A Fuzzy-Based Approach for Interpretable Spike Detection in Living Neural Biocomputers

Adham Aboulkheir, Hani Hagras, Michael Barros
The Artificial Intelligence Group
Unconventional Communications and Computing Lab (UC2)
School of Computer Science and Electronic Engineering
University of Essex
Colchester, United Kingdom
E-mail: aa23575@essex.ac.uk

***Abstract*—The convergence of neuroscience and artificial intelligence has led to the development of "living biocomputers," where biological neural networks are cultured on Micro-Electrode Array (MEA) chips. This creates hybrid biocomputing platforms that blur the boundaries between biology and technology. A critical challenge in this domain is the accurate and interpretable detection of neural spikes from MEA recordings. While black-box models like deep neural networks have achieved good accuracy, their lack of transparency hinders scientific inquiry. Revealing spatial-temporal patterns across electrode arrays can inform electrode placement strategies, stimulation protocols, and our understanding of how biological neural networks process information. This approach could bridge the gap between understanding the organization of neurons, their activity and the computing task in hand and ultimately allow living neural biocomputers to have a solution to the current "black-box" biocomputing approaches. We adopt a fuzzy logic framework because it can handle the biological uncertainty and signal variability inherent in MEA recordings, enabling more faithful spike discrimination than rigid threshold-based or crisp machine learning approaches. This paper presents a fuzzy rule-based classifier that combines high performance with human-readable interpretability for spike detection in neural computing. Our approach integrates ANNIGMA (Artificial Neural Network for Input Gain and Measurement Analysis)-based feature selection to identify the most salient spike characteristics and a genetic algorithm to optimize the fuzzy rule base. We validated our model using multi-chip pooled validation protocol, achieving a peak F1 score on testing data of 97.74% on a two-chip experiment and 96.12% on a six-chip experiment. The resulting fuzzy rules are linguistically interpretable, allowing researchers to understand the underlying logic of spike detection. This work can be a step forward towards building trustworthy neural biocomputers with direct applications to drug discovery, disease modeling, and understanding the fundamentals of neural computation.**

Keywords: Biocomputing, synthetic biological intelligence, fuzzy classifiers.

I. INTRODUCTION

The emergence of living neural biocomputers (one example depicted in Fig.1) represents one of the most transformative developments at the intersection of synthetic biology, neuroscience, and computing. Unlike traditional silicon-based processors, these neurons-on-a-chip systems harness the remarkable computational properties of biological neural networks self-organization, adaptive learning, massive parallelism, and extraordinary energy efficiency [1], [2], [3]. Implemented through Micro-Electrode Array (MEA) technology, living neural computers integrate cultured neurons with electronic interfaces, creating hybrid biocomputing platforms that blur the boundaries between biology and technology [4], [5]. These systems exhibit spontaneous network formation, synaptic plasticity, and emergent computational capabilities that arise from the collective dynamics of thousands of interconnected neurons [1], [2], [3]. The potential applications are profound: accelerating drug discovery for neurological diseases by testing compounds on functional human neural tissue, developing next generation neuroprosthetics that seamlessly interface with biological nervous systems, and unlocking fundamental insights into neural computation that could revolutionize artificial intelligence [1], [6]. As synthetic biology advances, living neural biocomputers are poised to complement traditional computing architectures, offering biological solutions to problems where conventional approaches struggle such as pattern recognition, adaptive control, and real-time learning in dynamic environments [1], [3], [6].

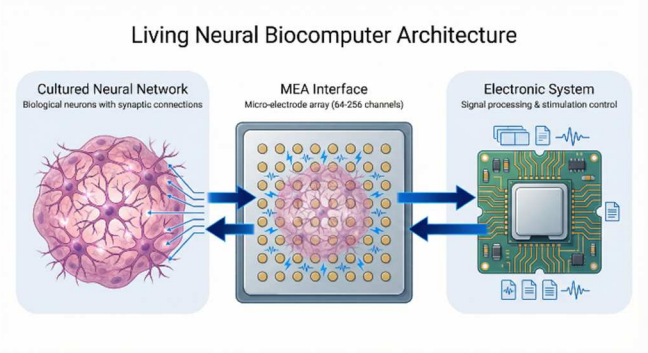

Fig. 1. Conceptual architecture of a living neural biocomputer showing cultured neurons interfaced with a micro-electrode array (MEA) and electronic signal processing system.

