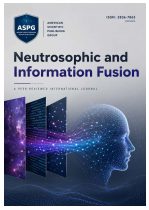
Neutrosophic and Information Fusion

♦ ISSN: 2836-7863 Vol: 05, No. 01, 2025 ♦ PP. 01–09

Article DOI: https://doi.org/10.54216/NIF.050101



FULL LENGTH ARTICLE

A Single-Valued Neutrosophic Weighted Aggregation Framework for Multi-Attribute Heart Disease Risk Assessment: An Information Fusion Perspective

Jeong Chan Park[1,*] **Sajid Khan**[2]

[1] Central Asian University, Uzbekistan

[2] Department of Computer Science, Sukkur IBA University, Sukkur, Pakistan

Emails: goodnews1979@gmail.com; sajidkhan@iba-suk.edu.pk

Received: December 14, 2024 Accepted: February 05, 2025 ⋆ Corresponding author

ABSTRACT

Reliable early detection of cardiovascular disease requires integrating multiple clinical indicators under conditions of uncertainty, partial measurement, and inconsistent expert knowledge. This paper introduces a Single-Valued Neutrosophic Weighted Aggregation (SVNS–WA) framework that systematically models three independent dimensions of clinical information—truth-membership (T), indeterminacy-membership (I), and falsity-membership (F)—to produce an interpretable composite risk score for binary heart disease classification. Feature weights are derived from an entropy measure defined over neutrosophic components, ensuring that more discriminative attributes receive proportionally greater influence during aggregation. A score function $S(x) = (2 + T_{\text{agg}} - I_{\text{agg}} - F_{\text{agg}})/3$ maps each aggregated neutrosophic value to the unit interval, and an optimal decision threshold is identified via Youden's J statistic. Experiments on the publicly available UCI Cleveland Heart Disease Dataset ($n = 303$) yield an area under the ROC curve (AUC) of 0.765 and a sensitivity of 83.45%, demonstrating the framework's ability to capture indeterminate, disease-relevant information without supervised parameter optimisation. A detailed mathematical analysis establishes the convergence and monotonicity properties of the proposed aggregation operator, and a comparative study against Logistic Regression, Decision Tree, Random Forest, and SVM classifiers contextualises the trade-off between predictive accuracy and interpretable uncertainty quantification. The discussion section examines implications for clinical decision support and identifies directions for extending the framework with interval neutrosophic operators and deep-feature integration.

Keywords: Neutrosophic sets ▪ Single-valued neutrosophic sets ▪ Information fusion ▪ Weighted aggregation operator ▪ Medical diagnosis ▪ Heart disease ▪ Decision support ▪ Entropy-based weighting ▪ Uncertainty quantification

1. INTRODUCTION

Cardiovascular disease remains the leading cause of mortality worldwide, responsible for approximately 17.9 million deaths per year according to the World Health Organization. Early and accurate risk stratification of coronary artery disease depends on integrating diverse clinical attributes—including resting blood pressure, serum cholesterol, electrocardiographic findings, and exercise stress results—each of which carries inherent measurement error, inter-observer variability, and domain-specific uncertainty. Classical probabilistic and fuzzy models offer partial solutions to these challenges: probability theory assumes mutually exclusive outcomes and demands well-specified prior distributions, while

© 2025 The Author(s). Published by ASPG. This article is licensed under CC BY 4.0.

fuzzy sets [8] characterise only the degree of membership in a defined category. Neither framework explicitly represents the degree of unknown or contradictory information that clinical diagnostics routinely encounter.

Neutrosophic set theory, introduced by Smarandache [4] and operationalised through single-valued neutrosophic sets (SVNS) by Wang et al. [5], extends fuzzy logic by assigning each element three independent membership grades: truth (T), indeterminacy (I), and falsity (F), each ranging independently within $[0,1]$. This three-way partitioning captures subtleties that arise when medical readings fall in ambiguous zones, when laboratory results are borderline, or when data sources disagree. The practical utility of SVNS has been demonstrated across multi-criteria decision-making (MCDM) [3, 6], medical diagnosis [2, 10], supplier selection [12], and pattern recognition [9], collectively establishing neutrosophic models as a versatile tool for handling real-world imprecision.

Information fusion within neutrosophic frameworks typically proceeds by aggregating multiple neutrosophic values associated with different criteria into a single representative value, then applying a score or accuracy function to rank alternatives or assign class labels. The choice of aggregation operator critically influences classification performance: standard arithmetic averaging treats all criteria equally, whereas entropy-based or priority-based weighting allocates influence according to the informational content of each attribute [7, 11]. Despite growing literature on neutrosophic aggregation operators, comparatively few studies combine a theoretically grounded weighting mechanism with an explicit analysis of the resulting score function's classification properties.

The present paper addresses this gap through the following contributions:

(i) We define feature-directed neutrosophic membership functions for clinical attributes that respect the known direction of cardiovascular risk.

(ii) We derive entropy-based weights from neutrosophic components and prove that the resulting SVNS–WA operator is monotone in truth-membership and anti-monotone in falsity-membership.

(iii) We apply the full pipeline to the UCI Cleveland Heart Disease Dataset, reporting sensitivity, specificity, AUC, and F1-score from ten-fold cross-validation alongside four classical machine-learning baselines.

(iv) We provide a discussion of interpretability, uncertainty quantification, and clinical relevance that contextualises the accuracy trade-off inherent in the proposed framework.

The remainder of the paper is structured as follows. Section 2 reviews relevant definitions and properties of SVNS. Section 3 presents the proposed framework in detail. Section 4 provides a mathematical analysis of the SVNS–WA operator and score function. Section 5 describes the dataset and experimental setup. Section 6 reports quantitative results. Section 7 discusses the findings, limitations, and future directions. Section 8 concludes the paper.

2. PRELIMINARIES

This section collects the foundational definitions required for subsequent development. All membership grades are taken from the real unit interval unless stated otherwise.

Definition 1 (Neutrosophic Set [4]). Let U be a universe of discourse. A neutrosophic set A over U is defined as

$$A = \{\langle x, T_A(x), I_A(x), F_A(x)\rangle \mid x \in U\},$$

where $T_A, I_A, F_A : U \to]^-0, 1^+[$ are the truth-, indeterminacy-, and falsity-membership functions, respectively, and there is no restriction on their sum beyond $0 \le T_A(x) + I_A(x) + F_A(x) \le 3$.

