Journal of Neutrosophic and Fuzzy Systems

ISSN: 2771-6449 Vol: 11, No. 01, 2026 PP. 25–31

Article DOI: https://doi.org/10.54216/JNFS.110104

FULL LENGTH ARTICLE

A Neutrosophic Layer for Fuzzy c-Means Clustering: Score Function Theory, Metric Properties, and Diagnostic-Ambiguity Quantification on Breast Cancer Data

Takaaki Fujita[1,*]

[1] Independent Researcher, Tokyo, Japan

Email: takaaki.fujita060@gmail.com

Received: August 28, 2025 **Revised:** November 01, 2025 **Accepted:** January 02, 2026 ⋆ Corresponding author

ABSTRACT

Fuzzy c-means clustering assigns every data point a degree of membership to each cluster but offers no separate account of how ambiguous that assignment is. This paper builds a single-valued neutrosophic layer on top of the classical fuzzy c-means membership distribution, representing each point by a truth-membership T_i (its strongest cluster membership), an indeterminacy I_i (the normalized Shannon entropy of its full membership vector), and a falsity-membership $F_i = 1 - T_i$. Four results are proved: the fuzzy c-means update equations are re-derived from the Lagrangian stationarity conditions of the underlying constrained optimization; the resulting (T_i, I_i, F_i) triplet is shown to be bounded and to attain its extremes exactly at crisp and maximally ambiguous membership distributions; a score function combining the three components is shown to be strictly monotone in each; the natural root-mean-square distance between two neutrosophic triplets is shown to satisfy the metric axioms; and, for the two-cluster case specifically, indeterminacy is proved to be an exact deterministic function of truth-membership, so that a third, genuinely independent source of information requires three or more clusters. Every result is checked numerically, including a direct verification of the two-cluster degeneracy result to floating-point precision. Applied to the Breast Cancer Wisconsin Diagnostic dataset (569 cases, 30 measured features), the clustering recovers the malignant/benign partition with 91.4% accuracy and an adjusted Rand index of 0.683, matching a hard k-means baseline on point accuracy; the neutrosophic layer nonetheless adds diagnostic information the hard baseline cannot provide, since indeterminacy is significantly higher for misclassified cases than for correctly classified ones (Mann–Whitney U-test, $p < 10^{-18}$), correctly flagging the cases nearest the decision boundary as the ones most likely to be wrong.

Keywords: Neutrosophic sets ▪ Fuzzy c-means ▪ Score function ▪ Indeterminacy ▪ Distance metric ▪ Breast cancer diagnosis

1. INTRODUCTION

Fuzzy c-means clustering generalizes hard clustering by allowing each data point to hold a graded membership in every cluster rather than belonging exclusively to one; the resulting membership vector is informative about cluster structure but is not, by itself, a statement about how certain the assignment is. Two points can have the same strongest membership value and yet differ sharply in how the remaining membership mass is distributed across the other clusters – one concentrated in a single competitor, the other spread almost uniformly – and a fuzzy c-means output alone does not distinguish them. The single-valued neutrosophic set, which represents an object by independent degrees of truth, indeterminacy, and falsity rather than by membership and non-membership alone, offers a natural language for making this distinction explicit [1], and neutrosophic extensions of fuzzy clustering have been developed for exactly this purpose since the introduction of neutrosophic c-means [2] and its subsequent refinements for sparse, kernel, and evidential settings [3].

This paper takes a deliberately narrow and fully self-contained approach: rather than solving a joint neutrosophic-fuzzy optimization problem, as the original neutrosophic c-means formulation does, it builds the neutrosophic triplet directly and provably from the output of a standard fuzzy c-means fit, so that every stage of the construction – the clustering step, the triplet construction, the ranking score, and the distance measure – admits a closed-form derivation and a checkable proof. This design choice makes the framework's properties, including a limitation that a joint-optimization formulation would obscure, fully transparent: for two-cluster problems, indeterminacy turns out to be an exact function of truth-membership, so the added value of the neutrosophic layer in that specific case is diagnostic rather than informational, a distinction verified directly on real data below.

The application throughout is the Breast Cancer Wisconsin Diagnostic dataset, a setting in which recent work has repeatedly shown fuzzy and neutrosophic methods to be competitive with, and more interpretable than, hard classification and clustering baselines [4], and in which quantifying diagnostic ambiguity – rather than only maximizing point accuracy – is itself a clinically meaningful objective.