However, realizing this vision requires robust analytical tools to decode neural activity in these biological computing platforms. These systems encode information through neuron membrane potential, which is described as an action potential and has its peak called a spike. Spike detection within neuronal populations is challenging, since the MEAs have

one-many mapping with neurons, thus becoming the foundational step for understanding how living neural networks process and transmit information [7], [8]. Yet biological systems present unique challenges including signal variability across neurons, fluctuating noise levels and inter-chip differences that reflect the inherent stochasticity of living tissue — uncertainties that require adaptive and robust handling methods [9]. However, current methods face a critical trade-off between accuracy and interpretability, with high-performing machine learning approaches operating as black boxes that obscure the underlying neural features driving classification decisions. This paper addresses this challenge by presenting an interpretable fuzzy rule-based classifier for spike detection that maintains high accuracy while generating human-readable rules. Fuzzy logic systems have demonstrated strong capabilities in building transparent, adaptive decision-making systems across complex real-world domains [10], [11], making them a natural fit for the interpretability demands of living neural biocomputer analysis and advancing the development of transparent analysis tools essential for the continued progress of living neural computer technology.

Section II presents a high-level overview on biological living computers. Section III presents an overview of the proposed fuzzy based approach for interpretable spike detection in living neural computers. Section IV presents the experiments and results while Section V presents the conclusions and future work.

II. A High-level Overview on Biological Living Computers

Micro-Electrode Array (MEA) technology represents a revolutionary interface between biological neural networks and electronic systems, consisting of silicon chips embedded with microelectrodes capable of recording extracellular potentials and delivering electrical stimulation to guide network development [4], [5]. When neurons are cultured on MEA substrates, they spontaneously form functional networks exhibiting synaptic connections, synchronized bursting activity, and emergent computational properties, including activity-dependent plasticity and homeostatic regulation [1], [2], [3]. These platforms support transformative applications, from screening drug candidates on human neural tissue to developing hybrid systems where biological networks perform pattern recognition tasks [1], [4], [6]. Extracting meaningful information from these systems requires accurate spike detection — identifying short-duration voltage deflections associated with neuronal firing amidst background noise [7], [8]. Spike detection methods have evolved from threshold-based approaches and template matching to machine learning models such as convolutional neural networks, which achieve high accuracy but operate as black boxes, preventing the scientific validation that neuroscience applications demand. Fuzzy rule-based systems address this interpretability gap by generating transparent IF-THEN linguistic rules with measurable quality metrics such as support and confidence, revealing the spatial-temporal electrode patterns that characterize spike events. A further challenge is achieving robust generalization across different MEA chips, where variability in neuron density, synapse formation, and electrode-neuron coupling reflects the stochastic nature of biological development [1], [2]; multi-chip pooled validation strategies address this by exposing classifiers to diverse biological preparations during both training and testing [1].

III. A Proposed Fuzzy-Based Approach for Interpretable Spike Detection in Living Neural Biocomputers

We are proposing a Fuzzy-based approach [12] for interpretable spike detection in living neural biocomputers (illustrated in Fig. 2). Our approach is divided into two main stages: a Learning Stage, where the fuzzy rule-based classifier is trained, and a Recognition Stage, where the trained classifier is used to detect/predict spikes in new MEA recordings. This two-stage architecture, inspired by the work of Sanz et al. [13], [14], allows for a clear separation between model training and deployment. Our methodology prioritizes interpretability, ensuring that the generated fuzzy rules can be understood and validated by the related experts.