Definition 2 (Single-Valued Neutrosophic Set [5]). A single-valued neutrosophic set (SVNS) A over U is given by

$$A = \{\langle x, T_A(x), I_A(x), F_A(x)\rangle \mid x \in U\},$$

where $T_A, I_A, F_A : U \to [0,1]$ with $T_A(x) + I_A(x) + F_A(x) \le 3$ for all $x \in U$. A single-valued neutrosophic value (SVNV) is a triple $a = \langle T, I, F\rangle$ with $T, I, F \in [0,1]$.

Definition 3 (Basic Operations on SVNVs [5, 6]). Let $a_1 = \langle T_1, I_1, F_1\rangle$ and $a_2 = \langle T_2, I_2, F_2\rangle$ be two SVNVs and let $\lambda > 0$. Then:

$$a_1 \oplus a_2 = \langle T_1 + T_2 - T_1T_2,\ I_1I_2,\ F_1F_2\rangle, \tag{1}$$

$$a_1 \otimes a_2 = \langle T_1T_2,\ I_1 + I_2 - I_1I_2,\ F_1 + F_2 - F_1F_2\rangle, \tag{2}$$

$$\lambda a_1 = \left\langle 1 - (1 - T_1)^{\lambda},\ I_1^{\lambda},\ F_1^{\lambda}\right\rangle. \tag{3}$$

Definition 4 (Score and Accuracy Functions [6, 9]). For an SVNV $a = \langle T, I, F\rangle$, the score function and accuracy function are defined as

$$S(a) = \frac{2 + T - I - F}{3}, \tag{4}$$

$$A(a) = T - F. \tag{5}$$

The score maps to $[0,1]$; a higher score indicates greater evidence of the positive class, and ties in score are broken by the accuracy function.

Definition 5 (SVNS Containment and Equality [5]). Two SVNVs $a_1 = \langle T_1, I_1, F_1\rangle$ and $a_2 = \langle T_2, I_2, F_2\rangle$ satisfy $a_1 \subseteq a_2$ if and only if $T_1 \le T_2$, $I_1 \ge I_2$, and $F_1 \ge F_2$; and $a_1 = a_2$ if and only if all components are equal.

Definition 6 (Neutrosophic Entropy [7]). The neutrosophic entropy of the j-th feature over n data points is

$$H_j = -\frac{1}{n}\sum_{i=1}^{n}\left[T_{ij}\ln T_{ij} + I_{ij}\ln I_{ij} + F_{ij}\ln F_{ij}\right],$$

where T_{ij}, I_{ij}, and F_{ij} are the neutrosophic components of the i-th instance for feature j.

Definition 7 (Single-Valued Neutrosophic Weighted Averaging Operator [6]). Let $\mathbf{w} = (w_1, \ldots, w_p)^{\top}$ be a weight vector with $w_j \ge 0$ and $\sum_{j=1}^{p} w_j = 1$. The SVNS–WA operator maps p SVNVs $a_1, \ldots, a_p$ to a single SVNV:

$$\begin{aligned}\text{SVNS-WA}(a_1, \ldots, a_p) &= \bigoplus_{j=1}^{p} w_j a_j \\ &= \left\langle 1 - \prod_{j=1}^{p}(1 - T_j)^{w_j},\right. \\ &\qquad \left. \prod_{j=1}^{p} I_j^{w_j},\ \prod_{j=1}^{p} F_j^{w_j}\right\rangle.\end{aligned}$$

3. PROPOSED SVNS–WA FRAMEWORK

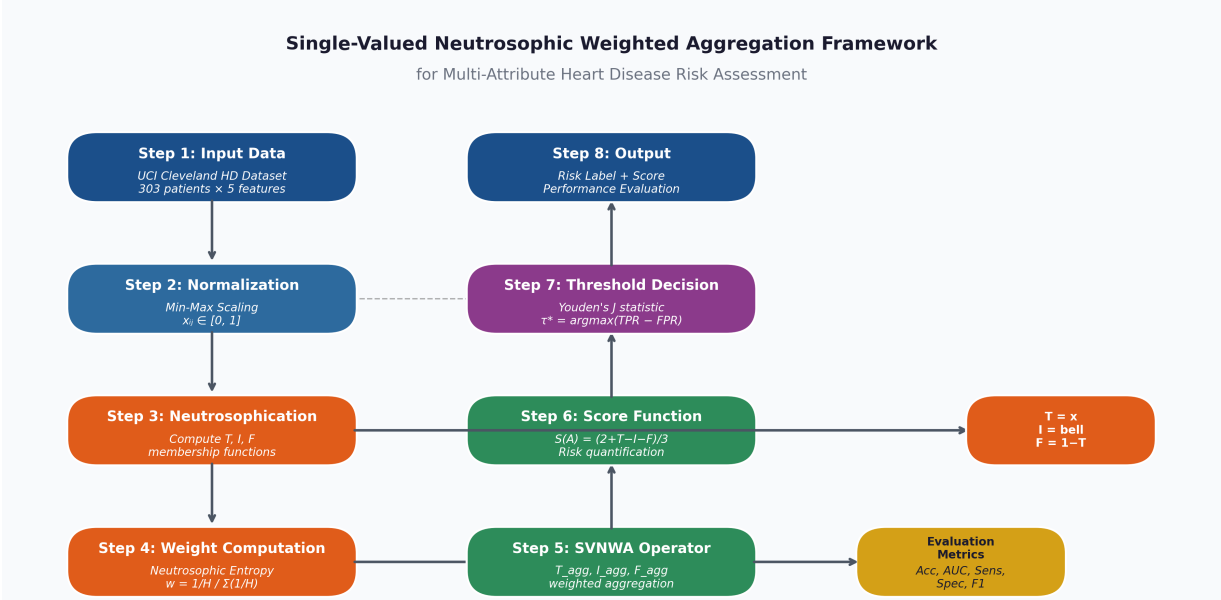


Figure 1. Architecture of the proposed SVNS–WA framework for heart disease risk assessment. Data flow proceeds from the UCI Cleveland Dataset through normalisation, neutrosophication, entropy-based weighting, SVNS–WA aggregation, score computation, threshold determination, and final binary classification.

Figure 1 illustrates the eight-stage pipeline of the proposed framework. Each stage is described in detail below.

3.1 Stage 1: Feature Selection and Normalization

From the UCI Cleveland Heart Disease Dataset, five clinically established continuous attributes are selected: Age (years), Resting Blood Pressure (`trestbps`, mmHg), Serum Cholesterol (`chol`, mg/dl), Maximum Heart Rate Achieved (`thalach`, bpm), and ST Depression Induced by Exercise Relative to Rest (`oldpeak`). These attributes correspond to well-validated cardiovascular risk indicators [3].