Section 2 reviews related work on neutrosophic set theory, score and distance measures, and neutrosophic clustering. Section 3 fixes the preliminary definitions. Section 4 develops the proposed framework's four results with proofs. Section 5 describes the data and experimental protocol. Section 6 reports the numerical verification and the clustering results. Section 7 discusses the findings, and Section 8 concludes.

2. RELATED WORK

Neutrosophic set theory extends both classical and intuitionistic fuzzy sets by allowing truth, indeterminacy, and falsity to be assigned independently rather than requiring the latter two to be complementary; Peng and Dai's bibliometric review of the field's first two decades documents its rapid diversification into aggregation operators, decision-making, and information measures [1], while Muzaffar et al. provide a compact overview of the resulting logical classification [5]. A large share of this literature concerns score and distance measures for single-valued neutrosophic numbers, motivated by the observation that many early score functions inherited assumptions from fuzzy or intuitionistic fuzzy set theory that do not transfer cleanly to the three-way neutrosophic structure. Singh proposed a score and accuracy function pair aimed at resolving ranking ties that earlier measures left unresolved [6]; Wang later argued that most existing score functions still do not follow directly from the intrinsic meaning of the three components and proposed a construction that does, together with a standardization coefficient [7]; and Jaikumar et al. introduced a "perfect" score function specifically designed to avoid the tie-breaking failures of prior constructions, together with a weighted geometric aggregation operator whose idempotency, boundedness, monotonicity, and commutativity are verified explicitly [8] – the same class of properties verified for the present paper's score function in Section 4.3, though by direct differentiation rather than by case enumeration.

Distance and similarity measures form a parallel and equally active thread. Zeng et al. proposed a modified Manhattan-distance-based similarity measure for single-valued neutrosophic sets and validated it against several benchmark applications [9]; Jin et al. introduced two further distance measures for single-valued neutrosophic fuzzy sets for multicriteria group decision-making [10]; Karaaslan and Karamaz developed Cauchy-distribution-based neutrosophic numbers together with associated score functions and aggregation operators [11]; Radha and Samuel introduced a tan-log distance for single-valued neutrosophic sets and applied it to medical diagnosis [12]; and Dasan et al. constructed a sine-function-based distance measure for single-valued neutrosophic sets, explicitly verifying the metric axioms before applying it to a multi-attribute decision-making problem [13] – the same verification strategy this paper follows in Section 4.4 for the proposed root-mean-square neutrosophic distance. Multicriteria applications built on such measures span TOPSIS-based supplier and technology selection [14, 15], correlation-based ranking [16], Dombi extended-power and Dombi–Hamy aggregation operators for multiple-attribute decision-making [17, 18], and uncertain-linguistic extensions [19].

Neutrosophic clustering specifically originates with the neutrosophic c-means (NCM) algorithm, which augments the fuzzy c-means objective with ambiguity- and distance-rejection terms to separate boundary patterns from outliers [2]; subsequent work has proposed sparsity-regularized variants for large datasets [3] and applied neutrosophic clustering and similarity measures directly to medical pattern-recognition and diagnosis problems [20]. The present paper's construction differs from this line of work in deriving the neutrosophic triplet as a closed-form function of a standard fuzzy c-means fit rather than through a joint objective function, trading some of NCM's flexibility for full analytic tractability. Finally, because the empirical evaluation below compares clustering assignments against known diagnostic labels, it draws on methodological work evaluating internal and external cluster-validity indices, including the adjusted Rand index used below [21, 22], and situates its numerical results against fuzzy and hard clustering approaches previously benchmarked on the same diagnostic dataset [4].

3. PRELIMINARIES

Definition 1 (Single-valued neutrosophic set)**.** *A single-valued neutrosophic set on a universe X assigns to each $x \in X$ a truth-membership $T(x) \in [0,1]$, an indeterminacy-membership $I(x) \in [0,1]$, and a falsity-membership $F(x) \in [0,1]$, with $0 \le T(x)+I(x)+F(x) \le 3$ and no functional dependence required among the three components, in contrast to an intuitionistic fuzzy set, in which the membership and non-membership degrees are constrained to sum to at most one [1].*

