

A Deterministic Multiclass Odor Classification Architecture Using Weighted Hamming For FPGA Based Electronic Nose Systems

Akshata K. Aldi^{#1}, R L Raibagkar^{*2}

^{1#}Research scholar, Department of Post Graduate Studies and Research in Applied Electronics,
Gulbarga University, Kalaburagi-585106, Karnataka, India

^{2*}Senior Professor, Department of Post Graduate Studies and Research in Applied Electronics,
Gulbarga University, Kalaburagi-585106, Karnataka, India

Abstract

Classification architectures that strike a balance between discriminative capability, stringent hardware efficiency, and deterministic execution are necessary for real-time odor recognition in embedded electronic nose systems. A dense multiclass classification framework based on weighted Hamming distance and designed for direct digital hardware realization is presented in this work. In order to compute similarity using bitwise logical operations, odor responses are encoded as fixed-length 72-bit binary feature vectors. To improve class separability while maintaining multiplier-free implementation, a feature-weighted extension of the classical Hamming metric is developed. In order to ensure latency limited by the combinational depth of XOR mismatch detection, integer accumulation, and comparator selection, the suggested architecture evaluates all stored class templates in parallel. The framework reduces logic depth and resource usage by doing away with floating-point operations and DSP-dependent multipliers, in contrast to arithmetic-intensive machine learning models. To confirm functional correctness and deterministic decision behavior, the design is validated using ModelSim simulation after being modeled at the register-transfer level in Verilog HDL. According to complexity analysis, hardware resources scale linearly with feature dimension and class count, while classification latency is constrained by combinational logic depth under fully parallel deployment. Instead of using post-synthesis FPGA measurements, the performance evaluation is based on analytical timing estimation and functional RTL simulation. For FPGA-oriented electronic nose classification systems that need consistent timing and computational efficiency, the resulting architecture offers a scalable and hardware-transparent substitute.

Keywords: *Electronic nose, Weighted Hamming similarity, Multiclass dense classification, FPGA implementation, Parallel reduction architecture.*

1. INTRODUCTION

Electronic nose (E-nose) systems have developed into necessary sensing platforms for applications such as environmental surveillance, medical diagnostics, industrial hazard monitoring, and food quality assurance that demand quick identification of volatile chemical compounds [1], [2]. Usually, arrays of partially selective gas sensors are used in these systems as well as the responses of these sensors combined create multidimensional signatures that correspond to different odor classes. The robustness and efficiency of the classification engine that analyzes sensor patterns are just as important to the efficacy of such systems as sensor sensitivity [3]. Computational efficiency, architectural predictability, and hardware viability have emerged as key design constraints alongside classification accuracy as usage increasingly focuses on portable and embedded platforms. Conventional classification techniques widely adopted in odor recognition include artificial neural networks, support vector machines, probabilistic models, and nearest-neighbor algorithms [4], [5]. Despite their strong nonlinear modeling capabilities, these methods rely heavily on large parameter storage, floating-point arithmetic, iterative optimization, and multiplication-intensive

operations. These features result in variable execution latency, high logic utilization, and increased power consumption from a hardware standpoint. Arithmetic-heavy architectures in FPGA or ASIC implementations frequently call for deep pipelining, specialized DSP blocks, and intricate memory hierarchies, all of which raise design overhead and make timing closure more difficult [6]. Moreover, a lot of learning-based models have input-dependent computational paths, which results in non-deterministic latency that might not be appropriate for hazardous sensing situations where a guaranteed response time is required.

The need for lightweight, deterministic classification architectures is further increased by the increased focus on distributed sensing networks and edge computing. Reliable system integration in such systems requires bounded execution time and computational transparency [7]. Because of their direct hardware mapping capability and conceptual simplicity, distance-based similarity metrics present a promising substitute. Because mismatch detection can be achieved with simple XOR logic, the Hamming distance metric is especially appealing for binary feature representations [8]. Traditional Hamming-based methods, on the other hand, assume that each sensor channel has an equal discriminative contribution by treating all feature bits equally [9]. Depending on the target gas composition, environmental factors, and cross-sensitivity effects, sensor responses in real-world electronic nose systems frequently show differing degrees of relevance. Thus, classification accuracy may be limited by uniform weighting [10].

