UnimNeuron: Automatic Connective Selection for Interpretable Neuro-Fuzzy Rules

Paulo Vitor de Campos Souza
NOVA IMS, Universidade Nova de Lisboa
Email: psouza@novaims.unl.pt

***Abstract*—In evolving data streams, neuro-fuzzy systems must adapt to non-stationary distributions while preserving interpretability and numerical stability. However, most evolving fuzzy models rely on fixed logical connectives or implicitly encode logical behavior, limiting transparency and flexibility.**

This paper introduces the *UnimNeuron*, a neuro-fuzzy neuron whose antecedent connective is automatically selected from data using an order-invariant unimum-based aggregation. The proposed neuron dynamically exhibits conjunctive (AND), disjunctive (OR), or compensatory (COMP) behavior without operator search or manual tuning, and provides an explicit index to inspect the logical regime of each rule.

Based on this formulation, two evolving neuro-fuzzy classifiers are proposed and evaluated under a prequential protocol on real and synthetic data streams with concept drift. Experimental results show that the proposed models achieve predictive performance statistically comparable to state-of-the-art evolving fuzzy systems, while maintaining a compact rule base, zero numerical instabilities, and interpretable fuzzy rules. An ablation study further confirms the relevance of automatic connective selection across diverse scenarios.

***Index Terms*—Fuzzy logic, neuro-fuzzy systems, uninorms, interpretability, aggregation operators**

I. INTRODUCTION

In recent years, evolving systems have become increasingly important in scenarios where data continuously evolves, such as in real-time decision-making or dynamic environments. These systems must not only adapt to the evolving nature of data but also preserve the ability to make predictions and maintain transparency, especially in safety-critical applications. In these contexts, interpretability becomes as crucial as predictive performance, as understanding the rationale behind the decisions made by a model is essential for user trust and system validation.

Most existing approaches to evolving neuro-fuzzy systems suffer from a fundamental limitation: the aggregation operators used for combining inputs and generating fuzzy rules are either fixed or manually tuned, making the underlying logical semantics opaque and difficult to interpret. Moreover, many systems rely on pre-defined logical connectives, such as AND or OR, which may not be flexible enough to represent the complexities inherent in real-world reasoning. For example, situations where compensatory reasoning is required—where evidence does not follow a strict conjunctive or disjunctive relationship—are often poorly handled by traditional models.

To address this limitation, we propose the UnimNeuron, a novel neuro-fuzzy neuron that automatically adapts the logical connective of the antecedent based on the input data. The UnimNeuron utilizes the unimum operator, a generalization of uninorms, to dynamically select between three logical regimes: AND-like, OR-like, and compensatory (COMP). This approach eliminates the need for manual tuning of control parameters while preserving interpretability, as the logical regime of each rule is explicitly determined by the data rather than fixed hyperparameters.

The main objectives of this paper are as follows:

- To introduce the UnimNeuron, the first neuro-fuzzy neuron based on the unimum operator, capable of automatic connective selection.
- To show how the UnimNeuron combines flexibility with interpretability in an evolving neuro-fuzzy system, addressing the challenge of logical adaptability in non-stationary environments.
- To validate the performance of the UnimNeuron through experiments on a diverse set of real and synthetic data streams, demonstrating its robustness, stability, and interpretability compared to existing methods.

The primary contribution of this paper to the state of the art is the introduction of a novel, interpretable fuzzy neuron that combines automatic connective selection with the preservation of logical semantics. This enables more flexible reasoning in evolving systems without sacrificing interpretability. By employing the unimum operator, the UnimNeuron can seamlessly switch between different logical regimes (AND, OR, and COMP) based on the nature of the input data, providing a more dynamic and robust approach to modeling real-world decision-making processes.

The rest of this paper is organized as follows: Section II reviews the related work in fuzzy operators and logic neurons. Section III provides a formal definition of the UnimNeuron and discusses its key properties. Section IV introduces the two proposed evolving neuro-fuzzy models based on the UnimNeuron. Section V presents the experimental setup and results, followed by a discussion in Section VI. Finally, Section VII concludes the paper and outlines future research directions.

II. Related Work

A. Fuzzy Operators

Fuzzy aggregation operators are the mathematical core of fuzzy logic systems, enabling the construction of logical connectives over the unit interval. Classical operators include t-norms (generalizing fuzzy conjunctions) and t-conorms (generalizing disjunctions) [1]. Although powerful, they restrict the logical behavior to either purely conjunctive or purely disjunctive regimes.

To overcome this rigidity, generalizations have been introduced. Uninorms [2] extend t-norms and t-conorms by introducing a neutral element $g \in (0, 1)$, which allows interpolation between AND-like and OR-like behaviors. Similarly, nullnorms [3] rely on an annihilating element to provide greater expressive capacity. Further extensions, such as n-uninorms [4], add additional degrees of freedom but also increase complexity.

Most recently, Gürdal et al. [5] proposed the *unimum* operator (also known as unima), which goes beyond uninorms by explicitly introducing a third regime: a compensatory connective (COMP). This provides a principled mathematical framework where conjunctive, disjunctive, and compensatory behaviors coexist, depending on the relation of the input values to the neutral element. This innovation motivates the design of new neuro-fuzzy neurons that can exploit the richer semantics of the unimum.

B. Logic Neurons

The use of fuzzy operators has naturally extended to the design of fuzzy neurons, interpretable computational units that implement fuzzy IF/THEN rules [6]. Early proposals include the AND-neuron and OR-neuron [7], which aggregate inputs using t-norms and t-conorms, respectively. While transparent and simple, these neurons impose a fixed logical connective, limiting flexibility.

