s memory consumption is reduced up to 6 order of magnitude (Adult and Letter) compared against standard WiSARD and approximately 7.7 times (Banana) when compared against dictionary WiSARD. The memory resources can be further reduced by increasing the false positive rate and the accuracy can be increased by tuning the hash functions to capture essential aspects of the data, which we leave as subject for future work.

5 Conclusion

In this work we propose the Bloom WiSARD model which extends WiSARD by implementing RAM nodes as Bloom filters. By using Bloom filters, memory resources are significantly reduced and for pattern recognition purposes we experimentally found that Bloom filters can build robustness into the system. Our experiments show that the model provides good accuracy and requires low training and testing times. In addition, it consumes up to 6 orders of magnitude less resources than standard WiSARD and about 7.7 times less resources than WiSARD implemented with dictionaries. Future work will focus on extending Bloom filter operations such as frequency counts of elements stored, in order to enable Bloom WiSARD to use improved techniques such as DRASiW [13] or the Bloom filter false free zone [14]. More broadly, we envision that this work is one step further towards the use of Bloom filters for machine learning [4, 15].

References

- [1] I. Aleksander, M. De Gregorio, F. Maia Galvão França, P. Machado Vieira Lima, and H. Morton. A brief introduction to weightless neural systems. In *ESANN 2009, 17th European Symposium on Artificial Neural Networks*, 2009.
- [2] I. Aleksander, W.V. Thomas, and P.A. Bowden. Wisard-a radical step forward in image recognition. *Sensor Review*, 4(3):120–124, 1984.
- [3] Burton H. Bloom. Space/time trade-offs in hash coding with allowable errors. *Commun. ACM*, 13(7):422–426, July 1970.
- [4] Lailong Luo, Deke Guo, Richard T. B. Ma, Ori Rottenstreich, and Xueshan Luo. Optimizing Bloom filter: Challenges, solutions, and comparisons. *IEEE Communications Surveys and Tutorials*, 2019.
- [5] Peter C. Dillinger and Panagiotis Manolios. Bloom filters in probabilistic verification. In Alan J. Hu and Andrew K. Martin, editors, *Formal Methods in Computer-Aided Design*, pages 367–381, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg.
- [6] P. Sterne. Efficient and robust associative memory from a generalized bloom filter. *Biological Cybernetics*, 106(4):271–281, Jul 2012.
- [7] A. Kirsch and M. Mitzenmacher. Less hashing, same performance: Building a better Bloom filter. In Yossi Azar and Thomas Erlebach, editors, *Algorithms – ESA 2006*, 2006.
- [8] Wikipedia. Murmurhash fuction, 2019. <https://en.wikipedia.org/wiki/MurmurHash>.
- [9] R.J. Mitchell, J.M. Bishop, and P.R. Minchinton. Optimising memory usage in n-tuple neural networks. *Mathematics and Computers in Simulation*, 40(5):549 – 563, 1996.
- [10] Yann LeCun. The mnist database of handwritten digits, 1998. <http://yann.lecun.com/exdb/mnist/>.
- [11] G. Huang, H. Zhou, X. Ding, and R. Zhang. Extreme learning machine for regression and multiclass classification. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 42(2):513–529, April 2012.
- [12] Dua Dheeru and Efi Karra Taniskidou. UCI machine learning repository, 2017.
- [13] Massimo De Gregorio and Maurizio Giordano. Cloning DRASiW systems via memory transfer. *Neurocomputing*, 192:115–127, 2016.
- [14] Sándor Z. Kiss, Éva Hosszu, János Tapolcai, Lajos Rónyai, and Ori Rottenstreich. Bloom filter with a false positive free zone. In *IEEE INFOCOM*, 2018.
- [15] Moustapha M Cisse, Nicolas Usunier, Thierry Artieres, and Patrick Gallinari. Robust Bloom filters for large multilabel classification tasks. In *Advances in Neural Information Processing Systems*, pages 1851–1859, 2013.