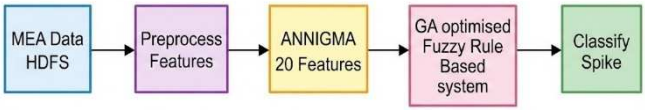

Fig. 2. The overall architecture of the proposed fuzzy rule-based spike detection system, showing the complete pipeline from MEA data input to final classification output.

A. Learning Stage

The Learning Stage involves several steps to construct the fuzzy rule-based classifier from a labeled training dataset.

1) Data Preprocessing and Feature Extraction

The MEA recordings are preprocessed offline to extract spike rate features from the raw electrode signals. Each sample in our dataset represents the neural network's response to a single stimulation event, captured across 142 electrodes in the MEA chip. For each electrode, the firing rate during the response window (30-80ms post-stimulation) is computed, which captures the primary spike activity following stimulation. This rate encoding represents the average firing rate of each electrode during this critical period, resulting in a 142-dimensional feature vector per sample (one feature per electrode). To reduce dimensionality and improve interpretability, we applied ANNIGMA feature selection to identify the 20 most informative electrodes [15]. This reduces the feature space from 142 to 20 dimensions while preserving discriminative information for spike detection.

2) ANNIGMA Feature Selection

To reduce the dimensionality of the feature space and improve the interpretability of the fuzzy rules, we employ the ANNIGMA (Artificial Neural Network for Input Gain and Measurement Analysis) feature selection algorithm. Our implementation of ANNIGMA leverages the internal weights

of a trained neural network to assess feature importance. The algorithm first trains a simple multi-layer perceptron (MLP) with one hidden layer of 50 neurons (using ReLU activation) on the classification task. The MLP is trained for a maximum of 500 iterations using the Adam optimizer with early stopping (patience=10 on a 20% validation split). Training typically converges within 100-200 iterations when early stopping is triggered. The importance of each input feature is then calculated as the sum of the absolute values of the weights connecting it to all neurons in the hidden layer. The importance score I for a feature f_i is calculated as the sum of the absolute values of the product of weights connecting it to all neurons in the hidden layer and their connections to the output layer [15]:

$$I(f_i) = \sum_{j=1,k=1}^{n} \left| w_{ij} \times w_{jk} \right| \quad (1)$$

where w_{ij} is the weight from the i_{th} input neuron (representing feature f_i) to the j_{th} *neuron* in the hidden layer. The scores are normalized, and the top 20 features with the highest importance scores are selected for the subsequent stages.

3) Fuzzy Rule Generation

Once the most informative features are selected, we generate a set of fuzzy rules. Each feature's domain is partitioned into three fuzzy sets: Low, Medium, and High, using uniformly distributed triangular membership functions. The triangular membership function parameters are uniformly distributed across each feature's range. A fuzzy rule is an IF-THEN statement, for example: IF Peak_Amplitude is HIGH AND Spike_Width is LOW THEN Class is Spike. Each sample in the training set generates a candidate rule. The matching degree of a sample to a rule's antecedent is calculated using the product T-norm [13], [14]:

$$match(x, R_j) = mA_j(x_p) = \prod_{i=1}^{n} \mu A_{i,j}(x_i) \quad (2)$$

where $\mu A_{i,j}(x_i)$ is the membership degree of the sample's i_{th} feature in the fuzzy set A_i of the rule R_j, *n is the number of rule antecedents*. To manage the rule base size of the rule, we filter rules based on their support and confidence metrics [13]. The support of a fuzzy rule $A_j \rightarrow C_j$ measures the coverage of the rule over the training examples belonging to class C_j [13]:

$$Support(A_j \rightarrow C_j) = \frac{\left(\sum_{p \in Class_{C_j}} mA_j(x_p)\right)}{|N|} \quad (3)$$

where N is the total number of training samples. The confidence is defined as follows [13], [14]:

$$Confidence(A_j \rightarrow C_j) = \frac{\left(\sum_{p \in Class_{C_j}} mA_j(x_p)\right)}{\left(\sum_{p=1}^{N} \mu m(x_p)\right)} \quad (4)$$