To enhance expressiveness, Yager et al. introduced the Unineuron [2], which leverages the uninorm framework. By means of the neutral element g, the Unineuron can interpolate between conjunctive and disjunctive regimes. Calvo et al. [3] extended this concept with the Nullneuron, relying on nullnorms and an annihilating element u to further expand the space of logical connectives. These neurons operate in two levels: a local transformation of inputs and weights, followed by global aggregation through the respective operator [8]. While powerful, these models suffer from reduced interpretability: the connective of a fuzzy rule depends on hyperparameters (g or u), which must be manually chosen or tuned.

C. Limitations of Existing Neurons

Although unineurons and nullneurons increased flexibility, they introduced interpretability challenges. Determining whether a given rule behaves as AND-like or OR-like requires analyzing operator parameters, making the extracted knowledge less transparent. Moreover, compensatory behaviors, which are common in real-world reasoning, are not explicitly modeled by these structures. This motivates the development of neurons where logical regimes emerge automatically from data, rather than being imposed by parameter tuning.

D. Research Gap and Contribution

Despite significant progress, no neuro-fuzzy neuron has been defined directly on top of the unimum operator. As a logical generalization of uninorms, the unimum naturally supports AND, OR, and COMP regimes, ensuring that rule connectives are determined by input values relative to the neutral element. This paper introduces the *UnimNeuron*, the first neuro-fuzzy unit based on the unimum. By construction, the UnimNeuron achieves *automatic connective selection*, thereby combining flexibility with interpretability in a principled way.

III. The UnimNeuron: Formal Definition and Properties

A. The UnimNeuron: Formal Definition and Properties

This section formalizes the UnimNeuron, a fuzzy neuron whose antecedent connective is an *order-invariant unimum* derived from recent Uninorm–Unima theory [5]. Its design separates (i) a local relevancy transform p and (ii) a global n-ary connective $M_e^{(n)}$, both governed by a data-driven neutral element $e \in (0, 1)$.

a) Preliminaries.: Let $\mathbf{a} = [a_1, \ldots, a_n] \in [0, 1]^n$ denote fuzzy inputs (e.g., Gaussian memberships), and $\mathbf{w} = [w_1, \ldots, w_n] \in [0, 1]^n$ be connection weights. A relevancy-enhancing map p (Eq. (1)) produces local signals

$$\flat_i = p(a_i, w_i, e) \in [0, 1], \qquad i = 1, \ldots, n, \tag{1}$$

where e acts as the connective's neutral reference.

b) Binary unimum (prototype).: For $x, y \in [0, 1]$ and fixed $e \in (0, 1)$, define

$$M_e(x, y) = \begin{cases} T(x, y), & x \le e,\ y \le e \quad (\textit{AND-like}), \\ S(x, y), & x \ge e,\ y \ge e \quad (\textit{OR-like}), \\ C(x, y), & \text{otherwise} \qquad (\textit{Compensatory}), \end{cases} \tag{2}$$

with the interpretable choices

$$\begin{aligned} T(x, y) &= \min\{x, y\}, \\ S(x, y) &= \max\{x, y\}, \\ C(x, y) &= \tfrac{x+y}{2}. \end{aligned} \tag{3}$$

The operator behaves as a pure conjunctive or disjunctive connective in the homogeneous regimes, and as a compensatory operator when inputs cross the neutral boundary. Note that C is not idempotent; idempotency holds only in the extremal branches (as desired for clear interpretability).

c) Order-invariant n-ary unimum.: Following Unima theory, we define the n-ary aggregation *directly*—avoiding non-associative binary folding:

$$M_e^{(n)}(\mathbf{z}) = \begin{cases} \min_i z_i, & z_i \le e\ \forall i, \\ \max_i z_i, & z_i \ge e\ \forall i, \\ \frac{1}{n}\sum_{i=1}^{n} z_i, & \text{otherwise}, \end{cases} \qquad \mathbf{z} = (z_1, \ldots, z_n) \in [0, 1]^n. \tag{4}$$

This operator is symmetric (permutation-invariant), internal, and reproduces the three logical regimes of the binary prototype without order dependence.

d) UnimNeuron output.: With $\flat_i = p(a_i, w_i, e)$, the UnimNeuron output is

$$z = UNIM(\mathbf{w}; \mathbf{a}; e) = M_e^{(n)}\big(\flat_1, \ldots, \flat_n\big). \tag{5}$$

e) Automatic connective selection.: Define

$$\mathcal{I}_{\leq} = \{i : \flat_i \leq e\}, \qquad \mathcal{I}_{\geq} = \{i : \flat_i \geq e\}.$$

Then

$$\begin{aligned} &\text{AND-like if } \mathcal{I}_{\leq} = \{1, \ldots, n\}, \\ &\text{OR-like if } \mathcal{I}_{\geq} = \{1, \ldots, n\}, \\ &\text{COMP otherwise.} \end{aligned} \tag{6}$$

Thus, the logical regime of the rule antecedent is *data-driven*.

f) Interpretability index.: We quantify the connective tendency using

$$\rho \;=\; \tfrac{1}{n} \sum_{i=1}^{n} \left[\nVdash\{\flat_i \geq e\} - \nVdash\{\flat_i \leq e\}\right] \in [-1, 1], \tag{7}$$

so that $\rho \approx +1$ (OR-like), $\rho \approx -1$ (AND-like), and $\rho \approx 0$ (compensatory / mixed evidences).