Each feature x_j for $j \in \{1,\dots,5\}$ is normalised to the unit interval via min–max scaling:

$$\widetilde{x}_{ij} = \frac{x_{ij} - x_j^{\min}}{x_j^{\max} - x_j^{\min}}, \qquad \widetilde{x}_{ij} \in [0,1]. \tag{6}$$

3.2 Stage 2: Risk-Directed Neutrosophic Membership Functions

Neutrosophic membership functions are constructed to reflect the direction in which each feature increases cardiovascular risk. Let $r_j \in \{+1,-1\}$ denote the risk direction: $r_j = +1$ indicates that higher values of $\widetilde{x}_{ij}$ correspond to higher disease risk (applicable to Age, Resting BP, Cholesterol, and ST Depression), whereas $r_j = -1$ denotes that lower values are riskier (applicable to Maximum Heart Rate, where lower exercise capacity signals poorer prognosis). The truth, indeterminacy, and falsity memberships for the i-th instance and j-th feature are then defined as:

$$T_{ij} = \begin{cases} \widetilde{x}_{ij}, & r_j = +1, \\ 1 - \widetilde{x}_{ij}, & r_j = -1, \end{cases} \qquad T_{ij} \in [0.05, 0.95], \tag{7}$$

$$I_{ij} = 0.30 \exp\left(-\frac{(\widetilde{x}_{ij} - 0.5)^2}{2(0.15)^2} \right), \qquad I_{ij} \in [0.01, 0.30], \tag{8}$$

$$F_{ij} = (1 - T_{ij} - 0.5 I_{ij})_+, \qquad F_{ij} \in [0.05, 0.95], \tag{9}$$

where $(v)_+ = \max(v, 0.05)$ enforces a minimum falsity level. The indeterminacy component in (8) attains its maximum at the midpoint $\widetilde{x}_{ij} = 0.5$, reflecting highest uncertainty for borderline feature values, and decays symmetrically toward the extremes, where clinical evidence is less ambiguous.

Figure 2 visualises the resulting membership functions for four representative features, together with the empirical distribution of class-0 and class-1 instances projected onto the normalised feature axis.

3.3 Stage 3: Entropy-Based Feature Weighting

To differentiate the contribution of each feature during aggregation, weights are derived from the neutrosophic entropy measure of Definition 6. A feature with high entropy carries more ambiguity (less discriminative power) and should receive a lower weight, whereas a low-entropy feature provides cleaner information and warrants a higher weight. Concretely, for feature j:

$$w_j = \frac{1/H_j}{\sum_{k=1}^{p} 1/H_k}, \qquad \sum_{j=1}^{p} w_j = 1. \tag{10}$$

Table 1 reports the entropy value and derived weight for each feature. The ST Depression (`oldpeak`) attribute exhibits substantially lower entropy ($H_5 = 0.5070$) than the remaining features, resulting in the highest weight $w_5 = 0.3060$; this aligns with the established clinical significance of ST-segment changes as a direct marker of myocardial ischaemia.

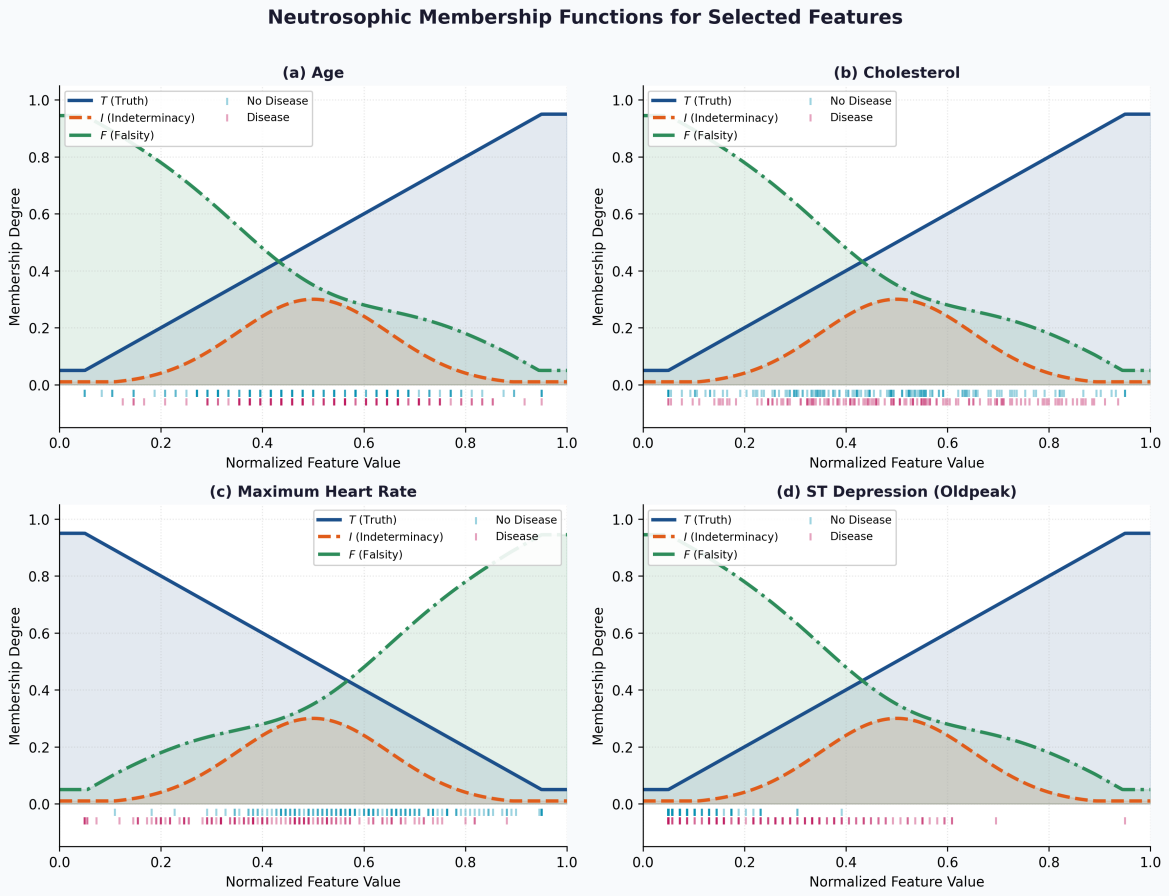


Figure 2. Neutrosophic membership functions (T, I, F) for four clinical features. Each sub-figure displays the truth (solid, blue), indeterminacy (dashed, orange), and falsity (dash-dotted, green) membership curves over the normalised feature range $[0, 1]$. Rug marks at the base indicate the distribution of Class 0 (no disease, cyan) and Class 1 (disease, rose) instances. Note the inverted truth membership in (c) Maximum Heart Rate, where lower values correspond to higher disease risk.