The dominance of a rule measures how much better the rule predicts its assigned class compared to other classes [13], [14]:

$$Dominance(A_j \rightarrow C_j) = Confidence(A_j \rightarrow C_j) - \max_{k \neq j} Confidence(A_j \rightarrow C_k) \quad (5)$$

where k ≠ c_j

The initial weight of each rule is computed as the product of its support and confidence [13], [14]:

$$w(R_j) = Support(A_j \rightarrow C_j) \times Confidence(A_j \rightarrow C_j) \quad (6)$$

We use minimum thresholds to filter the rule base: support ≥ 0.03, confidence ≥ 0.75, and dominance ≥ 0.50. Rules that do not meet these criteria are discarded.

4) Genetic Algorithm Optimization

After generating the initial set of fuzzy rules, we use a genetic algorithm (GA) [14], [16] to optimize the rule base. The GA operates on a population of 30 candidate solutions over 30 generations. While multiobjective approaches like NSGA-II [17] could simultaneously optimize multiple performance metrics, we employ a single-objective GA that optimizes the geometric mean to balance both classes. The fitness of a solution is evaluated using the geometric mean (GM) of the true positive rate (TPR) and the true negative rate (TNR) [18] over training data:

$$GM = \sqrt{TPR \times TNR} \quad (7)$$

The GA uses tournament selection [16] to select parents for reproduction:

$$Select(P) = \underset{x_i \in T}{\operatorname{argmax}} f(x_i) \quad (8)$$

$where\ i \in T, T \subset P, |T| = k$, P is the population, T is a randomly selected tournament of size $k=3$, and $f(x_i)$ is the fitness of individual i. For crossover, we apply BLX-α crossover [19] (α = 0.5, probability = 0.9) to real-valued components (membership function parameters and rule weights) [13], [14]:

$$c_i \sim U([min(P1_i, P2_i) - \alpha \cdot d_i, max(P1_i, P2_i) + \alpha \cdot d_i]) \quad (9)$$

where $d = |P1_i - P2_i|$ is the distance between parent values, and U denotes the uniform distribution. For binary components (rule selection), we use uniform crossover with the same probability (0.9). For mutation, we apply Gaussian mutation [16] (probability = 0.1) to real-valued components:

$$x_i' = x_i + N(O, \sigma^2) \quad (10)$$

where σ = 0.1 is the mutation strength. For binary components, bit-flip mutation is applied with probability 0.1 and a flip rate of 0.05 (5% of bits). The top 2 individuals are preserved through elitism. While our approach focuses on spike detection (binary classification of spike vs. non-spike events), future extensions could incorporate spike classification methods [20] to distinguish single unit from multi-unit activity, enabling more detailed analysis of neural population dynamics.

B. Prediction Stage

In the prediction/classification stage, the trained fuzzy classifier is used to classify new, unseen MEA data. For a given test sample, the classification decision is made by aggregating the outputs of all fuzzy rules using a weighted voting scheme. The final classification score for each class is the weighted sum of the match degrees for all rules predicting that class [13], [14]:

$$Score(C) = \sum_{R_i \to C} \left(w_i \times match(x, R_i) \right) \qquad (11)$$

where the sum is over all rules R_i that predict class C, and w_i is the optimized weight of rule R_i. The sample is assigned to the class with the highest score.