1) Fundamental Properties:

Lemma 1 (Internality)**.** *For all* $x, y \in [0, 1]$*,* $\min\{x, y\} \leq M_e(x, y) \leq \max\{x, y\}$*. Likewise,* $\min_i z_i \leq M_e^{(n)}(\mathbf{z}) \leq \max_i z_i$*.*

Lemma 2 (Monotonicity and symmetry)**.** *Both* M_e *and* $M_e^{(n)}$ *are nondecreasing in each argument and permutation-invariant.*

Lemma 3 (Neutral-element behavior)**.** *For all* $x \in [0, 1]$*,* $M_e(x, e) = M_e(e, x) = x$*. In the* n*-ary operator,* e *acts as a regime-separating pivot.*

Proposition 1 (Regime decomposition)**.** *If* $z_i \leq e$ *for all* i*, then* $M_e^{(n)}(\mathbf{z}) = \min_i z_i$*. If* $z_i \geq e$ *for all* i*, then* $M_e^{(n)}(\mathbf{z}) = \max_i z_i$*. Otherwise,* $M_e^{(n)}(\mathbf{z}) = \frac{1}{n}\sum_i z_i$*.*

Corollary 1 (Extremal limits)**.** *As* $e \uparrow 1$*,* $M_e^{(n)}(\mathbf{z}) \to \min_i z_i$ *(AND limit). As* $e \downarrow 0$*,* $M_e^{(n)}(\mathbf{z}) \to \max_i z_i$ *(OR limit).*

a) Remark (continuity).: Within each branch of (4), $M_e^{(n)}$ is continuous and internal. At regime boundaries ($z_i = e$), the value is well-defined, although global continuity is not enforced by the $\{\min, \max, \text{mean}\}$ construction. Generator-based uninorms may be substituted for full continuity, as discussed in [5].

IV. Proposed Evolving Neuro-Fuzzy Models

This section introduces two novel evolving neuro-fuzzy classifiers based on uninorm-driven fuzzy neurons and ADPA-style prototype evolution. Both models share a common architectural backbone grounded in the ENFS-Uni0 framework and the interpretability fundamentals discussed previously, such as the use of Gaussian evidences, uni-null based logical connectives and regime preservation (AND/OR/COMP) in the antecedent layer. The two proposed variants differ only in their *local update principles*, yielding distinct stability–plasticity behaviors while maintaining full compatibility with the unified rule structure.

A. Common Architecture of the Proposed Models

Both models follow the same generic evolving rule-based structure. Each rule is associated with a data cloud generated by an autonomous partitioning procedure derived from ADPA. Let $x \in \mathbb{R}^d$ be the current sample. A new rule is created when x is sufficiently distant from the existing prototypes with respect to the Chebyshev metric and the global scatter statistics. Otherwise, the nearest prototype is updated incrementally. This fully on-line strategy ensures single-pass learning and local adaptation to drifting concepts.

1) Antecedent Layer: For each rule r, the antecedent part computes per-feature Gaussian evidences

$$a_{r,j}(x_j) = \exp\!\left(-\frac{(x_j - c_{r,j})^2}{2\sigma_r^2}\right), \tag{8}$$

where $c_{r,j}$ is the j-th component of the rule center and σ_r is a scalar rule radius.

a) Feature relevance weighting.: In the proposed framework, feature relevance weights w_j play different roles depending on the learning strategy adopted. In the ENF–Unim$_{\text{Safe}}$ variant, the weights are adapted online using a local separability criterion, inspired by Lughofer's evolving fuzzy systems, allowing each rule to emphasize the most informative variables while preserving semantic completeness. Conversely, in the ENF–Unim$_{\text{PA}}$ model, feature weights are kept fixed ($w_j = 1$), and discrimination is handled exclusively by the passive–aggressive consequent learner. This design avoids competing adaptation mechanisms and ensures numerical stability while retaining interpretable logical behavior.

These evidences are combined using a uninorm-based neuro-fuzzy connective. Following the interpretability principles established in the original UnimNeuron formulation, a uni-null uninorm is used to construct a neuron whose output

$$y_r(x) = \mathcal{U}(a_{r,1}(x_1), \ldots, a_{r,d}(x_d);\, w_r, g_{1,r}, g_{2,r}, z_r) \tag{9}$$

represents the logical aggregation regime (AND/OR/COMP) associated with rule r. The parameters $(w_r, g_{1,r}, g_{2,r}, z_r)$ yield a compact and interpretable characterization of local behavior, as previously discussed in the manuscript.

2) Consequent Layer: Given R rules, the global consequent is computed through a zero-order Takagi–Sugeno model:

$$s(x) = \phi(x)^\top W, \qquad \phi(x) = [1,\, y_1(x), \ldots, y_R(x)]^\top, \tag{10}$$

where $W \in \mathbb{R}^{(R+1)\times C}$ holds the class scores for C classes. The standard ENFS-Uni0 uses Recursive Least Squares (RLS) with forgetting factor to update W. In the proposed models, this component is replaced or modified to improve numerical stability and robustness, as detailed next.

3) Rule Pruning and Structural Control: A lightweight similarity-based pruning operator is employed to merge newly created rules that are redundant, maintaining interpretability and avoiding unbounded growth. Additionally, a maximum rule budget $R_{\max}$ enforces compactness—whenever exceeded, the least relevant rule is removed. These mechanisms preserve a stable, small and readable rule base across long data streams.

B. ENF-Unim-Safe: Safe Local Update Mechanism

The first proposed model, ENF-Unim-Safe, introduces a *safe local update* strategy inspired by the stability principles advocated for evolving TS models. While the general architecture follows the structure above, the Safe variant modifies *how* antecedent and consequent parameters are updated in order to prevent abrupt, unjustified shifts in rule behavior.