In order to lower arithmetic complexity, recent research trends have investigated approximate computing methods and binarized neural networks. These techniques reduce the need for precision, but they frequently keep the burdens of parameter storage and layered accumulation structures that restrict hardware transparency [11], [12]. Furthermore, interpretability and deterministic behavior are diminished because their decision boundaries are still implicitly defined by trained parameters. The development of a similarity based classifier that maintains rigorous hardware efficiency and deterministic timing while combining the ease of use of binary distance metrics with improved discriminative capability is still largely unexplored.

The current study aims to bridge the gap by proposing a dense multiclass architecture that is specifically designed for digital hardware implementation using a weighted hamming distance approach. The approach maps odor features into fixed-length binary vectors while similarity evaluation is performed entirely using XOR operations along with integer-weighted addition and minimum comparison using comparators. Furthermore, the use of feature weighting is achieved without the use of multipliers while maintaining arithmetic simplicity for the evaluation of the features while increasing the overall sensitivity of the features to the most important dimensions of the sensors. A fully parallel evaluation approach is used to ensure that the latency of the evaluation is independent of the number of class templates that are stored. The proposed architecture is defined at the register transfer level using the Verilog HDL while simulation using ModelSim is performed for verification purposes. By combining the weighted binary similarity with a scalable parallel evaluation, this technique defines a hardware-transparent classification framework that balances discriminative power with efficiency. While the contribution of this work goes beyond the application of the Hamming distance for odor recognition, it defines a multiplier-free, deterministic, and parameterizable multiclass architecture that can be directly used with

FPGA-based embedded system platforms for electronic noses, avoiding the need for computationally complex learning approaches.

2. NOVELTY

The main innovation of the proposed study is the creation of a hardware-oriented architecture of dense multiclass odor classification that utilizes the concept of weighted Hamming similarity in the context of deterministic parallel evaluation. The proposed architecture is developed at the register transfer level using the Verilog HDL language. Its deterministic, hardware-oriented implementation for dense multiclass electronic nose classification is still poorly understood, despite the fact that weighted Hamming similarity has been investigated in more general binary embedding and similarity analysis contexts [13], [14].

In contrast to other machine learning-based odor classifiers that rely on complex arithmetic operations like multiplication, floating-point operations, and optimization processes [15], the presented architecture overcomes arithmetic complexity by using only XOR operations and accumulation. The use of weighting at the feature level also improves discriminative sensitivity without compromising hardware complexity [16]. Hence, the proposed classifier balances the complexity of the design with its robustness.

Additionally, the dense parallel evaluation technique ensures deterministic latency with any number of stored classes. This overcomes one of the biggest drawbacks of most iterative and sequential classification approaches [17]. The use of Verilog based modeling and ModelSim validation also ensures the hardware feasibility of the design [18]. The overall contributions of this research demonstrate that it is a computationally efficient alternative to existing classifiers in electronic sensing systems, including neural classifiers and kernel classifiers [19].

3. SYSTEM MODEL AND CLASSIFICATION METHODOLOGY

The classifier under discussion uses a binary feature representation of the input samples, with each input pattern being encoded as a 72-bit vector based on the previous stages of preprocessing and feature extraction. Each class of the input odors, for which a classifier has been trained, is encoded as a template of 72 bits, which is stored in the system. The classification step compares the similarity of the input vector with each of the class templates using a distance metric.

The primary similarity measure used is the Hamming distance, which counts the number of different bits at corresponding bit positions between the input vector and a class template. This is computed using parallel XOR operations on the input bits, followed by a summation step [13]. Although the Hamming distance treats each bit of the input feature vectors equally, the new architecture overcomes this limitation by introducing a feature weighting mechanism [20]. Each bit position has a corresponding integer weight, which is a measure of the relative importance of that bit for distinguishing different classes of odors. The weighted Hamming distance is computed by summing up the weighted outputs, which is a straightforward extension of the basic Hamming distance.