IV. Experiments and Results

A. Dataset Description

Data were recorded using the MaxOne high-density MEA system (MaxWell Biosystems, Switzerland); however, our spike detection methodology is device-independent and applicable to any MEA recording platform. Our experiments were conducted on the dataset of MEA recordings from five different experimental conditions. MEA chip contains a subset of electrodes (ranging from 141 to 154 electrodes) arranged in a grid pattern (see microscopic view in Fig. 3) The dataset includes recordings from different numbers of chips (2 and 6 chips pooled together), mature neural network activity measured in days in vitro (DIV21 and DIV22), and different stimulation protocols (amplitude-based and frequency-based). For each stimulation event, the neural network's response is captured as a 142-dimensional feature vector, where each dimension represents the firing rate of one electrode during the 30-80ms post-stimulation response window. Each experiment uses a balanced distribution of spike and non-spike events to ensure unbiased classifier training. The raw dataset exhibits class imbalance that varies across chips, reflecting the biological diversity of neural cultures [21]. For example, our 6-chip pooled dataset contains 25.2% spike events and 74.8% non-spike events. This imbalance arises from differences in neuron density, synapse formation patterns, and spontaneous activity levels across different biological preparations, representing the inherent stochasticity of living neural networks.

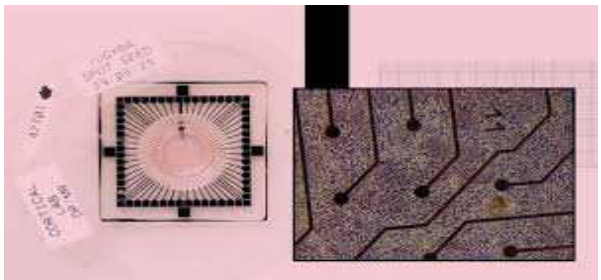

Fig. 3. Microscopic view of the living neural network. A culture of mouse cortical neurons growing over the high-density microelectrode array (MEA) grid. The dark circles represent the electrodes that capture the extracellular electrical activity (spikes) used in the dataset [22].

TABLE I: SUMMARY OF THE DATASETS IS PROVIDED.

Experiment	Chips	Total Samples	Train	Test	Features	Rules	Protocol
2chips	2	11,916	9,532	2,384	20	549	Amplitude
6chips	6	28,060	22,448	5,612	20	868	Amplitude
frequency_0	4	7,236	5,788	1,448	20	761	Frequency
DIV22_only	4	2,150	1,720	430	20	662	Amplitude
DIV21_only	3	25,910	20,728	5,182	20	726	Amplitude

The dataset sizes range from 2,150 to 28,060 samples across the five experimental conditions (see Table I for details).

B. Multi-Chip Pooled Validation Protocol

To ensure the generalization capability of our model, we employed a multi-chip pooled validation protocol with class balancing. For each experiment, data from MEA chip (see experimental platform in Fig. 4) were combined into a single dataset. The raw dataset exhibits class imbalance (25.2% spike events in the 6-chip dataset). We partitioned the data into training (80%) and testing (20%) sets using stratified sampling. Undersampling was applied exclusively to the training set to achieve 50-50 class balance, while the test set was kept at its natural imbalanced distribution for out-of-sample evaluation under realistic conditions. This approach evaluates the classifier's ability to generalize across samples from multiple chips with diverse characteristics. The random stratified split ensures that both training and testing sets contain representative samples from all chips, allowing the classifier to learn from diverse chip characteristics. The overall performance of the classifier is tested on out of sample unseen samples from the same chip population.

Fig. 4. MaxOne MEA recording platform (MaxWell Biosystems) with single-well High-Density Microelectrode Array (HD-MEA) electrophysiology platform for recording and stimulating electrogenic cells in vitro [23].

To ensure reproducibility, all stochastic components of our pipeline (train/test split, ANNIGMA [15] neural network initialization and training, and genetic algorithm operations) were initialized with a fixed random seed (seed=42). This ensures that our results are deterministic and can be replicated by other researchers.

C. Evaluation Metrics

We evaluated the performance of our classifier using a comprehensive set of metrics [18] including geometric mean (GM) and F1-score (as shown in Table II). These metrics provide a complete picture of the model's performance, including its ability to correctly identify both spikes and non-spike events.