Let $\Delta\theta_r$ denote the candidate update of any antecedent parameter of rule r (e.g., c_r, σ_r, or neuron parameters). The Safe model accepts the update only if

$$\|\Delta\theta_r\| \leq \eta_s \quad \text{and} \quad |y_r(x) - y_r^{\text{old}}(x)| \leq \gamma_s, \tag{11}$$

where η_s bounds the parameter drift and γ_s bounds the change in logical activation. This constraint preserves the interpretability regime of the rule, ensuring that AND-like rules remain AND-like and preventing sudden logical inversions unrelated to genuine distributional shifts.

At the consequent level, a damped RLS update

$$W \leftarrow W + \alpha_s K\, e^\top \tag{12}$$

is performed, where e is the instantaneous residual and $\alpha_s < 1$ attenuates numerically aggressive corrections. This results in a stable, locally consistent adaptation behavior.

C. ENF-Unim-PA: Passive–Aggressive Neuro-Fuzzy Update

The second proposed model replaces the RLS update entirely with a Passive–Aggressive (PA) regression step. This eliminates the need for matrix inversion and drastically increases numerical robustness in rapidly changing data streams.

Let $y \in \{1, \ldots, C\}$ be the true class and $\hat{s}_y(x)$ the predicted score for the correct class. The PA update solves

$$\min_W \frac{1}{2}\|W - W_{\text{old}}\|^2 \quad \text{s.t.} \quad \hat{s}_y(x) - \hat{s}_{\neg y}(x) \geq 1, \tag{13}$$

yielding the closed-form step

$$W \leftarrow W + \tau(t - s(x))\phi(x)^\top, \tag{14}$$

where t is the one-hot target and the step size is

$$\tau = \min\left(C_{\text{PA}},\, \frac{1 - (t - s(x))^\top \phi(x)}{\|\phi(x)\|^2}\right). \tag{15}$$

The constant C_{PA} controls the aggressiveness of adaptation. When the margin is already satisfied, $\tau = 0$, ensuring a purely passive response.

This PA-based adaptation is naturally robust to abrupt drifts, avoids covariance explosions typical of RLS in noisy streams, and keeps the model fast even for large rule bases. Importantly, the antecedent interpretability mechanisms remain unchanged, ensuring compatibility with the unified uninorm-based rule semantics.

V. Experimental Setup

A. Experiment 3: Ablation Study on UnimNeuron Components

This experiment isolates and quantifies the contribution of the main components of the proposed UnimNeuron-based models by evaluating internal design choices of the antecedent, rather than comparative performance against other evolving fuzzy systems (Experiment 2). Two variants, ENF_Unim_Safe and ENF_Unim_PA, are assessed under five ablation settings: FULL (complete UnimNeuron with adaptive connective and feature relevance weights), FIXED_AND, FIXED_OR, and FIXED_COMP (antecedent aggregation fixed to min, max, and mean, respectively), and NO_W (feature relevance weights disabled, $w_j = 1$). All remaining components, including rule creation, merging, update mechanisms, and hyperparameters, are kept identical to Experiment 2 to ensure a controlled comparison. Experiments are conducted on diverse real and synthetic data streams using the same prequential protocol and evaluation metrics described in Section II.

B. Full stream evaluation

This section describes the streaming benchmarks, competing models, parameter configurations, and evaluation protocol employed in the large-scale study of evolving fuzzy systems (EFS). All experiments follow the same implementation pipeline defined in the public experiment script provided with this work (cf. source code in [9]).

1) Data Streams: The models are evaluated on a diverse set of real and synthetic non-stationary data streams implemented in the `river` framework. Real datasets include Elec2 (binary classification with stochastic supply–demand drift), CreditCard (highly imbalanced fraud detection), Phishing (heterogeneous web page classification), and Bananas (a classic nonlinearly separable benchmark for rule interpretability). Synthetic streams provide controlled drift scenarios and include ConceptDrift-Agrawal (gradual drift), Agrawal (stationary nonlinear relations), RandomRBFDrift (fast stochastic drift), and Sine (smooth periodic drift). All datasets are evaluated in complete mode with a maximum number of samples ranging from 5k to 40k, depending on the stream.

2) Models Under Comparison: We compare six state-of-the-art evolving fuzzy methods—ePL (Evolving Product Logic), ePL+ (extended variant with online adaptation), exTS (Evolving Takagi–Sugeno fuzzy system), Simpl-eTS (lightweight exTS with simplified consequents), and eMG (evolving multimodal Gaussian clusters)—with two proposed UnimNeuron-based architectures: ENF-Unim-Safe, which employs UnimNeuron antecedents with safe update restrictions, and ENF-Unim-PA, which integrates UnimNeuron antecedents with Perceptron-like adaptive consequents (PA-I) and merge-aware structural control.

All competing models are used exactly as provided by the original authors through the evolvingfuzzysystems library [1]. All proposed models share the same ADPA-style premises and

[1]https://github.com/kaikerochaalves/evolvingfuzzysystems

Gaussian adaptation presented in Sec. III, differing only in their update logic and connective selection mechanism.

C. Parameter Configuration

The two proposed models were configured to stabilize rule growth and remain competitive under severe drift. ENF-Unim-Safe uses $max\_rules = 5$, $safe\_threshold = 0.07$ to enable selective local learning, $min\_firing\_for\_update = 0.05$, a merge threshold $\tau_{merge} = 0.62$, and a drift sensitivity of 0.35. ENF-Unim-PA also enforces $max\_rules = 5$, with a reduced passive–aggressive parameter $C = 0.25$, PA-I updates enabled (use_PA_I = True), a margin threshold of 0.8, an exponential decay rate of 10^{-2}, and a merge threshold $\tau_{merge} = 0.68$. All hyperparameters of baseline models were kept at their default settings.