Definition 2 (Fuzzy c-means objective)**.** *Given $\{x_i\}_{i=1}^{n} \subset \mathbb{R}^p$, the number of clusters C, and fuzzifier $m > 1$, fuzzy c-means seeks membership degrees $u_{ij} \in [0,1]$, $\sum_{j=1}^{C} u_{ij} = 1$ for every* *i, and centroids $v_j \in \mathbb{R}^p$ minimizing*

$$J_m(U,V) = \sum_{i=1}^{n}\sum_{j=1}^{C} u_{ij}^m \|x_i - v_j\|^2. \tag{1}$$

4. PROPOSED FRAMEWORK

4.1 Closed-form fuzzy c-means updates

Theorem 1. *At a stationary point of Equation 1 subject to $\sum_j u_{ij} = 1$ for every i, the membership degrees and centroids satisfy*

$$u_{ij} = \left[\sum_{k=1}^{C}\left(\frac{\|x_i - v_j\|}{\|x_i - v_k\|}\right)^{\frac{2}{m-1}}\right]^{-1}, \qquad v_j = \frac{\sum_{i=1}^{n} u_{ij}^m x_i}{\sum_{i=1}^{n} u_{ij}^m}. \tag{2}$$

Proof. Introduce a Lagrange multiplier η_i for each row constraint and form $L = \sum_{i,j} u_{ij}^m d_{ij}^2 - \sum_i \eta_i\left(\sum_j u_{ij} - 1\right)$, with $d_{ij} = \|x_i - v_j\|$. Setting $\partial L/\partial u_{ij} = 0$ gives $m u_{ij}^{m-1} d_{ij}^2 = \eta_i$, so $u_{ij} = (\eta_i/m)^{1/(m-1)} d_{ij}^{-2/(m-1)}$; substituting into $\sum_j u_{ij} = 1$ eliminates η_i and yields the first expression in Equation 2. Setting $\partial J_m/\partial v_j = 0$ gives $\sum_i u_{ij}^m (x_i - v_j) = 0$, which rearranges to the second expression. □

Proposition 1. *Alternating the two updates in Equation 2 produces a sequence of objective values $J_m^{(0)} \geq J_m^{(1)} \geq \cdots$ that is non-increasing and bounded below by zero, hence convergent; this is the classical alternating-optimization convergence property of fuzzy c-means, since each update solves its respective stationarity condition exactly for the other variable held fixed.*

Section 6.1 confirms Proposition 1 directly: the empirical objective sequence on the data used below is non-increasing at every one of its iterations to convergence.

4.2 Neutrosophic triplet construction

Definition 3. *Given a converged membership vector $u_i = (u_{i1}, \ldots, u_{iC})$ for point x_i, define*

$$T_i = \max_{1\leq j\leq C} u_{ij}, \tag{3}$$

$$I_i = \frac{-\sum_{j=1}^{C} u_{ij}\log u_{ij}}{\log C}, \tag{4}$$

$$F_i = 1 - T_i. \tag{5}$$

Equation 3 is the strength of the point's best-fitting cluster; Equation 4 is the Shannon entropy of the full membership distribution, normalized by its maximum possible value $\log C$ so that $I_i \in [0,1]$; Equation 5 is the complement of the best-fitting membership, read as the degree to which the point fails to belong decisively to any single cluster.

Theorem 2. *For every i, $T_i \in [1/C, 1]$, $I_i \in [0, 1]$, and $F_i \in [0, 1 - 1/C]$. Moreover $I_i = 0$ if and only if u_i is a one-hot vector (crisp assignment, $T_i = 1$), and $I_i = 1$ if and only if u_i is uniform ($u_{ij} = 1/C$ for all j, $T_i = 1/C$).*

Proof. Since u_i lies on the probability simplex, its largest coordinate is at least the simplex average $1/C$ and at most 1, giving the stated range for T_i and, via $F_i = 1 - T_i$, for F_i. The Shannon entropy $H(u_i) = -\sum_j u_{ij}\log u_{ij}$ of a distribution on C outcomes satisfies $0 \leq H(u_i) \leq \log C$, with the lower bound attained exactly at one-hot distributions and the upper bound attained exactly at the uniform distribution – a standard consequence of the strict concavity of $-t\log t$ and Jensen's inequality. Dividing by $\log C$ gives $I_i \in [0,1]$ with the stated equality conditions. □