Our proposed fuzzy rule-based classifier achieved excellent performance across all five experiments, demonstrating its robustness and generalization capabilities. This performance is particularly notable when compared to standard Logistic Regression used in similar biological computing setups, which have reported recognition accuracies of approximately 80% on complex pattern recognition tasks [24], [25]. Unlike these models, our fuzzy classifier provides transparent, interpretable rules that explain the underlying spike detection decisions.

TABLE II: THE ACHIEVED KEY RESULTS.

Experiment	GM	F1 (Macro) (%)	Rules
2chips	0.9773	97.74	549
6chips	0.9611	96.12	868
frequency_0	0.9389	93.90	761
DIV22_only	0.9512	95.11	662
DIV21_only	0.9540	95.40	726
Mean ± Std	0.9565 ± 0.0140	95.66 ± 1.27	-

The classifier achieved a mean F1-score of 95.69% ± 1.24% for spike detection and 95.62% ± 1.30% for no-spike detection, with a macro-average F1-score of 95.66% ± 1.27%. These balanced F1-scores across both classes demonstrate that the model does not exhibit bias towards either spike or no-spike events, which is crucial for reliable neural activity analysis. Notably, the 2chips experiment, which had 11916 samples, achieved the highest F1-score of 97.74%.

The slight decrease in F1 score with more chips reflects the increased diversity and variability across different MEA chips, demonstrating that our approach can generalize well across multiple chips despite their individual characteristics.

D. Rule Interpretability and Analysis

A key advantage of our fuzzy rule-based classifier is its interpretability. The model generates a set of human-readable IF-THEN rules (with small number of antecedents) that provide insights into the classification logic. The complexity of the rule base, as shown in Fig. 5, remains manageable across the different experiments, with the number of rules ranging from 549 to 868. The rule density, measured as the number of rules per 1000 samples, varies from 27.5 to 308.5, with the DIV22_only experiment having the highest density due to its smaller sample size.

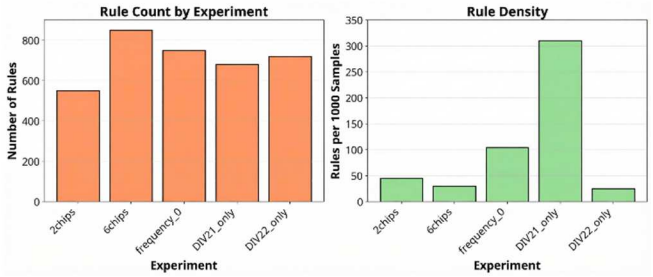

Fig. 5. Analysis of the rule complexity, showing the number of rules generated and the rule density per 1000 samples for each experiment.

To illustrate the interpretability of our approach, we present several examples of fuzzy rules extracted from the 2chips experiment:

Rule 1 (No-Spike, Weight: 0.45):
IF Electrode_6 is Medium AND Electrode_15 is Medium
THEN No-Spike
(Confidence: 99.1%, Support: 34.1%)

Rule 2 (No-Spike, Weight: 0.39):
IF Electrode_15 is Medium AND Electrode_17 is Medium
THEN No-Spike
(Confidence: 99.4%, Support: 29.6%)

Rule 3 (Spike, Weight: 0.10):
IF Electrode_13 is Medium AND Electrode_17 is Medium
THEN Spike
(Confidence: 100.0%, Support: 9.2%)

Rule 4 (Spike, Weight: 0.09):
IF Electrode_6 is Medium AND Electrode_17 is Medium
THEN Spike
(Confidence: 100.0%, Support: 9.2%)

These rules are easy for stakeholders to create associations between spike detection and electrode state, allowing biocomputer engineers to verify the computer task performance against the distribution of electrode states set. This is a step forward in decoding the living biocomputers performance towards functional verification stages, in the direction closer to reliability in conventional silicon-based computers. The high confidence values indicate consistent patterns in the training data. These interpretable rules enable researchers to validate the classifier's logic (e.g., Rules 1-2 characterize baseline states while Rules 3-4 indicate spike generation), generate hypotheses about network topology (e.g., Electrode_17 appears in multiple rules, suggesting a critical functional role), and facilitate debugging by identifying potential issues in feature extraction or training data quality.