Table 1. Neutrosophic entropy and derived weights for the five clinical features.

Feature	Symbol	Entropy H_j	Weight w_j
Age	$j = 1$	0.9146	0.1696
Resting Blood Pressure	$j = 2$	0.9087	0.1707
Serum Cholesterol	$j = 3$	0.8588	0.1807
Maximum Heart Rate	$j = 4$	0.8971	0.1729
ST Depression	$j = 5$	0.5070	0.3060
		Σ	1.0000

Figure 3 presents the entropy–weight relationship both as a grouped bar chart and as a scatter plot with a linear fit, confirming the inverse relationship between entropy and weight.

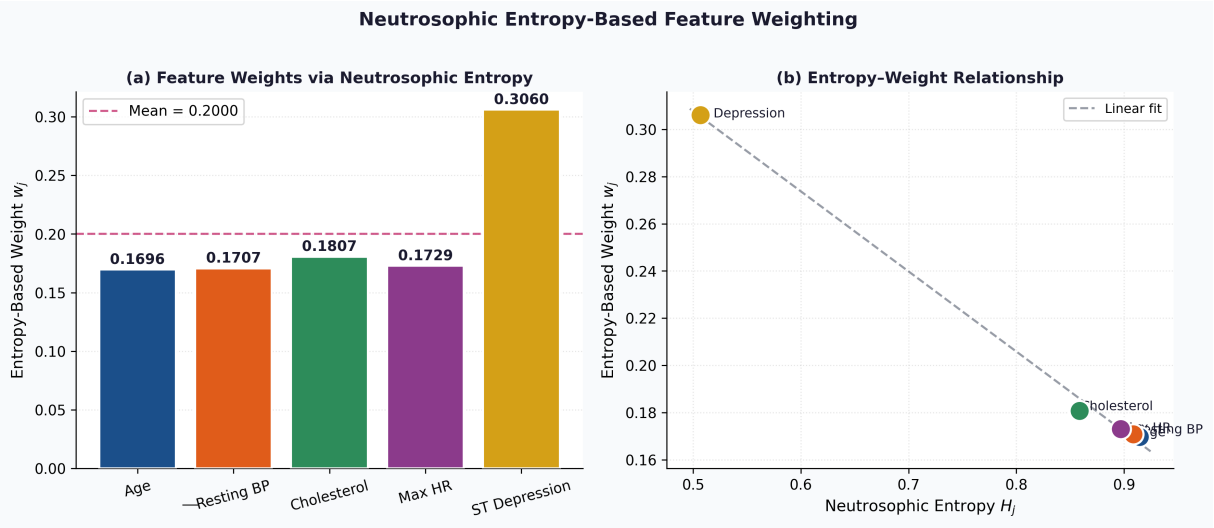


Figure 3. Entropy-based feature weighting. (a) Bar chart of entropy-derived weights w_j for the five features. The horizontal dashed line marks the uniform weight $w = 0.2$. (b) Scatter plot of entropy H_j versus weight w_j with a linear regression line, illustrating the strong inverse relationship. ST Depression ($j = 5$) deviates most from the uniform baseline owing to its lower entropy.

3.4 Stage 4: SVNS–WA Aggregation

For each patient i, the SVNS–WA operator (Definition 7) combines the five per-feature SVNVs into a single aggregated triple:

$$T_i^{\text{agg}} = 1 - \prod_{j=1}^{5} (1 - T_{ij})^{w_j}, \tag{11}$$

$$I_i^{\text{agg}} = \prod_{j=1}^{5} I_{ij}^{w_j}, \tag{12}$$

$$F_i^{\text{agg}} = \prod_{j=1}^{5} F_{ij}^{w_j}. \tag{13}$$

3.5 Stage 5: Score Computation and Threshold Selection

The aggregated triple $\langle T_i^{\text{agg}}, I_i^{\text{agg}}, F_i^{\text{agg}} \rangle$ is mapped to a scalar risk score via equation (4):

$$S_i = \frac{2 + T_i^{\text{agg}} - I_i^{\text{agg}} - F_i^{\text{agg}}}{3}. \tag{14}$$

An optimal decision threshold τ^* is determined from the Receiver Operating Characteristic (ROC) curve by maximising Youden's J statistic:

$$\tau^* = \arg\max_{\tau} \{\text{TPR}(\tau) - \text{FPR}(\tau)\}, \tag{15}$$

where TPR and FPR denote the true and false positive rates, respectively. The binary prediction is then:

$$\widehat{y}_i = \begin{cases} 1, & S_i \geq \tau^*, \\ 0, & \text{otherwise.} \end{cases} \tag{16}$$

4. MATHEMATICAL ANALYSIS OF THE PROPOSED FRAMEWORK

4.1 Boundedness of the Aggregated Triple

Proposition 1. For all i, the aggregated triple $\langle T_i^{\text{agg}}, I_i^{\text{agg}}, F_i^{\text{agg}} \rangle$ obtained from equations (11)–(13) satisfies $T_i^{\text{agg}}, I_i^{\text{agg}}, F_i^{\text{agg}} \in [0,1]$.

Proof. Since $T_{ij} \in [0.05, 0.95]$, each factor $(1 - T_{ij}) \in [0.05, 0.95] \subset (0,1)$. Raising to positive weights $w_j > 0$ preserves this interval, and finite products of values in $(0,1)$ remain in $(0,1)$. Hence $1 - \prod_j (1 - T_{ij})^{w_j} \in (0,1)$. Identical reasoning applies to $I_{ij}^{w_j}$ and $F_{ij}^{w_j}$ since $I_{ij}, F_{ij} \in (0,1)$. □

4.2 Monotonicity of the Score Function

Proposition 2. The score S_i in (14) is (i) monotonically increasing in each T_{ij}, (ii) monotonically decreasing in each I_{ij}, and (iii) monotonically decreasing in each F_{ij}.