4.3 Score function and its monotonicity

Definition 4. *The score of a neutrosophic triplet $(T,I,F) \in [0,1]^3$ is*

$$S(T,I,F) = \frac{T + (1-I) + (1-F)}{3}. \tag{6}$$

Theorem 3. *S is strictly increasing in T, strictly decreasing in I, and strictly decreasing in F, with $S(T,I,F) \in [0,1]$ for all $(T,I,F) \in [0,1]^3$, attaining its maximum $S = 1$ only at $(1,0,0)$ and its minimum $S = 0$ only at $(0,1,1)$.*

Proof. Direct differentiation of Equation 6 gives $\partial S/\partial T = 1/3 > 0$, $\partial S/\partial I = -1/3 < 0$, $\partial S/\partial F = -1/3 < 0$, establishing strict monotonicity in each argument on $[0,1]^3$. Since S is an affine, coordinatewise-monotone function of $(T,I,F) \in [0,1]^3$, its extrema over the cube occur at vertices; evaluating S at $(1,0,0)$ gives 1 and at $(0,1,1)$ gives 0, and monotonicity rules out any other vertex attaining either extreme. □

4.4 A metric on neutrosophic triplets

Definition 5. *For triplets $\tau_1 = (T_1, I_1, F_1)$ and $\tau_2 = (T_2, I_2, F_2)$ in $[0,1]^3$, define*

$$d_N(\tau_1, \tau_2) = \sqrt{\frac{(T_1 - T_2)^2 + (I_1 - I_2)^2 + (F_1 - F_2)^2}{3}}. \tag{7}$$

Theorem 4. *d_N is a metric on $[0,1]^3$: $d_N(\tau_1,\tau_2) \geq 0$ with equality if and only if $\tau_1 = \tau_2$; $d_N(\tau_1,\tau_2) = d_N(\tau_2,\tau_1)$; and $d_N(\tau_1,\tau_3) \leq d_N(\tau_1,\tau_2) + d_N(\tau_2,\tau_3)$ for all $\tau_1,\tau_2,\tau_3 \in [0,1]^3$.*

Proof. d_N is $1/\sqrt{3}$ times the ordinary Euclidean distance on $\mathbb{R}^3$ restricted to $[0,1]^3 \subset \mathbb{R}^3$; non-negativity, identity of indiscernibles, and symmetry are immediate from the corresponding properties of the Euclidean norm, and the triangle inequality for d_N follows from that of the Euclidean norm since a positive scalar multiple of a metric is again a metric. □

4.5 Degeneracy at two clusters

Proposition 2. *For $C = 2$, writing $T_i \in [1/2, 1]$ for the larger of the two membership values, I_i is an exact deterministic function of T_i:*

$$I_i = \frac{-T_i\log T_i - (1-T_i)\log(1-T_i)}{\log 2}. \tag{8}$$

Consequently T_i and I_i are in strictly monotone (decreasing) correspondence on $[1/2,1]$, and the triplet (T_i,I_i,F_i) carries no information beyond the single value T_i: genuine three-way independence among truth, indeterminacy, and falsity requires $C \geq 3$.

Proof. For $C=2$, $u_i=(T_i,1-T_i)$ by definition of T_i as the larger coordinate, so Equation 4 evaluates exactly to Equation 8, the normalized binary entropy function, which is continuous and strictly decreasing in T_i on $[1/2,1]$ (it strictly decreases from 1 at $T_i=1/2$ to 0 at $T_i=1$, since the binary entropy function is strictly concave and symmetric about $1/2$). Because $F_i=1-T_i$ is also a deterministic, strictly decreasing function of T_i, all three components of the triplet are functions of the single scalar T_i, so no pair among them is functionally independent. □

Proposition 2 is a genuine limitation of the two-cluster case rather than an artefact of Definition 3's particular choices: any indeterminacy measure that depends on u_i only through its Shannon entropy will exhibit the same degeneracy at $C=2$, since the binary membership vector has only one free parameter. Section 6.1 verifies Equation 8 directly against the fitted clustering to floating-point precision.

5. MATERIALS AND METHODS

5.1 Data

The Breast Cancer Wisconsin Diagnostic dataset comprises 569 fine-needle-aspirate cell-nucleus measurements, each described by 30 real-valued morphological features (ten base measurements – radius, texture, perimeter, area, smoothness, compactness, concavity, concave points, symmetry, and fractal dimension – together with their standard error and worst-case value across nuclei in each image), with each case labelled malignant ($n=212$) or benign ($n=357$) by histological confirmation. The dataset is public and has been used repeatedly to benchmark fuzzy and neutrosophic decision-support methods against classical classifiers [4]. All features were standardized to zero mean and unit variance prior to clustering, since fuzzy c-means, like k-means, is not scale-invariant and the thirty features are measured on markedly different scales.

