Recurrence Shadow Mapping under Neutrosophic Clinical Evidence: An Uncertainty-Oriented Model for Post-Treatment Healthcare Decision Support

Murodbek Ahrorov[1,*], Ahmed Aziz[2]
[1]Medical School, Central Asian University, Uzbekistan
[2]Central Asian University, Uzbekistan
Emails: m.akhrorov@centralasian.uz; aziz.bfci@gmail.com

Abstract

Post-treatment follow-up in differentiated thyroid cancer requires a decision model that is not limited to binary recurrence prediction. Patients may present with partially reassuring anatomical findings, incomplete biochemical response, hetero-geneous pathological subtype, or contradictory clinical history. These situations are better described as a triadic state composed of support for recurrence, support against recurrence, and unresolved indeterminacy. This paper proposes a recurrence shadow mapping model based on single-valued neutrosophic clinical evidence. The model transforms clinico-pathologic descriptors into truth, indeterminacy, and falsity memberships; aggregates evidence through entropy-contrast weighting; and produces a recurrence-shadow index that separates stable, observation, alert, and high-alert follow-up states. The proposed method is designed for healthcare decision support rather than automatic replacement of clinical judgment. Its mathematical contribution is a bounded neutrosophic score that penalizes inconsistent evidence without suppressing clinically meaningful warning signals. Experimental evaluation demonstrates that recurrence-oriented evidence sources can be expressed in a transparent mathematical form, and that indeterminacy itself becomes an interpretable clinical quantity. The findings support the use of neutrosophic information fusion for medical cases where uncertainty is structural rather than merely statistical.

Keywords: Neutrosophic healthcare decision support; Recurrence shadow; Differentiated thyroid cancer; Single-valued neutrosophic set; Clinical evidence fusion; Indeterminacy-aware risk stratification

1. Introduction

Healthcare decision support increasingly depends on models that combine heterogeneous clinical evidence. In oncology follow-up, however, the evidence rarely follows a clean binary pattern. A patient can show low anatomical disease burden while still having an indeterminate treatment response, or a favorable pathological subtype while carrying metastatic or nodal indicators. Classical predictive models usually collapse these conditions into a probability of recurrence. Although probability is useful, it does not explicitly separate affirmative evidence, contradictory evidence, and unresolved evidence. This limitation is important in post-treatment cancer surveillance, where follow-up intensity is shaped not only by the predicted event but also by the credibility and conflict of the underlying evidence.

Differentiated t hyroid c ancer i s c linically s uitable f or t his p roblem b ecause r ecurrence r isk i s i nfluenced by a mixture of demographic, pathological, anatomical, and treatment-response variables. The public cohort associated with the UCI Differentiated Thyroid Cancer Recurrence resource was donated in 2023 and contains 383 patient records, 16 features, and no missing values, with follow-up of at least ten years across a 15-year period r (o,o). Recent machine-learning studies on thyroid cancer recurrence show the value of computational stratification, but most are optimized for classification accuracy or interpretability of conventional predictors rather than explicit representation of indeterminacy a (l); m (h). Clinical guidelines also emphasize risk-adapted follow-up, which motivates models that distinguish high-confidence stability from uncertain observation states d (a).

DOI: https://doi.org/10.54216/NIF.050201
Received: March 01, 2025 Accepted: June 28, 2025

Neutrosophic sets provide a mathematically natural language for this setting. Instead of assigning each patient to a single certainty scale, a single-valued neutrosophic representation assigns three degrees: truth membership T, indeterminacy membership I, and falsity membership F. Recent developments in single-valued neutrosophic aggregation, dynamic operators, EDAS methods, and healthcare service evaluation show that the triadic representation is effective when evidence is incomplete, conflicting, or distributed across multiple criteria r (a,a); v (r). Yet healthcare applications often use neutrosophic operators as generic multi-criteria ranking tools, while disease follow-up requires disease-specific mapping between clinical meaning and neutrosophic geometry.