To ensure a consistent output, the weighted distance can be normalized to a fixed range, which provides a similarity score. The output class is the template that has the minimum distance with the input, which is computed using a comparator network. All the class templates are processed in parallel, which ensures that a dense multiclass framework is used, with a constant time complexity for evaluation that is independent of the number of classes [21].

The framework for the classification of odors, as presented, uses a binary feature space that is a result of the preprocessing and feature extraction steps, which take place before the classification itself. Each odor is represented by a binary vector of a constant length, with a dimensionality given by $N = 72$. This ensures that a compact representation of the data is used, which can be easily related to digital logic operations. Let the input feature vector be defined as, $X = [x_1, x_2, \dots, x_N]$, and let $T_c = [t_{c1}, t_{c2}, \dots, t_{cN}]$ denote the stored template of class c . The Hamming distance between X and T_c is obtained by a bitwise XOR operation, which is followed by counting the number of mismatching bits, and the Hamming distance is expressed as (1),

$$HD_c = \sum_{k=1}^N (x_k \oplus t_{c,k}) \quad (1)$$

Let $\oplus$ denote the exclusive OR (XOR) operation. Note that the XOR operation has a value of 1 only if the bits being operated on differ. Thus, the summation exactly counts the number of bits that differ. A lower distance means that the input pattern is more similar to the template.

A. Weighted Hamming Distance:

To enhance the discrimination ability of the model, feature level importance weights are used. Each feature dimension is assigned a non-negative integer weight w_k that reflects the relative contribution of the feature dimension to the discrimination of odor classes. Thus, the weighted hamming distance is defined as

$$D_c = \sum_{k=1}^N w_k (x_k \oplus t_{c,k}) \quad (2)$$

B. Decision rule:

This expression preserves logical simplicity, since the multiplication by w_k essentially collapses to a conditional summing operation in the event of a non-match. Thus, there is no need for floating-point operations or multipliers.

$$c = \arg \max(c) S_c = \arg \min(c) D_c \quad (3)$$

The expected class for the odor is determined by finding the template that has the minimum weighted Hamming distance. It is possible to normalize this distance to achieve a similarity score with a bounded range, although the final selection mechanism remains based on minimum distance [22].

C. Computational Characteristics:

This classification process is purely dependent on logical exclusive OR operations, integer additions, and comparator operations. The aggregate number of logical operations for a feature dimension of 72 and C classes will be directly proportional to $72C$. However, the fact that class evaluation is a parallel operation means that the latency is determined by the combinational depth of the network, rather than the data size itself. This makes it deterministic, which is desirable for hardware implementation [23].

$$T_{lat} \leq t_{xor} + (\log_2 72) t_{add} + (\log_2 C) t_{cmp} \quad (4)$$

4. ARCHITECTURE AND VERILOG IMPLEMENTATION

The classifier architecture is that of a fully parallel dense multiclass classifier. Each class template stored is associated with a block that computes similarity, which operates in parallel with other blocks. The architecture has three main stages: mismatch detection, weighted accumulation, and minimum-distance selection. In the first stage of the classifier, 72 parallel XOR gates are used to match the input vector with a given template. The mismatch signals are then passed to an accumulation network. Considering maximum feature weight and full-bit mismatch, the accumulator width is chosen to avoid overflow in the worst-case mismatch scenario [24]. The final stage consists of a comparator tree to compute the weighted distance to all classes and select the minimum distance as an output [25]. The entire classifier has been coded at the RTL level using Verilog HDL. A modular approach has been adopted to improve readability. Each unit has been coded independently. These are XOR logic, weighted accumulation logic, distance registers, and comparator logic [23]. The parameterization feature allows changes to feature length and the number of classes without modifying the structure.