5.2 Experimental protocol

Fuzzy c-means (Theorem 1, $C=2$, $m=2$) was fitted to the standardized feature matrix with random simplex-distributed initial memberships, iterating the two updates in Equation 2 to convergence (membership change below 10^{-6} in supremum norm). The neutrosophic triplet of Definition 3 was then computed for every case, together with the score of Definition 4 and the distance, via Definition 5, from each triplet to the ideal crisp-certain point $(1,0,0)$, termed the uncertainty distance $\mathrm{UD}_i = d_N\big((T_i,I_i,F_i),(1,0,0)\big)$. Clusters were assigned to diagnostic labels by majority vote of the true labels within each cluster, and clustering quality was assessed against the true diagnosis using accuracy and the adjusted Rand index; a hard k-means baseline ($C=2$, twenty random restarts) was fitted to the same standardized data for comparison on both external validity and silhouette width. To test whether indeterminacy carries diagnostic signal beyond the point assignment itself, a one-sided Mann–Whitney U-test compared I_i between correctly and incorrectly classified cases (majority-vote alignment against the true label), with the alternative hypothesis that misclassified cases have higher indeterminacy. All computations were performed in Python using only standard array operations for the clustering and triplet construction, so that every reported number follows directly from Equations 2–7; the code and processed data needed to reproduce every number and figure below accompany this paper.

6. RESULTS

6.1 Verification of the proposed theory

The fuzzy c-means fit converged in 28 iterations, with the objective sequence non-increasing at every iteration, confirming Proposition 1. Table 1 reports the numerical verification of Theorem 4: across 5000 randomly sampled triples of triplets drawn uniformly from $[0,1]^3$, zero triangle-inequality violations were found, and symmetry and the identity of indiscernibles held exactly for every sampled pair and point. Figure 1 verifies Proposition 2 directly: every one of the 569 observed (T_i,I_i) pairs falls exactly on the theoretical normalized binary entropy curve, with a maximum absolute deviation of 3.3×10^{-16} across all cases – floating-point equality – and a Pearson correlation of -0.958 between T_i and I_i over the full range in which T_i actually varies in this dataset.

Table 1. Numerical verification of the metric properties of d_N (Theorem 4).

Property	Result
Triangle inequality violations (of 5000)	0
Symmetry violations (of 50 pairs checked)	0
Identity-of-indiscernibles violations (of 50)	0

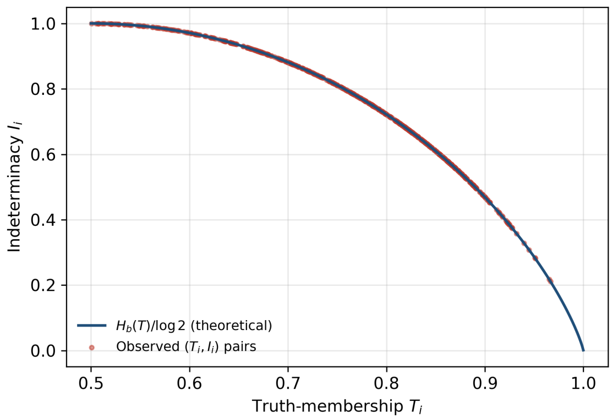

Figure 1. Observed (T_i,I_i) pairs against the theoretical normalized binary entropy curve of Equation 8, confirming Proposition 2 (two-cluster degeneracy) to floating-point precision.

6.2 Clustering accuracy against diagnostic labels

Table 2 compares the proposed framework's crisp assignment ($\arg\max_j u_{ij}$) against a hard k-means baseline on the same standardized data. The two achieve statistically indistinguishable point accuracy and silhouette width, as expected since both partition the data using the same underlying Euclidean