V. Conclusions and Future Work

In this paper, we presented a fuzzy rule-based classifier for the detection of neural spikes in MEA recordings. Our approach successfully combines the high accuracy of modern machine learning techniques with the interpretability of fuzzy based systems. By integrating ANNIGMA-based feature selection and a genetic algorithm for weight optimization, we have developed a system that achieves high accuracy while generating a set of human-readable fuzzy rules. Like interpretable fuzzy classifiers in medical diagnosis [26], [27], our approach provides transparency in the decision-making process, which is crucial for building trust in AI-assisted neuroscience research. This reflects the broader paradigm of building interpretable, human-understandable intelligent systems across diverse domains [28]. Crucially, this interpretability is a prerequisite for advancing "living computer" technology from experimental curiosities to reliable computational platforms, allowing bioengineers to distinguish between true biological computation and experimental artifacts and guiding the engineering of more stable and predictable biocomputing architectures. Future work will explore the use of type-2 fuzzy logic to better handle the uncertainties inherent in MEA data, as well as a real-time implementation for online spike detection and closed-loop experiments within living computers.

References

[1] J. Küchler et al., "Engineered biological neuronal networks as basic logic operators," Frontiers in Computational Neuroscience, vol. 19, Article 1559936, April 2025.

[2] B. Genocchi et al., "Astrocytes induce desynchronization and reduce predictability in neuron-astrocyte networks cultured on microelectrode arrays," Royal Society Open Science, vol. 11, no. 10, Article 240839, October 2024.

[3] S. M. Potter, "Distributed processing in cultured neuronal networks," in Progress in Brain Research: Advances in Neural Population Coding, vol. 130, M. A. L. Nicolelis, Ed. Amsterdam, The Netherlands: Elsevier, 2001, pp. 49-62.

[4] Y. Nam and B. C. Wheeler, "In vitro microelectrode array technology and neural recordings," Critical Reviews in Biomedical Engineering, vol. 39, no. 1, pp. 45-61, January 2011.

[5] M. E. J. Obien, K. Deligkaris, T. Bullmann, D. J. Bakkum, and U. Frey, "Revealing neuronal function through microelectrode array recordings," Frontiers in Neuroscience, vol. 8, Article 423, January 2015.

[6] T. B. DeMarse, D. A. Wagenaar, A. W. Blau, and S. M. Potter, "The neurally controlled animat: Biological brains acting with simulated bodies," Autonomous Robots, vol. 11, no. 3, pp. 305-310, November 2001.

[7] S. Gibson, J. W. Judy, and D. Marković, "Spike sorting: The first step in decoding the brain," IEEE Signal Processing Magazine, vol. 29, no. 1, pp. 124-143, January 2012.

[8] J. O. Muthmann, H. Amin, E. Sernagor, A. Maccione, D. Panas, L. Berdondini, U. S. Bhalla, and M. H. Hennig, "Spike detection for large neural populations using high density multielectrode arrays," Frontiers in Neuroinformatics, vol. 9, Article 28, 2015.

[9] N. Sahab, H. Hagras, "Adaptive Non-singleton Type-2 Fuzzy Logic Systems: A Way Forward for Handling Numerical Uncertainties in Real World Applications", International Journal of Computers, Communications and Control, Vol.5, No.3, pp.503-529, December 2011.

[10] K. Almohammadi, H. Hagras, B. Yao, A. Alzahrani, D.Alghazzawi, G. Aldabbagh,"A Type-2 Fuzzy Logic Recommendation System for Adaptive Teaching" Soft Computing, Vol.21, No.4. pp.965-967, 2017.

[11] F. Doctor, H. Hagras, "Neuro type-2 fuzzy based method for decision making". US Patent 8515884.