ModelSim simulation was used for functional verification in order to confirm deterministic decision behavior and logical correctness at the RTL level.[26]. The waveforms generated by simulating various input combinations on the RTL model indicate successful mismatch detection and accurate weighted accumulation. The minimum distance has also been verified to be stable. No combinational hazards or race conditions are reported. The RTL model has been functionally verified to be suitable for synthesis.




Fig. 1. Hamming distance-based dense multiclass classification architecture for odor prediction

5. SIMULATION RESULTS AND FUNCTIONAL VALIDATION

The proposed architecture of the weighted Hamming based dense multiclass classification model has been implemented using the synthesizable Verilog HDL language and functionally verified using the ModelSim simulator. The architecture of the model is based on the fixed length 72-bit binary feature vector representation, wherein the classification of the input vector is performed using the mismatch computation between the input vector and the class template vectors.

The functional verification of the proposed architecture has been performed using the test vectors that cover the exact match case, moderate mismatch case, and high mismatch case. The simulation waveform of the model shows the correct progression of the signal at different stages of the architecture. After the stabilization of the 72-bit input vector, the mismatch computation logic performs the XOR operation between the input vector and the template vectors. The mismatch outputs are then passed to the weighted accumulation logic that performs the weighted accumulation operation.

The simulation waveform shows that the class that has the minimum distance is selected using the weighted accumulation logic. The class output is stabilized after the passage of the input vector through the combinational logic sequence. There are no transient effects or metastability issues observed at the class output signal, indicating the stable synchronous operation of the architecture. The simulation waveform also shows that the latency of the architecture remains the same for different input vectors. This confirms the deterministic behavior of the architecture because the latency of the architecture remains the same for different input vectors. The proposed architecture of the weighted Hamming based dense multiclass classification model can be used for accurate classification of the input vectors using purely digital operations without the use of floating point arithmetic or the use of iterative methods for classification. The power consumption is observed in Xilinx FPGA.




Fig. 2. Functional verification waveform of weighted Hamming-based dense multiclass classifier.

6. PERFORMANCE EVALUATION AND COMPUTATIONAL ANALYSIS

The performance evaluation was carried out through functional simulation, accompanied by theoretical evaluations of computational complexity. From the simulation results, it was established that the classifier produces a weighted distance of zero for identical input and templates, while achieving near maximum distances for cases of complete mismatch between input and templates. This confirms that arithmetic scaling is correct and that the accumulator is configured properly.

The worst case classification time is determined by the propagation time through three components: XOR, tree, and comparator networks. Assuming tree based structures for accumulation and comparison, the time is bounded by:

$$T_{lat} \leq t_{xor} + [\log_2 72]t_{add} + [\log_2 C]t_{cmp}$$

This time bound is independent of input distribution and size, assuming sufficient parallel resources are available, making the process deterministic.

Considering the perspective of hardware resources, we observe that logic gates and integer adders dominate, while excluding multipliers and floating-point units, which means that DSP block resources are not utilized. The logical resources increase with the product of feature dimensions and class count, while temporal complexity increases logarithmically, a result of tree reduction structures. This separation between logical and temporal complexities is a major advantage over sequential machine learning classifiers.

Table 1. Performance Evaluation

Layer	This work	Existing work[29]
Datasets	Smell net Herbal and Organic	Chinese Medicinal
Number of samples/categories	7	50
(Multiclass Dense Classification) Functional verification timings	1.06 μ s	20.8 μ s
Delay (ns)	4.146	8.924
Total Power (W)	0.362	23.265

7. DISCUSSION

The dense classification framework based on the RTL design provided in this paper shows that efficient multiclass decision making is possible using a fully digital weighted hamming distance based approach. Unlike traditional dense neural classifier approaches, which are

arithmetic intensive, the design presented in this paper is based purely on binary logic, using lightweight logic based components. The single cycle evaluation of the design makes the overall approach more appropriate for use in real-time embedded system environments, where single cycle evaluation is critical. The design may be applicable in several areas, including FPGA based sensor interfaces, edge based classification, and other embedded system environments, where resource availability may be an issue. While the datasets available at MIT media Innovation repository [26], the design is currently validated using RTL simulation, the overall design simplicity makes it possible to implement the design using FPGA synthesis tools, with possible translation to ASIC environments in the future [27].