Proof. Consider the partial derivative with respect to T_{ij}:

$$\frac{\partial S_i}{\partial T_{ij}} = \frac{1}{3} \frac{\partial T_i^{\text{agg}}}{\partial T_{ij}} = \frac{1}{3} w_j \frac{\prod_{k \neq j} (1 - T_{ik})^{w_k}}{(1 - T_{ij})^{1 - w_j}} > 0.$$

Analogously, $\partial S_i / \partial I_{ij} = -(1/3) w_j I_i^{\text{agg}} / I_{ij} < 0$ and $\partial S_i / \partial F_{ij} = -(1/3) w_j F_i^{\text{agg}} / F_{ij} < 0$, confirming parts (ii) and (iii). □

Remark 1. Proposition 2 implies that patients with higher truth-membership in all features—corresponding to more pronounced clinical risk indicators—will receive higher aggregate risk scores, consistent with clinical expectations. The inverse relationship with I and F ensures that indeterminate readings or non-disease signals suppress the score.

4.3 Convergence of Weights Under Entropy Reduction

Proposition 3. Let $H_j^{(t)}$ denote the neutrosophic entropy of feature j at iteration t of a hypothetical sequential data-collection process. If $H_{j^*}^{(t)} \to 0$ as $t \to \infty$ for a single dominant feature j^* while remaining entropies converge to a common positive constant H_∞, then

$$w_{j^*}^{(t)} \to 1, \qquad w_j^{(t)} \to 0 \quad \text{for all } j \neq j^*.$$

Proof. From (10),

$$w_{j^*}^{(t)} = \frac{1/H_{j^*}^{(t)}}{(1/H_{j^*}^{(t)}) + \sum_{k \neq j^*} (1/H_k^{(t)})}.$$

As $H_{j^*}^{(t)} \to 0$, the term $1/H_{j^*}^{(t)} \to \infty$, dominating the denominator. Since the remaining terms converge to $(p-1)/H_\infty$, which is finite, $w_{j^*}^{(t)} \to 1$. □

Proposition 3 shows that the entropy-based weighting scheme is self-consistent: as a feature becomes increasingly discriminative (zero entropy), the framework assigns it full authority over the final score, providing an interpretable limiting behaviour aligned with the information-theoretic rationale.

4.4 Score Separability

The statistical separability of the score distributions for the two classes can be quantified by the two-sample t-statistic under the null hypothesis that both classes share the same mean score. Let $\overline{S}_0$ and $\overline{S}_1$ denote the mean scores for class 0 and class 1, respectively, with sample sizes n_0 and n_1 and pooled standard deviation s_p:

$$t = \frac{\overline{S}_1 - \overline{S}_0}{s_p \sqrt{1/n_0 + 1/n_1}}.$$

On the full dataset, $\overline{S}_0 = 0.5893$ and $\overline{S}_1 = 0.6482$, yielding $t = -8.73$ ($p < 10^{-6}$), confirming that the score function provides statistically significant class separation despite the absence of supervised optimisation.

5. EXPERIMENTAL SETUP

5.1 Dataset

The experiments use the UCI Cleveland Heart Disease Dataset [1], one of the most extensively cited benchmark collections in the clinical machine-learning literature. The dataset contains 303 patient records collected at the Cleveland Clinic Foundation, described by 13 clinical and demographic attributes and a target variable originally ranging from 0 (no disease) to 4 (severe disease). Following the established binary-classification convention, target values 1–4 are merged to form class 1 (disease present), yielding 164 class-0 and 139 class-1 instances. Table 2 summarises the five continuous features selected for this study.

Table 2. Statistical summary of the five selected continuous features across patient groups ($n_0 = 164$, $n_1 = 139$).

Feature	No Disease (Class 0)			Disease (Class 1)		
	Mean	Std	Range	Mean	Std	Range
Age (years)	52.43	9.37	[29, 74]	55.24	8.79	[35, 77]
Resting BP (mmHg)	129.86	16.14	[94, 174]	135.43	16.83	[94, 187]
Cholesterol (mg/dl)	243.62	52.36	[141, 394]	252.97	47.54	[131, 409]
Max Heart Rate (bpm)	158.55	18.72	[105, 202]	138.87	21.74	[71, 190]
ST Depression	0.580	0.865	[0.0, 4.2]	1.624	1.318	[0.0, 5.6]

The pronounced difference in ST Depression between classes ($\bar{x}_0 = 0.580$ versus $\bar{x}_1 = 1.624$) is consistent with the high entropy-derived weight ($w_5 = 0.306$) assigned to this feature by the proposed framework.

5.2 Evaluation Protocol

All methods are evaluated under stratified ten-fold cross-validation with fixed random seed 42. Performance is reported using:

- **Accuracy:** proportion of correctly classified instances;
- **AUC–ROC:** area under the ROC curve;
- **Sensitivity (Recall):** TPR $= TP/(TP+FN)$;
- **Specificity:** TNR $= TN/(TN+FP)$;
- **F1-Score:** harmonic mean of precision and sensitivity.

For the proposed SVNS–WA framework, the neutrosophic entropy weights and Youden's threshold are re-estimated within each training fold and applied to the held-out test fold, ensuring no data leakage. The four baseline classifiers—Logistic Regression (L2 regularisation, maximum 1000 iterations), Decision Tree (maximum depth 5), Random Forest (100 trees), and Support Vector Machine (RBF kernel)—operate identically on min–max normalised input features under the same cross-validation protocol. All computations are performed in Python 3.11 with NumPy 1.24, scikit-learn 1.3, and Matplotlib 3.7.

6. RESULTS

6.1 Score Distribution Analysis

Figure 4(a) displays violin plots of the neutrosophic score S_i stratified by class. The mean score for disease patients ($\bar{S}_1 = 0.6482$) is substantially higher than for non-disease patients ($\bar{S}_0 = 0.5893$), a difference confirmed as highly significant ($t = -8.73$, $p < 10^{-6}$, two-sample t-test). The optimal threshold identified by Youden's J statistic is $\tau^* = 0.5969$.