[12] L. A. Zadeh, "Fuzzy sets," Information and Control, vol. 8, no. 3, pp. 338-353, June 1965.

[13] J. A. Sanz, D. Bernardo, F. Herrera, H. Bustince, and H. Hagras, "A compact evolutionary interval-valued fuzzy rule-based classification system for the modeling and prediction of real-world financial applications with imbalanced data," IEEE Transactions on Fuzzy Systems, vol. 23, no. 4, pp. 973-990, August 2015.

[14] J. A. Sanz, A. Fernandez, H. Bustince, and F. Herrera, "A genetic tuning to improve the performance of fuzzy rule-based classification systems with interval-valued fuzzy sets: Degree of ignorance and lateral position," International Journal of Approximate Reasoning, vol. 52, no. 6, pp. 751-766, September 2011.

[15] C.-N. Hsu, H.-J. Huang, and S. Dietrich, "The ANNIGMA-wrapper approach to fast feature selection for neural nets," IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics, vol. 32, no. 2, pp. 207-212, April 2002.

[16] A. E. Eiben and J. E. Smith, Introduction to Evolutionary Computing, 2nd ed. Berlin, Germany: Springer, 2015.

[17] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," IEEE Transactions on Evolutionary Computation, vol. 6, no. 2, pp. 182-197, April 2002.

[18] D. M. W. Powers, "Evaluation: From precision, recall and F-measure to ROC, informedness, markedness and correlation," Journal of Machine Learning Technologies, vol. 2, no. 1, pp. 37-63, January 2011.

[19] F. Herrera, M. Lozano, and J. L. Verdegay, "Tackling real-coded genetic algorithms: Operators and tools for behavioural analysis," Artificial Intelligence Review, vol. 12, no. 4, pp. 265-319, August 1998.

[20] A. Tankus, Y. Yeshurun, and I. Fried, "An automatic measure for classifying clusters of suspected spikes into single cells versus multiunits," Journal of Neural Engineering, vol. 6, no. 5, Article 056001, October 2009.

[21] G. M. Weiss, "Mining with rarity: A unifying framework," ACM SIGKDD Explorations Newsletter, vol. 6, no. 1, pp. 7-19, June 2004.

[22] C. Criddle, "Inside the British lab growing a biological computer," Financial Times, 2025. [Online]. Available: https://www.ft.com/content/713eab47-a1f1-4477-a7de-f2b150e2faac

[23] MaxWell Biosystems, "MaxOne: Multi-well high-resolution microelectrode array system," MaxWell Biosystems, 2024. [Online]. Available: https://www.mxwbio.com/resources/maxone-guide

[24] K. Vulić and J. Küchler, "Dataset for engineered biological neuronal networks as basic logic operators," ETH Zurich, Research Collection, April 2025.

[25] J. Küchler, K. Vulić, H. Yao, C. Valmaggia, S. J. Ihle, S. Weaver, and J. Vörös, "Engineered biological neural networks as basic logic operators—Official code repository," GitHub, 2024. [Online]. Available: https://github.com/Hullimulli/engineered_biological_neural_networks_as_basic_logic_operators

[26] D. Nauck and R. Kruse, "Obtaining interpretable fuzzy classification rules from medical data," Artificial Intelligence in Medicine, vol. 16, no. 2, pp. 149-169, June 1999.

[27] I. Gadaras and L. Mikhailov, "An interpretable fuzzy rule-based classification methodology for medical diagnosis," Artificial Intelligence in Medicine, vol. 47, no. 1, pp. 25-41, September 2009.

[28] V. Callaghan, M. Colley, H. Hagras, J. Chin, F. Doctor, G. Clarke, "Programming iSpaces: A tale of two paradigms", "Intelligent Spaces: The Application of Pervasive ICT", (Eds: A. Steventon, S. Wright), Springer-Verlag, Chapter 24, pp.389-421, December 2005.