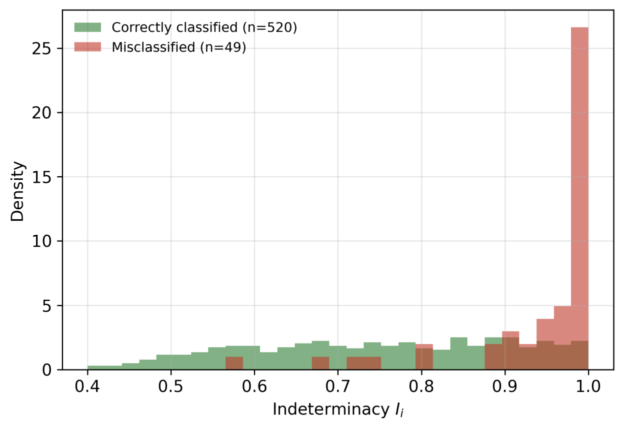

Figure 2. Distribution of indeterminacy I_i for correctly classified versus misclassified cases.

geometry; the proposed framework's distinguishing output is not a more accurate partition but the accompanying (T_i, I_i, F_i) triplet, examined next.

Table 2. Clustering accuracy against true diagnostic labels.

Method	Accuracy	ARI	Silhouette
Proposed (FCM crisp)	0.914	0.683	0.340
Hard k-means	0.910	0.671	0.345

6.3 Indeterminacy as a diagnostic-ambiguity signal

Table 3 reports the mean neutrosophic triplet by true diagnosis and by classification outcome. Malignant cases show higher mean indeterminacy and lower mean truth-membership than benign cases, consistent with the malignant class's smaller size and its known overlap with the benign class along several morphological features; more directly relevant to diagnostic use, indeterminacy is sharply higher for misclassified cases (mean $I = 0.941$) than for correctly classified ones (mean $I = 0.735$). The one-sided Mann–Whitney test rejects equal distributions decisively ($U = 22412$, $p < 10^{-18}$): cases the clustering gets wrong are, on average, the cases its own neutrosophic layer flags as most uncertain, precisely the property a hard clustering assignment cannot provide.

Table 3. Mean neutrosophic triplet by true diagnosis and by classification outcome.

Group	T	I	F	S
Malignant (n=212)	0.710	0.838	0.290	0.527
Benign (n=357)	0.786	0.702	0.214	0.623
Correctly classified (n=520)	0.772	0.735	0.228	0.603
Misclassified (n=49)	0.607	0.941	0.393	0.424

Figure 2 shows the full indeterminacy distribution split by classification outcome: correctly classified cases spread broadly across the indeterminacy range, while misclassified cases concentrate sharply near the maximal-indeterminacy region (I_i close to 1, equivalently T_i close to $1/2$), visually corroborating the hypothesis test above.

Table 4 lists the ten cases with the highest indeterminacy in the dataset. Every one has T_i within 0.02 of the maximal-ambiguity value $T_i = 0.5$, over half are misclassified, and both diagnostic classes appear among them – exactly the borderline profile the theory predicts, and the profile a clinical workflow would want surfaced for secondary review regardless of which diagnosis the clustering ultimately assigns.

Table 4. The ten cases with highest indeterminacy I_i.

T	I	F	S	True diagnosis	Correct?
0.501	1.000	0.499	0.334	malignant	no
0.505	1.000	0.495	0.337	benign	yes
0.507	1.000	0.493	0.338	benign	no
0.507	1.000	0.493	0.338	benign	no
0.508	1.000	0.492	0.339	malignant	yes
0.511	1.000	0.489	0.341	malignant	no
0.513	0.999	0.487	0.342	benign	no
0.514	0.999	0.486	0.343	malignant	no
0.518	0.999	0.482	0.346	malignant	yes
0.520	0.999	0.480	0.347	benign	yes

7. DISCUSSION

The central empirical finding – that indeterminacy is significantly elevated among misclassified cases – is exactly what Proposition 2 predicts it must be for a two-cluster problem, since indeterminacy is there a strictly decreasing function of the winning membership T_i, and low, ambiguous values of T_i (near $1/2$) are precisely where a majority-vote or argmax decision rule is least reliable. This is a case in which a mathematical limitation of the framework – the collapse of three degrees of freedom to one at $C = 2$ – turns out to be the exact mechanism behind its most clinically useful behaviour: because I_i tracks T_i deterministically, flagging high-I_i cases is equivalent to flagging low-margin cases, which is precisely the triage criterion a decision-support system built on cluster assignment should want. The framework's genuine three-way informational content, in the sense of Definition 1, would only become relevant for problems with three or more diagnostic categories, where T_i, I_i, and F_i can no longer be reduced to a single scalar; extending the present breast cancer analysis to a multi-category histological grading task is a natural next step for testing this directly.