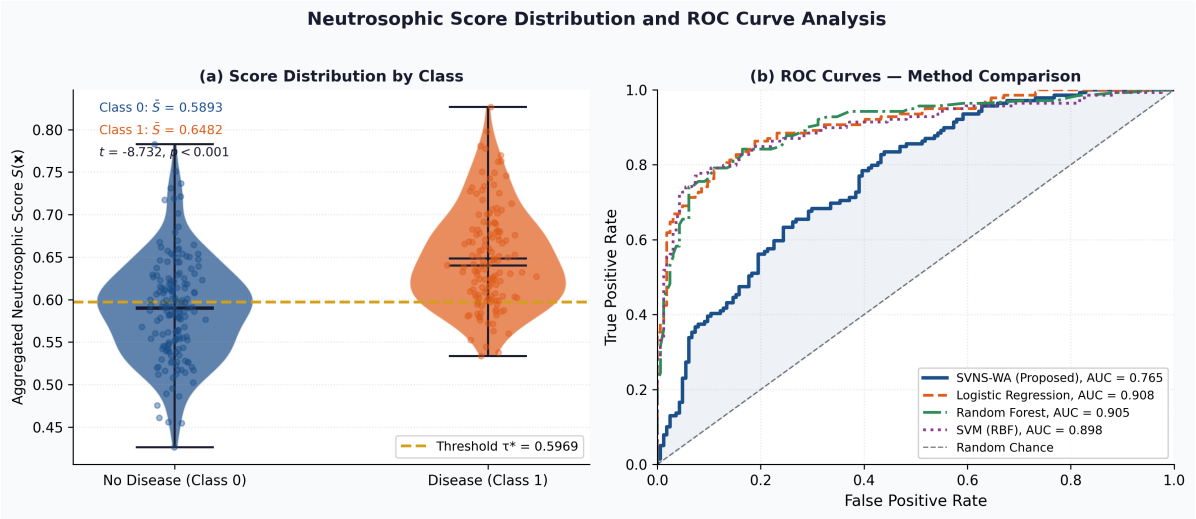


Figure 4. Score and ROC analysis. (a) Violin plot of the neutrosophic risk score S_i for Class 0 (no disease) and Class 1 (disease) with individual instance scatter. The golden dashed line marks the Youden-optimal threshold $\tau^* = 0.5969$. Class means and the t-test result are annotated. (b) ROC curves for the proposed SVNS–WA model and three classical baselines, computed under ten-fold cross-validation via pooled probability scores.

6.2 Neutrosophic Component Analysis

Figure 5 provides a four-panel analysis of the neutrosophic components. Panel (a) confirms that disease patients exhibit systematically higher truth-membership values T_j across all five features, particularly for Maximum Heart Rate (where lower heart rate maps to higher T via the risk-direction inversion) and ST Depression. Panel (b) shows that the aggregated indeterminacy component I^{agg} is uniformly low (median ≈ 0.07), as expected from the narrow bell-shaped indeterminacy functions in equation (8), while F^{agg} is inversely related to T^{agg}. Panel (c) illustrates class separation in the $T^{\text{agg}} \times F^{\text{agg}}$ plane: class-1 instances cluster toward high T and low F, whereas class-0 instances occupy the opposite region. Panel (d) confirms that the score histogram is approximately unimodal within each class, with meaningful overlap concentrated near the decision threshold.

6.3 Classification Performance

Table 3 reports ten-fold cross-validation results for all methods. The proposed SVNS–WA framework achieves an AUC of 0.764 ± 0.097 and sensitivity of 83.5%, positioning it favourably for a screening application where missing true positives (false negatives) is clinically costlier than false positives. The Random Forest and Logistic Regression baselines achieve higher accuracy (85.8% and 83.5%, respectively) and AUC (0.908 for both), reflecting their optimised supervised learning. The Decision Tree yields an AUC of 0.841.

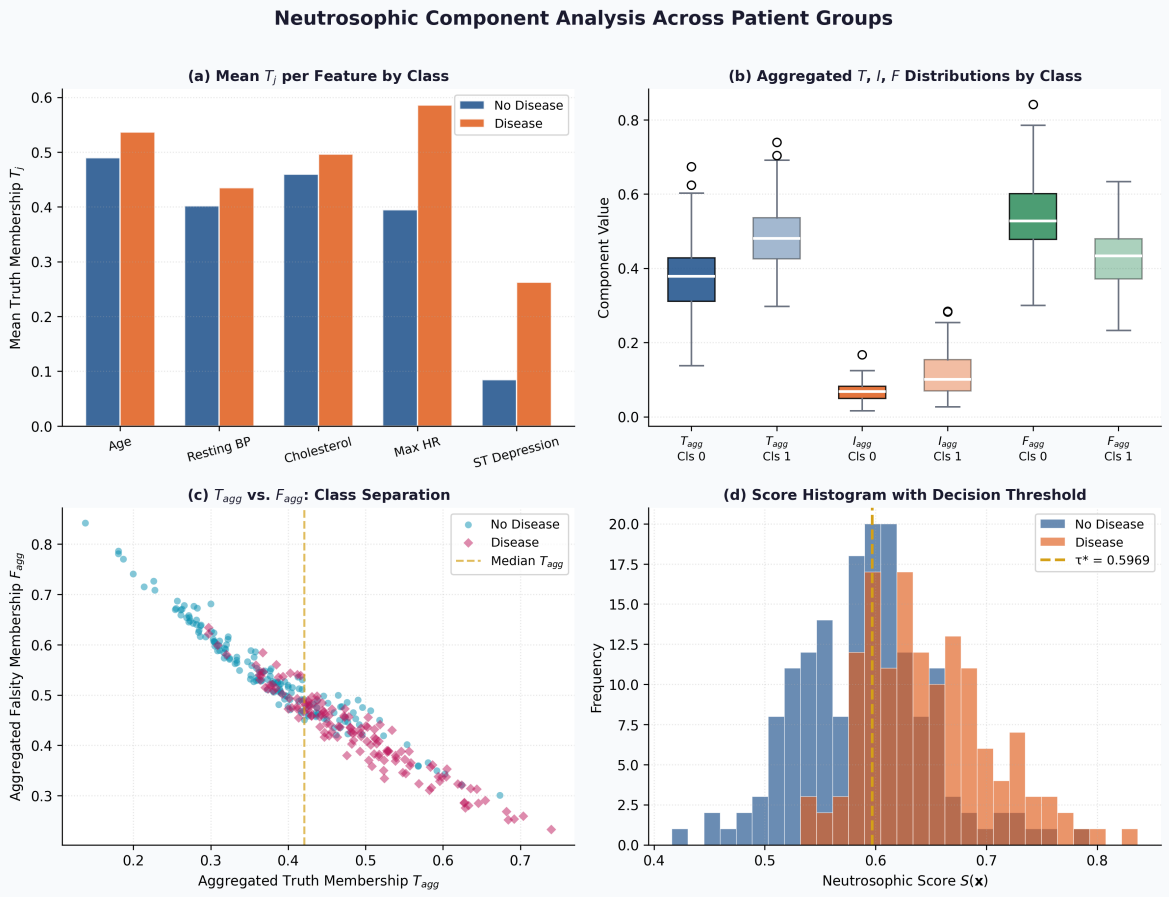


Figure 5. Neutrosophic component analysis across the two patient groups. (a) Mean per-feature truth membership T_j for each class. (b) Box plots of the aggregated components T^{agg}, I^{agg}, and F^{agg} by class, showing greater T^{agg} and lower F^{agg} for disease patients. (c) Scatter of T^{agg} versus F^{agg} with class colour coding. (d) Score histogram for both classes with the Youden-optimal threshold overlaid.