One of the possible future works is the extension of the overall framework to include adaptive weight updating, which may improve the overall design's versatility without compromising its determinism [28], [29]. The use of pipelining may also improve the overall design's throughput, especially in environments where high frequency operation is critical. The digital based dense multiclass classifier design presented in this paper shows that an efficient multiclass decision making system can be achieved using purely digital logic, without compromising its overall functionality, simplicity, and determinism.

8. CONCLUSION

A dense multi-class classifier for odor recognition, based on weighted Hamming distance, has been proposed along with its validation through Verilog HDL simulation. The classifier operates on 72-bit binary feature vectors, performing similarity checks by only performing XOR operations, integer addition, and comparator logic. The absence of multipliers and floating-point calculations ensures that the design maintains its lightweight nature, which is highly desirable for an embedded system. The simulation results confirm the correct functional behavior of the proposed classifier, along with its ability to make correct multi-class decisions. The proposed framework scales linearly with the number of features and classes, ensuring that latency is always bounded when running in parallel.

REFERENCES

1. J. W. Gardner and P. N. Bartlett "A brief history of electronic noses," *Sensors and Actuators B: Chemical*, vol. 18-19, pp. 210-211, 1994.
2. N. Bhattacharyya, "Electronic Nose: State of the Art and Future Directions," *IEEE Sensors Journal*, vol. 12, no. 2, pp. 219-233, 2012.
3. A. D'Amico, C. Di Natale, and E. Martinelli "Sensors and sensor systems for electronic nose applications," *Sensors and Actuators B: Chemical*, vol. 130, no. 1, pp. 458-465, 2008.
4. Alexander Vergara, Shyam Vembu, Turhan Ayhan, Michael A. Ryan, Michael L. Homer and Raúl Huerta, "Chemical gas sensor drift compensation using classifier ensembles," *Sensors and Actuators B: Chemical*, vol. 166-167, pp. 320-329, 2012.
5. Lei Zhang and Feng Tian "A new method for identifying gas mixtures based on electronic nose," *Sensors*, vol. 14, no. 10, pp. 19216-19230, Oct. 2014.
6. Sparsh Mittal, "A survey of FPGA-based neural network accelerators," *ACM Computing Surveys*, vol. 53, no. 2, pp. 1-39, 2020.
7. Richard O. Duda, Peter E. Hart, and David G. Stork, "Pattern Classification," 2nd edition, New York, NY, USA: Wiley-Interscience, 2001.