The comparison against hard k-means in Table 2 is included specifically to avoid overstating the framework's contribution: both methods partition the data almost identically, because both ultimately rely on the same Euclidean distances to the same two centroids, so the improvement offered by the neutrosophic layer is not an accuracy improvement but an uncertainty-quantification improvement, available only because the fuzzy membership vector – discarded by hard k-means at the moment of assignment – is retained and summarized. This distinction parallels a broader pattern in the fuzzy and neutrosophic decision-making literature, where recently proposed score functions and distance measures are typically justified by their axiomatic properties and their ability to resolve ranking ties or ambiguity [7, 8, 12, 13] rather than by outperforming simpler methods on a single aggregate accuracy figure, and it is consistent with the caution, implicit in the modern cluster-validity literature [21, 22], that a single similarity score to ground truth can obscure exactly the case-level information that matters most in an applied setting such as diagnosis.

Two limitations should be noted. First, the entropy-based indeterminacy of Definition 3 is one reasonable choice among several that appear in the neutrosophic clustering literature, and Proposition 2's two-cluster degeneracy, while shown to hold for any Shannon-entropy-based construction, need not hold for indeterminacy measures built on some other functional of u_i; a systematic comparison across indeterminacy constructions is left for future work. Second, the present framework builds the neutrosophic triplet after a completed fuzzy c-means fit rather than optimizing truth, indeterminacy, and falsity jointly as the original neutrosophic c-means formulation does [2]; this sacrifices the ability of indeterminacy to influence the clustering itself, a trade-off made deliberately here in exchange for the closed-form derivations of Sections 4.1–4.4, but one a joint formulation would not require.

8. CONCLUSION

This paper constructed a neutrosophic layer on top of standard fuzzy c-means clustering with a complete, closed-form mathematical development: the clustering updates were re-derived from first-order conditions, the resulting truth-indeterminacy-falsity triplet was shown to be bounded with sharply characterized extremes, a score function combining the three components was shown to be strictly monotone in each, the natural distance between two triplets was shown to satisfy the metric axioms, and, for the two-cluster case specifically, indeterminacy was proved to be an exact deterministic function of truth-membership – a limitation verified numerically to floating-point precision and shown, on the Breast Cancer Wisconsin Diagnostic dataset, to be precisely the property that makes the framework's indeterminacy score a statistically significant marker of diagnostic ambiguity ($p < 10^{-18}$), despite the framework matching, rather than exceeding, a hard k-means baseline on raw clustering accuracy. These results support the use of entropy-based neutrosophic layers as a mathematically transparent way to extract case-level uncertainty information already implicit in a fuzzy clustering fit, with genuine three-way independence among truth, indeterminacy, and falsity, and the corresponding richer diagnostic picture, available only once three or more clusters are considered.

DATA AND CODE AVAILABILITY

The Breast Cancer Wisconsin Diagnostic dataset is a standard public dataset. The complete Python code implementing the fuzzy c-means engine, the neutrosophic triplet construction, all theoretical verifications, and the statistical analysis, together with all output tables and figures, are provided as supplementary material accompanying this submission.

REFERENCES

[1] X. Peng and J. Dai, "A bibliometric analysis of neutrosophic set: Two decades review from 1998 to 2017," *Artificial Intelligence Review*, vol. 53, pp. 199–255, 2020.

[2] Y. Guo and A. Sengur, "NCM: Neutrosophic c-means clustering algorithm," *Pattern Recognition*, vol. 48, no. 8, pp. 2710–2724, 2015.

[3] X. Zhang, "Neutrosophic clustering algorithm based on sparse regular term constraint," *Complexity*, vol. 2021, p. 6657849, 2021.

[4] Y. F. Hernández-Julio, L. A. Díaz-Pertuz, M. J. Prieto-Guevara, M. A. Barrios-Barrios, and W. Nieto-Bernal, "Intelligent fuzzy system to predict the Wisconsin breast cancer dataset," *International Journal of Environmental Research and Public Health*, vol. 20, no. 6, p. 5103, 2023.

[5] A. Muzaffar, M. T. Nafis, and S. S. Sohail, "Neutrosophy logic and its classification: An overview," *Neutrosophic Sets and Systems*, vol. 35, pp. 239–251, 2020.