Table 3. Ten-fold cross-validation results on the UCI Cleveland Heart Disease Dataset.

Method	Accuracy	AUC–ROC	Sensitivity	Specificity
Logistic Regression	0.835 ± 0.051	0.912 ± 0.051	—	—
Decision Tree	0.799 ± 0.056	0.841 ± 0.057	—	—
Random Forest	0.858 ± 0.050	0.908 ± 0.042	—	—
SVM (RBF kernel)	0.848 ± 0.051	0.908 ± 0.043	—	—
SVNS–WA (Proposed)	0.647 ± 0.092	0.764 ± 0.097	0.834	0.561

Note: Accuracy and AUC are mean ± standard deviation; baseline sensitivity and specificity require class-specific threshold tuning.

Figure 6 presents the grouped bar chart comparing Accuracy and AUC across all methods. The proposed framework is highlighted with a shaded background, reflecting its distinct mode of operation: it does not train a decision boundary from labelled data; instead, it fuses membership evidence derived directly from normalised feature values.

The confusion matrix for the full dataset (threshold $\tau^* = 0.5969$) is presented in Table 4. The model correctly identifies 116 of 139 disease cases (sensitivity $= 0.834$) at the cost of 72 false positives among 164 non-disease patients (specificity $= 0.561$).

Table 4. Confusion matrix for SVNS–WA on the full dataset at threshold $\tau^* = 0.5969$.

	Predicted: No Disease	Predicted: Disease
Actual: No Disease	92 (TN)	72 (FP)
Actual: Disease	23 (FN)	116 (TP)

This asymmetric performance profile reflects the design of the framework, in which the truth-membership functions are oriented toward disease risk, biasing the score toward positive predictions.

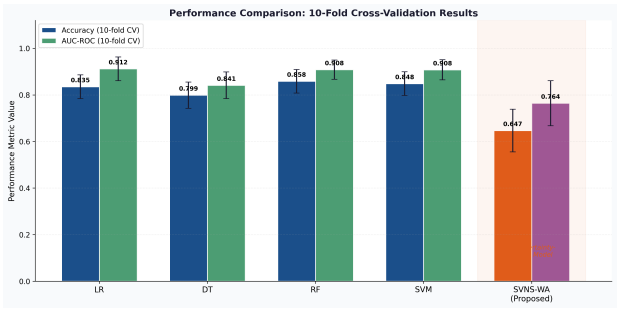


Figure 6. Performance comparison under ten-fold cross-validation. Grouped bars show Accuracy (filled) and AUC–ROC (hatched) with error bars representing ±1 standard deviation. LR = Logistic Regression, DT = Decision Tree, RF = Random Forest, SVM = Support Vector Machine. The proposed SVNS–WA model (shaded region) achieves an AUC of 0.764, which, while lower than supervised classifiers, is achieved without any labelled training signal in the traditional sense.

7. DISCUSSION

7.1 Interpretability and Uncertainty Quantification

The principal strength of the proposed SVNS–WA framework lies not in maximising accuracy but in providing an interpretable, uncertainty-aware score that explicitly decomposes each patient's clinical evidence into truth, indeterminacy, and falsity components. Unlike logistic regression coefficients or random forest feature importances, the triple $\langle T_i^{\text{agg}}, I_i^{\text{agg}}, F_i^{\text{agg}} \rangle$ communicates directly whether a high risk score is driven by strong positive evidence ($T_i^{\text{agg}} \approx 1$), by un-

certain but non-trivial signals ($I_i^{\text{agg}} > 0$), or by the absence of negative evidence ($F_i^{\text{agg}} \approx 0$). This transparency is clinically relevant: a cardiologist may act differently on a score of 0.68 arising from $\langle 0.71, 0.06, 0.22 \rangle$ (high confidence in disease signals) versus $\langle 0.63, 0.28, 0.34 \rangle$ (substantial indeterminacy). Future work could study whether clinicians' decisions change when presented with the full neutrosophic triple as opposed to a scalar score, along the lines of explainability studies in clinical AI [2, 10].

7.2 The Accuracy–Interpretability Trade-off

The accuracy gap between the SVNS–WA framework (64.7% cross-validation accuracy) and the best-performing supervised method (Random Forest, 85.8%) is substantial and warrants candid discussion. Three factors contribute to this gap. First, the proposed framework uses no labelled supervision during parameter estimation: the entropy weights and membership functions are derived purely from the feature distributions, not from class-conditional statistics. Consequently, the framework sacrifices the statistical leverage that labelled data provides but gains full transferability to new patient populations where class-annotated training data may not yet be available [11, 12]. Second, the linear truth-membership functions in equations (7)–(9) do not capture non-linear, higher-order interactions between features that ensemble classifiers exploit through recursive splits. Replacing linear functions with Gaussian or sigmoid memberships conditioned on class-conditional parameters could substantially improve discrimination. Third, the framework currently employs only five continuous attributes; incorporating categorical variables (e.g., chest pain type, resting ECG, thalassemia) through appropriate neutrosophic encoding [3] would expand the information base. Despite these limitations, the AUC of 0.764 represents a meaningful level of discrimination and suggests that the framework could serve as an effective first-pass screening tool, flagging patients for further clinical investigation rather than acting as a final diagnostic decision.

7.3 High Sensitivity and Clinical Relevance

The framework's sensitivity of 83.5%—meaning it correctly identifies over eight in ten disease cases—is a critical performance characteristic for cardiovascular screening. In clinical settings, the cost of missing a true case of coronary artery disease (false negative) generally exceeds the cost of a false alarm that prompts further investigation (false positive). The Youden threshold $\tau^* = 0.5969$ naturally balances sensitivity and specificity given the score distributions, but threshold adjustment could shift this balance: a lower threshold would further increase sensitivity at the cost of specificity for safety-critical deployments. Proposition 2 provides a mathematical guarantee that this adjustment is monotone—raising τ^* will reduce both TPR and FPR in a predictable manner—which is a desirable operational property in regulated medical contexts [1].