This paper develops such a mapping. The proposed recurrence shadow mapping model treats recurrence not as a single output label, but as a region in a neutrosophic clinical evidence space. The truth component quantifies evidence supporting recurrence, the falsity component quantifies evidence against recurrence, and the indeterminacy component quantifies contradiction between sources. This design is intended to make the model clinically inspectable: an elevated score due to anatomical burden differs from an elevated shadow state due to indeterminate response and inconsistent evidence. The novelty lies in converting follow-up ambiguity into a measurable shadow region rather than discarding it as model noise.

The main contributions are fourfold. First, the paper defines a source-wise neutrosophic encoding for clinical recurrence evidence that preserves the meaning of anatomical burden, response status, pathology, and patient-history pressure. Second, it proposes an entropy-contrast weighting rule that gives larger influence to evidence sources with higher separation capacity and lower distributional randomness. Third, it introduces a bounded recurrence-shadow index that penalizes indeterminacy while still allowing truth evidence to dominate when the clinical signal is strong. Fourth, it provides an implementation protocol with numerical analysis, ablation, sensitivity testing, and decision-zone interpretation. These elements make the work suitable for neutrosophic information fusion research and healthcare decision-support applications.

2. Related Work

The literature around this problem can be divided into three intersecting streams: computational thyroid-cancer decision support, neutrosophic aggregation and ranking, and uncertainty-aware healthcare evaluation. Borzooei et al. r (o) presented a machine-learning cohort study for well-differentiated thyroid cancer recurrence and showed that clinicopathologic variables can support risk stratification over long follow-up. Clark et al. a (l) further considered predictive analytics for thyroid cancer recurrence, while recent explainable AI work has explored feature-attribution approaches for recurrence prediction m (h). These studies support the use of computational models but typically represent uncertainty through model scores, not through an explicit truth-indeterminacy-falsity structure.

Clinical thyroid-cancer management is also shaped by guideline-based recurrence risk and surveillance decisions. The NCCN thyroid carcinoma guidelines provide a clinical framework for treatment and follow-up intensity d (a). Related work in pathology report processing has shown that structured and unstructured thyroid cancer data can be converted into risk categories through natural language processing o (o). These studies motivate transparent decision-support systems, especially when the data source is heterogeneous. Nevertheless, they leave open the question of how to mathematically represent a patient whose clinical signs are partly affirmative, partly reassuring, and partly unresolved.

The second stream concerns single-valued neutrosophic aggregation. Farid and Riaz r (a) proposed Einstein interactive aggregation operators that preserve interaction among truth, indeterminacy, and falsity. Their later dynamic aggregation work extended the idea to multi-period decision problems r (a). Ali et al. i (l) introduced parameterized neutrosophic aggregation operators that allow the influence of membership, indeterminacy, and non-membership components to be adjusted. These contributions are mathematically important because clinical follow-up also has a dynamic and parameter-sensitive nature.

Decision-making techniques under single-valued neutrosophic information have expanded rapidly. EDAS-based neutrosophic approaches have been developed for alternatives evaluation u (i), while decision-support formulations based on single-valued neutrosophic information were further examined in recent mathematical models t (a). Healthcare service quality evaluation under single-valued trapezoidal neutrosophic information has been addressed through an improved RAFSI method v (r). These works confirm that neutrosophic MCDM can handle subjective, incomplete, or imprecise healthcare evidence. The present paper differs by constructing a disease-specific recurrence shadow instead of ranking generic service alternatives.

The third stream concerns healthcare AI and uncertainty. AI techniques for thyroid cancer diagnosis and follow-up continue to evolve, including image-based diagnosis, predictive recurrence analysis, and decision-support tools b (a). Machine-learning models have also been applied to radiotherapy-related hypothyroidism complications in head and neck cancer e (e). Such studies show that clinical AI is not only a classification task; it is a workflow task requiring interpretability, caution, and medically meaningful uncertainty. Neutrosophic information fusion provides an additional mathematical layer because it treats uncertainty as an explicit component of the patient state rather than a residual error.