1) Evaluation Protocol: All streams are processed sample-by-sample under a *prequential* (test-then-train) protocol:

$$\hat{y}_t = f_{t-1}(\mathbf{x}_t), \qquad f_t \leftarrow \mathcal{U}(f_{t-1}, \mathbf{x}_t, y_t), \tag{16}$$

where $\mathcal{U}(\cdot)$ denotes the online update operator.

Input features are min–max normalized to $[0, 1]$, and labels are integer-encoded. For eFS baselines, learning follows a warmup phase of 200 samples and batch updates every 20 samples thereafter. All ENF-based models update online at each time step using the ADPA mechanism.

a) Metrics.: We report final and mean prequential accuracy, final rule count N_{rules}, and runtime. In addition, two robustness indicators are monitored.

b) Fallback predictions.: A fallback is recorded whenever a prediction is invalid (non-finite, empty, or exception raised). Let

$$I_t^{\text{fb}} = \begin{cases} 1, & \text{if prediction fails at time } t, \\ 0, & \text{otherwise.} \end{cases} \tag{17}$$

The total number of fallback predictions is

$$N_{\text{fb}} = \sum_{t=1}^{N} I_t^{\text{fb}}. \tag{18}$$

In such cases, a default label is assigned to allow uninterrupted stream processing.

c) Instability repairs.: Instability repairs count the number of internal corrective actions required to restore numerical stability during learning, such as exception handling in eFS models or resets of recursive estimators (RLS/PA) in ENF models. Let I_t^{inst} denote the corresponding indicator; the total count is

$$N_{\text{inst}} = \sum_{t=1}^{N} I_t^{\text{inst}}. \tag{19}$$

The full experimental codebase is available in [9].

VI. Discussion of Results

Table I summarizes the performance of all models across the seven streaming benchmarks. Several consistent patterns emerge.

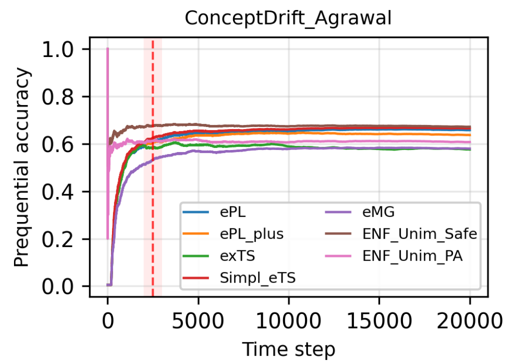

Fig. 1: Concept Drift- Accuracy trendlines evolution experiment. The vertical line indicates the drift moment.

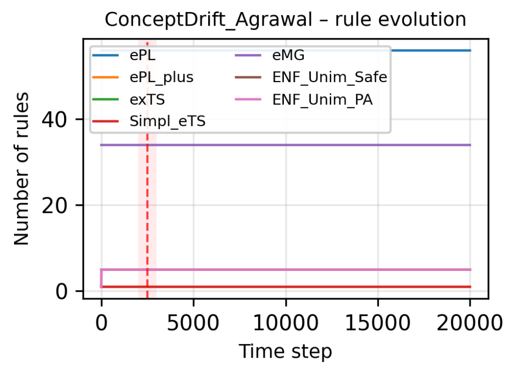

Fig. 2: Concept Drift- Rule evolution. The vertical line indicates the drift moment.

A. Overall Accuracy Trends

Across most datasets, ENF-Unim-PA achieves accuracy comparable to or exceeding the strongest classical EFS baselines while maintaining a strict upper bound of five rules. This reflects the stabilizing effect of the PA-I update combined with the adaptive AND/OR/COMP behavior of the UnimNeuron. ENF-Unim-Safe attains slightly lower accuracy but exhibits highly stable structural behavior, never exceeding five rules and showing zero instability events; its conservative updates favor robustness but reduce plasticity in highly separable streams such as Sine and Phishing. Among baselines, Simpl-eTS is consistently the strongest in stationary or slowly varying streams (e.g., Sine and Elec2) but degrades under fast drift (RandomRBF), while exTS achieves competitive accuracy at the cost of numerous numerical instabilities, highlighting the need for safety mechanisms in evolving TS models.

B. Dataset-by-Dataset Analysis

Across volatile, noisy, and drift-intensive streams, the proposed models remain competitive with the strongest baselines while enforcing a strict limit of five rules. ENF-Unim-PA consistently achieves high accuracy on volatile and smooth-drift streams (e.g., Elec2 and Sine), whereas ENF-Unim-

Safe favors stability and robustness under long-term or severe drift, achieving the best performance on the ConceptDrift benchmark. In noisy and nonlinear datasets (e.g., Phishing, Bananas, and Agrawal), both variants match or surpass baseline accuracy while maintaining a compact rule base. By contrast, strong baselines such as Simpl-eTS degrade under fast drift, and exTS attains competitive accuracy at the cost of severe numerical instability.

C. Structural and Stability Analysis

A key result is the constant rule count of the proposed models, with $N_{\text{rules}} \leq 5$ across all datasets. In contrast, ePL grows up to 62 rules, eMG exceeds 100 rules in some streams, and exTS exhibits severe numerical instability, with up to 1893 reset events. These results show that the UnimNeuron antecedent, combined with merge thresholds and PA-I learning, strongly regularizes structural growth without degrading predictive performance.

D. Statistical Analysis

To assess the statistical significance of the observed performance differences, a Friedman test (Table II) followed by a post-hoc Nemenyi test was conducted across all datasets and models.