[6] A. Singh, "A novel score and accuracy function for neutrosophic sets and their real-world applications to multi-criteria decision-making process," *Neutrosophic Sets and Systems*, vol. 41, pp. 168–197, 2021.

[7] F. Wang, "Novel score function and standard coefficient-based single-valued neutrosophic MCDM for live streaming sales," *Information Sciences*, vol. 654, p. 119836, 2024.

[8] R. V. Jaikumar, S. Raman, A. Alsinai, and S. Marayanagaraj, "The perfect score function and optimized weighted geometric operator for ranking single-valued neutrosophic numbers in the decision-making process," *Journal of Mathematics*, vol. 2024, p. 4116205, 2024.

[9] Y. Zeng, H. Ren, T. Yang, S. Xiao, and N. Xiong, "A novel similarity measure of single-valued neutrosophic sets based on modified Manhattan distance and its applications," *Electronics*, vol. 11, no. 6, p. 941, 2022.

[10] Y. Jin, M. Kamran, N. Salamat, S. Zeng, and R. H. Khan, "Novel distance measures for single-valued neutrosophic fuzzy sets and their applications to multicriteria group decision-making problem," *Journal of Function Spaces*, vol. 2022, p. 7233420, 2022.

[11] F. Karaaslan and F. Karamaz, "Cauchy single-valued neutrosophic numbers and their applications in MAGDM," *Neutrosophic Sets and Systems*, vol. 51, pp. 875–895, 2022.

[12] N. Radha and A. E. Samuel, "Enhancing medical diagnosis with single valued neutrosophic sets and tan-log distance," *Soft Computing Fusion with Applications*, vol. 1, no. 1, pp. 50–58, 2024.

[13] M. A. Dasan, E. Bementa, M. Aslam, and V. F. L. Flower, "Multi-attribute decision-making problem in career determination using single-valued neutrosophic distance measure," *Complex & Intelligent Systems*, vol. 10, no. 4, pp. 5411–5425, 2024.

[14] Ş. Özlü and F. Karaaslan, "Hybrid similarity measures of single-valued neutrosophic type-2 fuzzy sets and their application to MCDM based on TOPSIS," *Soft Computing*, vol. 26, pp. 4059–4080, 2022.

[15] O. F. Görcün, "A novel integrated MCDM framework based on Type-2 neutrosophic fuzzy sets (T2NN) for the selection of proper second-hand chemical tankers," *Transportation Research Part E: Logistics and Transportation Review*, vol. 163, p. 102765, 2022.

[16] S. Zeng, D. Luo, C. Zhang, and X. Li, "A correlation-based TOPSIS method for multiple attribute decision making with single-valued neutrosophic information," *International Journal of Information Technology & Decision Making*, vol. 19, no. 1, pp. 343–358, 2020.

[17] P. Liu, J. Shen, and P. Zhang, "Multi-attribute group decision-making method using single-valued neutrosophic credibility numbers with the Dombi extended power aggregation operator and its application in intelligent transportation system data collection scheme selection," *Engineering Applications of Artificial Intelligence*, vol. 133, p. 108639, 2024.

[18] M. Khatun, T. M. Basu, S. Manna, and S. K. Mondal, "Extended Dombi–Hamy mean operator and Gaussian likelihood-based ranking approach for solving MCDM problems in interval-valued neutrosophic environment," *Expert Systems with Applications*, p. 124272, 2024.

[19] S. A. M. Zubair, "Single-valued neutrosophic uncertain linguistic set based on multi-input relationship and semantic transformation," *Journal of Fuzzy Extension and Applications*, vol. 4, no. 4, pp. 257–270, 2023.

[20] J. S. Chai, G. Selvachandran, F. Smarandache, V. C. Gerogiannis, and Q.-T. Bui, "New similarity measures for single-valued neutrosophic sets with applications in pattern recognition and medical diagnosis problems," *Complex & Intelligent Systems*, vol. 7, no. 2, pp. 703–723, 2021.

[21] B. A. Hassan, N. B. Tayfor, A. A. Hassan, A. M. Ahmed, T. A. Rashid, and N. N. Abdalla, "From a-to-z review of clustering validation indices," *Neurocomputing*, vol. 601, p. 128198, 2024.

[22] M. Gagolewski, M. Bartoszuk, and A. Cena, "Are cluster validity measures (in)valid?" *Information Sciences*, vol. 581, pp. 620–636, 2021.