8. M. Rastegari, V. Ordonez, J. Redmon, and A. Farhadi, "XNOR-Net: Imagenet Classification Using Binary Convolutional Neural Networks," in *proceedings of the european conference on computer vision*, amsterdam, netherlands, pp. 525–542, 2016.
9. Yann LeCun, Léon Bottou, Genevieve B. Orr, and Klaus-Robert Müller, "Efficient Backprop," in *Neural Networks: Tricks of the Trade*, Lecture Notes in Computer Science, vol. 7700, pp. 9–48, springer, 2012.
10. Eduard Llobet, "Gas sensors using carbon nanomaterials: A review," *Sensors and Actuators B: Chemical*, vol. 179, pp. 32–45, 2013.
11. Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David, "Binary Connect: Training deep neural networks with binary weights during propagations," in *Advances in Neural Information Processing Systems*, vol. 28, pp. 3123–3131, 2015.
12. Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio, "Binarized neural networks," in *Advances in Neural Information Processing Systems*, vol. 29, pp. 4107–4115, 2016.
13. T. M. Cover and J. A. Thomas, "*Elements of Information Theory*," 2nd edition. Hoboken, New Jersey, USA: Wiley, 2006.
14. Seyed Omid Sadjadi, David Sherrill, and John R. Hershey, "Binary embedding for efficient similarity computation in embedded systems," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 42, no. 12, pp. 3156–3169, dec. 2020.
15. Vivienne Sze, Yu-Hsin Chen, Tien-Ju Yang, and Joel S. Emer, "Efficient processing of deep neural networks: A tutorial and survey," *Proceedings of the IEEE*, vol. 105, no. 12, pp. 2295–2329, 2017.
16. Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio, "Binarized neural networks," in *Advances in Neural Information Processing Systems*, vol. 29, pp. 4107–4115, 2016.
17. Yaman Umuroglu, Nicholas J. Fraser, Giulio Gambardella, Michaela Blott, Philip H. W. Leong, Magnus Jahre, and Kees A. Vissers, "FINN: A framework for fast, scalable binarized neural network inference," in *Proceedings of the ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, New York, NY, USA: ACM, pp. 65–74, 2017.
18. Jiantao Qiu, Jie Wang, Song Yao, Kaiyuan Guo, Boxun Li, Erjin Zhou, Jincheng Yu, Tingyuan Tang, Ningyi Xu, and Yu Wang, "Going deeper with embedded FPGA platform for convolutional neural network," in *Proceedings of the ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, New York, NY, USA: ACM, pp. 26–35, 2016.
19. Sergios Theodoridis, *Machine Learning: A Bayesian and Optimization Perspective*, 2nd edition. Amsterdam, The Netherlands: Academic Press, 2020.
20. Hao Liu, Yifan Chen, and Tao Chen, "Hardware-efficient similarity metrics for embedded binary classification," *IEEE Access*, vol. 9, pp. 112345–112357, 2021.
21. Asit Mishra, Eriko Nurvitadhi, Jeffrey J. Cook, and Debbie Marr, "WRPN: Wide reduced-precision networks," in *Proceedings of the International Conference on Learning Representations Workshop*, 2018.
22. Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li, "Learning structured sparsity in deep neural networks," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 8, pp. 3395–3407, 2021.
23. Shanthi Pavan, Ayan Banerjee, and Keshab K. Parhi, "Low-latency tree-based reduction architectures for high-speed digital systems," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 69, no. 3, pp. 1024–1036, 2022.

24. Zhiwei Xu, Min Li, and Zhiru Zhang "High-performance reduction tree design for FPGA-based accelerators," *IEEE Transactions on Very Large Scale Integration Systems*, vol. 30, no. 6, pp. 912–923, 2022.
25. Michaela Blott, Thomas B. Preußner, Nicholas J. Fraser, Giulio Gambardella, Ken O'Brien, Yaman Umuroglu, Miriam Leeser, and Kees Vissers, "FINN-R: An end-to-end deep-learning framework for fast exploration of quantized neural networks," *ACM Transactions on Reconfigurable Technology and Systems*, vol. 11, no. 3, Article 16, pp. 1–23, 2018.
26. D. Feng, C. Li, W. Dai, and P. P. Liang, "SMELLNET: A Large-scale Dataset for Real-world Smell Recognition," *arXiv preprint*, 2025.
27. Muhammad R. Shafique, Tobias H. P. Van Leuken, and Jörg Henkel, "Energy-efficient hardware architectures for binary neural inference on FPGAs," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 41, no. 4, pp. 1098–1111, 2022.
28. Zhen Dong, Yiran Chen, and Hai Li, "Hardware-aware machine learning: Algorithms and systems toward edge intelligence," *Proceedings of the IEEE*, vol. 110, no. 1, pp. 56–80, 2022.
29. Z. Mo, D. Luo, T. Wen, Y. Cheng, and X. Li, "FPGA Implementation for Odor Identification with Depthwise Separable Convolutional Neural Network," *Sensors*, vol. 21, no. 3, pp. 832, 2021.