Table III reports the pairwise p-values obtained from the Nemenyi procedure. At a significance level of $\alpha = 0.05$, no statistically significant differences are observed between the proposed models (ENF_Unim_Safe and ENF_Unim_PA) and the baseline evolving fuzzy systems. This indicates that the proposed approaches achieve performance statistically comparable to state-of-the-art methods.

TABLE II: Global statistical comparison via Friedman test (Acc$_{mean}$).

Statistic	p-value
$\chi_F^2 = 15.6956$	0.01548

TABLE III: Post-hoc Nemenyi test p-values (pairwise).

Model	ePL	ePL_plus	exTS	Simpl_eTS	eMG	ENF_Unim_Safe	ENF_Unim_PA
ePL	1.0000	0.9974	1.0000	0.2960	0.9840	0.3985	0.7763
ePL_plus	0.9974	1.0000	0.9999	0.6692	0.8081	0.7763	0.9772
exTS	1.0000	0.9999	1.0000	0.4356	0.9444	0.5523	0.8886
Simpl_eTS	0.2960	0.6692	0.4356	1.0000	0.0419	1.0000	0.9892
eMG	0.9840	0.8081	0.9444	0.0419	1.0000	0.0687	0.2654
ENF_Unim_Safe	0.3985	0.7763	0.5523	1.0000	0.0687	1.0000	0.9974
ENF_Unim_PA	0.7763	0.9772	0.8886	0.9892	0.2654	0.9974	1.0000

A single statistically significant difference is detected between Simpl_eTS and eMG ($p = 0.0419$), while all remaining comparisons yield p-values well above the significance threshold.

These results confirm that the proposed UnimNeuron-based models do not sacrifice predictive performance when compared to established evolving fuzzy classifiers, while offering additional advantages in terms of adaptive logical semantics and interpretability.

E. Ablation Study on UnimNeuron Components

The ablation study isolates the contribution of each component of the UnimNeuron-based antecedent by selectively disabling adaptive behaviors.

Fixing the logical connective reveals that no single regime universally dominates across datasets. In several streams, enforcing an AND-like behavior improves performance, particularly for the ENF-Unim-PA model, suggesting that increased selectivity in rule activation benefits margin-based learning. Conversely, OR- and COMP-like regimes perform better in scenarios where discriminative evidence is concentrated in a subset of features.

Removing feature relevance weights generally degrades interpretability and, in some cases, predictive performance, confirming that adaptive weighting plays a meaningful role beyond mere regularization.

Overall, the ablation results demonstrate that the full UnimNeuron design— combining adaptive logical regimes with feature-wise relevance weighting—offers the most consistent balance between accuracy, stability, and interpretability.

TABLE IV: Summary of ablation study results: FULL model vs. best-performing ablation (based on final accuracy). Full results are available in the public repository.

Dataset	Model	Variant	Acc$_{final}$	Acc$_{mean}$	Time [s]
Elec2	ENF$_{UnimSafe}$	NO_W	0.593	0.607	13.9
Elec2	ENF$_{UnimPA}$	FULL	0.797	0.803	17.4
Phishing	ENF$_{UnimSafe}$	FIXED_COMP	0.550	0.553	1.4
Phishing	ENF$_{UnimPA}$	FIXEDED_AND	0.578	0.576	1.5
CreditCard	ENF$_{UnimSafe}$	FULL	0.999	0.999	6.7
CreditCard	ENF$_{UnimPA}$	FULL	0.999	0.999	8.7
Bananas	ENF$_{UnimSafe}$	NO_W	0.547	0.555	6.7
Bananas	ENF$_{UnimPA}$	FIXED_AND	0.569	0.562	6.0
ImageSegments	ENF$_{UnimSafe}$	NO_W	0.138	0.147	3.1
ImageSegments	ENF$_{UnimPA}$	FIXED_AND	0.319	0.255	2.9
ConceptDrift	ENF$_{UnimSafe}$	FULL	0.672	0.673	25.8
ConceptDrift	ENF$_{UnimPA}$	FULL	0.607	0.608	33.1
RandomRBFDrift	ENF$_{UnimSafe}$	FULL	0.655	0.654	25.9
RandomRBFDrift	ENF$_{UnimPA}$	FIXED_AND	0.657	0.646	23.9
Sine	ENF$_{UnimSafe}$	NO_W	0.673	0.685	54.5
Sine	ENF$_{UnimPA}$	FIXED_AND	0.980	0.970	47.4

F. Interpretability Analysis

To illustrate the interpretability of the proposed UnimNeuron, we report a subset of rules learned on the *Elec2* dataset. Each rule is described by: (i) its support, (ii) the adaptive aggregation semantics expressed as the relative frequency of AND-, OR-, and COMP-like behavior, and (iii) a linguistic IF–THEN form over the original input variables.

Feature relevance weights w_j indicate the relative contribution of each dimension to the rule activation, while the adaptive connective reflects how the rule balances conjunctive, disjunctive, or compensatory reasoning over time.