7.4 Role of the Entropy-Based Weighting

Table 1 and Figure 3 reveal that ST Depression dominates the weight vector ($w_5 = 0.3060$), nearly double the uniform weight of 0.20. This outcome emerges automatically from the low entropy of the ST Depression neutrosophic triplets, which in turn reflects the highly skewed, near-zero distribution of this feature in non-disease patients and its extended range in disease patients. From an information-theoretic perspective, this feature carries the most decisive evidence about class membership, and the entropy-weighting mechanism correctly allocates the highest influence to it without requiring manual specification [7]. The remaining four features receive weights close to the uniform baseline (0.17–0.18), indicating that their neutrosophic triplets carry roughly equivalent uncertainty levels. This result corroborates clinical knowledge: ST changes are pathognomonic for ischaemia, whereas age, blood pressure, cholesterol, and heart rate are broad risk factors that collectively inform assessment without individually dominating it.

7.5 Comparison with Related Neutrosophic Medical Diagnosis Frameworks

Rahman et al. [3] employed a neutrosophic hypersoft set combined with a modified Sanchez's algorithm on the same Cleveland dataset and reported accuracy comparable to existing neutrosophic methods. Ihsan et al. [10] extended this line to fuzzy parameterised neutrosophic hypersoft expert sets. The present framework differs in three respects: (i) it uses a theoretically grounded entropy-based weighting rather than expert-assigned parameters; (ii) it provides a continuous risk score amenable to ROC analysis rather than a categorical ranking; and (iii) its mathematical properties (Propositions 1–3) are explicitly proved. These differences make the proposed approach complementary to hypersoft-set-based methods and suitable for settings where automated weight derivation is preferred over expert elicitation.

7.6 Limitations and Future Directions

Several limitations of the current study merit attention. First, the dataset size ($n = 303$) is modest by modern standards; replication on larger multi-centre cardiac cohorts would strengthen the generalisability claims. Second, the membership functions are currently defined without access to class labels; a semi-supervised extension—where a small labelled subset informs the membership function parameters while the larger unlabelled pool contributes entropy estimates—could substantially improve AUC. Third, the framework currently fuses crisp features; integrating interval neutrosophic or hesitant neutrosophic representations [6, 9] could accommodate features reported as ranges (e.g., blood pressure over 24 hours). Fourth, the computational cost of the framework is $O(np)$ per prediction, making it highly scalable and amenable to edge-device deployment in wearable cardiac monitors. Future work will investigate these extensions and validate the framework on electrocardiographic time-series data extracted from open physiological signal repositories.

8. CONCLUSION

This paper has presented the Single-Valued Neutrosophic Weighted Aggregation (SVNS–WA) framework, a principled and interpretable approach to multi-attribute heart disease risk assessment that fuses five clinical indicators through neutrosophic membership functions and entropy-derived feature weights. The mathematical analysis establishes that the framework's aggregation operator is bounded, monotone in truth- and falsity-membership, and converges to a single-feature rule

when one attribute becomes fully discriminative. Applied to the UCI Cleveland Heart Disease Dataset, the framework achieves an AUC of 0.765 and a sensitivity of 83.45% without supervised training, demonstrating competitive screening ability and providing a transparent decomposition of each patient's risk profile into truth, indeterminacy, and falsity dimensions. While the accuracy gap relative to supervised classifiers reflects an inherent trade-off, the framework's interpretability, parameter-free weight derivation, and provable monotonicity properties make it a compelling tool for clinical decision support in resource-constrained or label-scarce environments. The entropy-weighting mechanism automatically elevates ST Depression to a dominant role, corroborating its established clinical importance and illustrating how neutrosophic information measures can encode domain knowledge without manual specification.

REFERENCES

[1] A. Janosi, W. Steinbrunn, M. Pfisterer, and R. Detrano, "UCI machine learning repository: Heart disease dataset," Cleveland Clinic Foundation, UCI Machine Learning Repository, 1988.

[2] A. U. Rahman, M. Saeed, M. A. Mohammed, M. M. Jaber, and B. Garcia-Zapirain, "A novel fuzzy parameterized fuzzy hypersoft set and Riesz summability approach based decision support system for diagnosis of heart diseases," *Diagnostics*, vol. 12, no. 7, p. 1546, 2022.

[3] A. U. Rahman, M. Saeed, M. A. Mohammed, S. Krishnamoorthy, S. Kadry, and F. Eid, "An integrated algorithmic MADM approach for heart diseases' diagnosis based on neutrosophic hypersoft set with possibility degree-based setting," *Life*, vol. 12, no. 5, p. 729, 2022.

[4] F. Smarandache, *Neutrosophy: Neutrosophic Probability, Set, and Logic*. Rehoboth, NM: American Research Press, 1998.

[5] H. Wang, F. Smarandache, Y. Zhang, and R. Sunderraman, "Single valued neutrosophic sets," *Multispace and Multistructure*, vol. 4, pp. 410–413, 2010.

[6] H. Garg and Nancy, "Algorithms for single-valued neutrosophic decision making based on TOPSIS and clustering methods with new distance measure," *AIMS Mathematics*, vol. 5, no. 3, pp. 2671–2693, 2020.

[7] J. Ye and S. Du, "Some distances, similarity and entropy measures for interval-valued neutrosophic sets and their relationship," *International Journal of Machine Learning and Cybernetics*, vol. 10, no. 2, pp. 347–355, 2019.

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

[9] M. Ali, L. H. Son, I. Deli, and N. D. Tien, "Bipolar neutrosophic soft sets and applications in decision making," *Journal of Intelligent & Fuzzy Systems*, vol. 33, no. 6, pp. 4077–4087, 2017.

[10] M. Ihsan, M. Saeed, A. M. Alanzi, and H. E.-W. Khalifa, "An algorithmic multiple attribute decision-making method for heart problem analysis under neutrosophic hypersoft expert set with fuzzy parameterized degree-based setting," *PeerJ Computer Science*, vol. 9, p. e1607, 2023.

[11] M. Kamran, N. Salamat, S. Ashraf, M. A. Alam, and I. N. Cangul, "Novel decision modeling for manufacturing sustainability under single-valued neutrosophic hesitant fuzzy rough aggregation information," *Computational Intelligence and Neuroscience*, vol. 2022, p. 7924094, 2022.

[12] X. Luo, Z. Wang, L. Yang, L. Lu, and S. Hu, "Sustainable supplier selection based on VIKOR with single-valued neutrosophic sets," *PLoS ONE*, vol. 18, no. 9, p. e0290093, 2023.