1) **Rule 1.** *support*=747, σ=0.364. *semantics:* $\bar{\rho} = 0.818$, AND=0.0%, OR=47.2%, COMP=52.8%.
IF day is $Medium$ $(w = 1.00)$ $\wedge$ $nswprice$ is Low $(w = 0.22)$ $\wedge$ $date$ is Low $(w = 0.20)$ $\wedge$ $period$ is $Medium$ $(w = 0.20)$ $\wedge$ $nswdemand$ is $Medium$ $(w = 0.20)$ $\wedge$

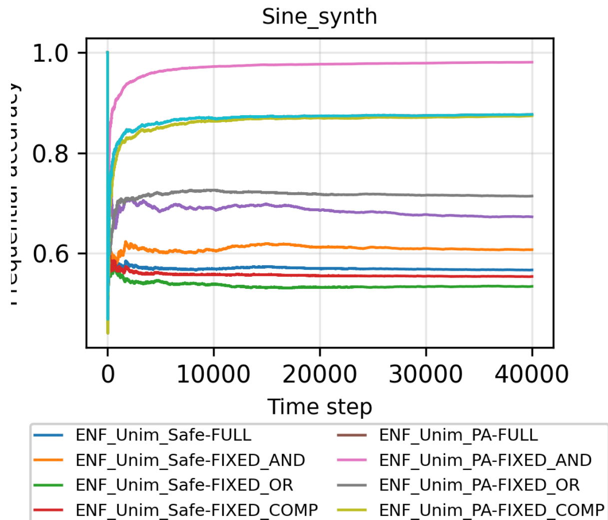

Fig. 3: Sine dataset: accuracy trendlines for the ablation experiment.

$vicprice$ is Low $(w = 0.20)$ **THEN** THEN class $= 0\,(p = 0.554)$.

2) **Rule 5.** *support*=38, σ=0.352. *semantics:* $\bar{\rho} = 0.687$, AND=0.0%, OR=19.4%, COMP=80.6%.
IF $date$ is Low $(w = 0.12) \wedge$ day is Low $(w = 0.12) \wedge$ $period$ is Low $(w = 0.12) \wedge$ $nswprice$ is Low $(w = 0.12) \wedge$ $nswdemand$ is Low $(w = 0.12) \wedge$ $vicprice$ is Low $(w = 0.12)$ **THEN** THEN class $= 0\,(p = 0.549)$.

These rules illustrate how the UnimNeuron captures structured and interpretable relationships by jointly adapting feature relevance and logical aggregation online. Variables such as *nswprice*, *nswdemand*, and *vicprice* receive distinct relevance weights w_j, indicating their differentiated contribution to rule activation, while less informative features are smoothly attenuated rather than eliminated, preserving semantic completeness. The emergent logical behavior of each rule—ranging from OR-like to compensatory (COMP)—reflects how evidences are integrated relative to the neutral element e, with permissive aggregation arising from a few dominant features and stricter conjunctive reasoning from more uniformly distributed relevance weights. Since both the connective and the feature weights are learned online, the resulting rules remain transparent and human-readable, explicitly conveying which variables matter and how they are logically combined, without sacrificing adaptability.

G. Summary of Results

The experimental results show that the proposed ENF-Unim models achieve competitive predictive performance while preserving strong structural and semantic interpretability. Across diverse real and synthetic data streams, including drift-intensive scenarios, the models attain accuracy statistically comparable to state-of-the-art TS- and PL-based approaches, maintain a compact and stable rule base (never exceeding five rules), and exhibit zero numerical instabilities or fallback predictions. They also demonstrate robust adaptation under severe concept drift and provide interpretable rule antecedents with automatically selected AND/OR/COMP semantics. The ablation study confirms that adaptive connective selection is critical: while fixed aggregation regimes may perform well in isolated cases, no static operator is consistently effective across datasets. Feature relevance weighting further enhances stability and compactness. Overall, these results validate the UnimNeuron as a principled building block for evolving neuro-fuzzy systems, offering a favorable trade-off between accuracy, interpretability, and structural simplicity in streaming environments.

VII. Conclusion

This paper introduced the UnimNeuron, a novel fuzzy neuron that automatically adapts its logical connective between AND, OR, and COMP regimes in an online and data-driven manner.

Building upon this concept, two evolving neuro-fuzzy models were proposed and evaluated under a rigorous prequential protocol on real and synthetic data streams with concept drift. The results demonstrate that the proposed models achieve competitive predictive performance while maintaining an exceptionally compact rule base, zero numerical instabilities, and fully interpretable rule structures.

Unlike traditional evolving fuzzy systems, where aggregation semantics are fixed or manually tuned, the UnimNeuron allows logical behavior to emerge directly from the data. Combined with feature-wise relevance weighting, this yields transparent rules that explain both which variables matter and how they are logically combined.

Future work will explore tighter integration of human-in-the-loop mechanisms, extensions to regression and forecasting tasks, and the incorporation of uncertainty-aware decision strategies in safety-critical streaming applications.

References

[1] E. P. Klement, *Triangular norms*. Springer Science & Business Media, 2013, vol. 8.

[2] R. R. Yager *et al.*, "Uninorms and their role in fuzzy logic," *Fuzzy Sets and Systems*, 1996, complete metadata to be added.

[3] T. Calvo *et al.*, "On a functional equation of nullnorms," *Fuzzy Sets and Systems*, 2001, complete metadata to be added.

[4] P. Akella, "Structure of n-uninorms," *Fuzzy Sets and Systems*, vol. 158, no. 15, pp. 1631–1651, 2007.

[5] U. Gürdal, Ömer Oğur, and S. Çetin, "Unima on [0,1]," *Fuzzy Sets and Systems*, vol. 516, p. 109439, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0165011425001782

[6] W. Pedrycz and F. Gomide, *Fuzzy systems engineering: toward human-centric computing*. John Wiley & Sons, 2007.

[7] K. Hirota and W. Pedrycz, "Or/and neuron in modeling fuzzy set connectives," *IEEE Transactions on Fuzzy Systems*, vol. 2, no. 2, pp. 151–161, 2002.

[8] A. Lemos, W. Caminhas, and F. Gomide, "New uninorm-based neuron model and fuzzy neural networks," in *2010 Annual Meeting of the North American Fuzzy Information Processing Society*, 2010, pp. 1–6.

[9] P. V. d. C. Souza, "Source code for the experimental evaluation of the unimneuron-based evolving neuro-fuzzy models," https://github.com/pdecampossouza/UnimNeuron-EvolvingNF, 2025, gitHub repository. Accessed: Month Year.

TABLE I: Prequential accuracy, rule growth and stability indicators on real and synthetic data streams.

Dataset	Model	N	Acc$_{final}$	Acc$_{mean}$	N^{final}_{rules}	Time [s]	Fallbacks	Instabilities
Elec2	ePL	10000	0.701	0.660	1	6.0	0	0
Elec2	ePL_{plus}	10000	0.737	0.690	1	6.8	0	0
Elec2	exTS	10000	0.729	0.701	7	32.1	0	421
Elec2	$Simpl_{eTS}$	10000	0.802	0.723	1	6.6	0	0
Elec2	eMG	10000	0.571	0.549	1	8.6	0	0
Elec2	$ENF_{UnimSafe}$	10000	0.591	0.603	5	13.2	0	0
Elec2	ENF_{UnimPA}	10000	0.797	0.803	5	17.2	0	0
CreditCard	ePL	5000	0.959	0.831	6	6.4	0	0
CreditCard	ePL_{plus}	5000	0.959	0.831	1	4.3	0	0
CreditCard	exTS	5000	0.959	0.831	3	62.4	0	160
CreditCard	$Simpl_{eTS}$	5000	0.959	0.831	1	4.1	0	0
CreditCard	eMG	5000	0.959	0.831	4	7.9	0	0
CreditCard	$ENF_{UnimSafe}$	5000	0.999	0.999	5	6.6	0	0
CreditCard	ENF_{UnimPA}	5000	0.999	0.999	5	8.5	0	0
Phishing	ePL	1250	0.741	0.476	44	8.2	0	0
Phishing	ePL_{plus}	1250	0.752	0.485	1	0.8	0	0
Phishing	exTS	1250	0.754	0.483	56	10.1	0	0
Phishing	$Simpl_{eTS}$	1250	0.751	0.485	1	0.8	0	0
Phishing	eMG	1250	0.725	0.464	116	44.9	0	0
Phishing	$ENF_{UnimSafe}$	1250	0.549	0.553	5	1.6	0	0
Phishing	ENF_{UnimPA}	1250	0.530	0.554	5	2.0	0	0
Bananas	ePL	5000	0.505	0.436	1	2.3	0	0
Bananas	ePL_{plus}	5000	0.512	0.440	1	2.8	0	0
Bananas	exTS	5000	0.512	0.450	3	19.8	0	153
Bananas	$Simpl_{eTS}$	5000	0.545	0.465	1	2.9	0	0
Bananas	eMG	5000	0.503	0.441	1	4.7	0	0
Bananas	$ENF_{UnimSafe}$	5000	0.541	0.544	5	6.5	0	0
Bananas	ENF_{UnimPA}	5000	0.526	0.535	5	8.6	0	0
ConceptDrift	ePL	20000	0.657	0.629	56	592.2	0	0
ConceptDrift	ePL_{plus}	20000	0.637	0.618	1	17.3	0	0
ConceptDrift	exTS	20000	0.576	0.570	5	80.1	0	897
ConceptDrift	$Simpl_{eTS}$	20000	0.665	0.638	1	13.4	0	0
ConceptDrift	eMG	20000	0.580	0.553	34	353.9	0	0
ConceptDrift	$ENF_{UnimSafe}$	20000	0.672	0.673	5	33.5	0	0
ConceptDrift	ENF_{UnimPA}	20000	0.607	0.609	5	39.0	0	0
Agrawal	ePL	20000	0.657	0.621	62	544.4	0	0
Agrawal	ePL_{plus}	20000	0.640	0.608	1	17.3	0	0
Agrawal	exTS	20000	0.599	0.582	4	92.5	0	893
Agrawal	$Simpl_{eTS}$	20000	0.666	0.633	1	14.6	0	0
Agrawal	eMG	20000	0.580	0.550	36	344.0	0	0
Agrawal	$ENF_{UnimSafe}$	20000	0.673	0.669	5	24.2	0	0
Agrawal	ENF_{UnimPA}	20000	0.610	0.607	5	34.6	0	0
RandomRBFDrift	ePL	20000	0.622	0.595	1	16.4	0	0
RandomRBFDrift	ePL_{plus}	20000	0.620	0.593	1	14.1	0	0
RandomRBFDrift	exTS	20000	0.632	0.619	4	87.4	0	907
RandomRBFDrift	$Simpl_{eTS}$	20000	0.756	0.718	1	14.2	0	0
RandomRBFDrift	eMG	20000	0.599	0.566	1	20.2	0	0
RandomRBFDrift	$ENF_{UnimSafe}$	20000	0.502	0.501	5	25.3	0	0
RandomRBFDrift	ENF_{UnimPA}	20000	0.502	0.503	5	31.6	0	0
Sine	ePL	40000	0.919	0.895	2	20.5	0	0
Sine	ePL_{plus}	40000	0.919	0.895	2	24.1	0	0
Sine	exTS	40000	0.909	0.888	4	85.2	0	1893
Sine	$Simpl_{eTS}$	40000	0.969	0.945	1	23.2	0	0
Sine	eMG	40000	0.915	0.891	4	71.5	0	0
Sine	$ENF_{UnimSafe}$	40000	0.560	0.563	5	50.4	0	0
Sine	ENF_{UnimPA}	40000	0.868	0.861	5	64.6	0	0