DOI: https://doi.org/10.54216/NIF.050201
Received: March 01, 2025 Accepted: June 28, 2025

3. Mathematical Model

Let $X = \{x_1, x_2, \dots, x_n\}$ be a set of patient records and let $S = \{s_1, s_2, s_3, s_4\}$ denote four clinical evidence sources: anatomical burden, guideline-response status, pathological pattern, and patient-history pressure. Each source is transformed into a normalized evidence value $e_{ij} \in [0, 1]$, where larger values indicate stronger support for recurrence. The patient-level evidence vector is

$$\mathbf{e}_i = (e_{i1}, e_{i2}, e_{i3}, e_{i4}). \tag{1}$$

The single-valued neutrosophic clinical state of patient x_i is

$$\mathcal{N}(x_i) = \langle T_i, I_i, F_i \rangle, \qquad T_i, I_i, F_i \in [0, 1]. \tag{2}$$

The truth membership is obtained through entropy-contrast weighted fusion:

$$T_i = \sum_{j=1}^{m} w_j e_{ij}, \qquad \sum_{j=1}^{m} w_j = 1, \quad w_j \geq 0. \tag{3}$$

For source s_j, the normalized histogram entropy is

$$H_j = -\frac{1}{\log B} \sum_{b=1}^{B} p_{jb} \log(p_{jb}), \tag{4}$$

where p_{jb} is the proportion of observations in bin b. The proposed weight is

$$w_j = \frac{1 - H_j + \epsilon}{\sum_{q=1}^{m} (1 - H_q + \epsilon)}, \tag{5}$$

with $\epsilon > 0$ preventing zero weights. Thus, a source receives a larger weight when its values are less random and more discriminative across the population.
Indeterminacy is defined as a contradiction-sensitive quantity:

$$I_i = \min\{1, \eta_0 + \eta_1 \sigma(\mathbf{e}_i) + \eta_2 \mathbb{1}_{R_i = \text{Indeterminate}} + \eta_3 \mathbb{1}_{C_i}\}, \tag{6}$$

where $\sigma(\mathbf{e}_i)$ is the standard deviation among source evidences, R_i is the response category, and C_i is a clinical conflict condition such as low guideline risk with non-excellent response. Falsity is not simply $1 - T_i$ because indeterminacy weakens the meaning of direct opposition. Therefore,

$$F_i = (1 - T_i)(1 - \rho I_i), \qquad 0 \leq \rho \leq 1. \tag{7}$$

This means that when indeterminacy is high, the evidence against recurrence is discounted because the case is less confidently stable.
The recurrence shadow index is

$$\Psi_i = T_i - F_i - \alpha I_i + \beta T_i (1 - I_i), \tag{8}$$

where α controls the penalty for unresolved evidence and β rewards strong truth evidence that remains coherent. The decision zones are

$$\mathcal{Z}(x_i) = \begin{cases} \text{stable shadow-low}, & \Psi_i < -0.30, \\ \text{observation shadow}, & -0.30 \leq \Psi_i < -0.05, \\ \text{recurrence-alert}, & -0.05 \leq \Psi_i < 0.25, \\ \text{high recurrence-alert}, & \Psi_i \geq 0.25. \end{cases} \tag{9}$$

Proposition 1. If $T_i, I_i, F_i \in [0, 1]$, $\alpha \in [0, 1]$, and $\beta \in [0, 1]$, then $\Psi_i \in [-2, 2]$. If additionally $F_i = (1 - T_i)(1 - \rho I_i)$ with $\rho \in [0, 1]$, then Ψ_i is monotone nondecreasing in T_i when I_i is fixed.
Proof. Since $T_i, F_i, I_i, T_i(1 - I_i) \in [0, 1]$, the lower bound follows from $T_i = 0$, $F_i = 1$, and $I_i = 1$, giving at least -2. The upper bound follows from $T_i = 1$, $F_i = 0$, $I_i = 0$, and $\beta = 1$, giving at most 2. For fixed I_i, substituting $F_i = (1 - T_i)(1 - \rho I_i)$ into (8) gives

$$\frac{\partial \Psi_i}{\partial T_i} = 1 + (1 - \rho I_i) + \beta(1 - I_i) \geq 0. \tag{10}$$

Thus, stronger recurrence-supporting evidence cannot reduce the recurrence shadow score. □
Lemma 1. For two patients x_a and x_b with equal T and F, if $I_a > I_b$ and $\alpha + \beta T > 0$, then $\Psi_a < \Psi_b$.

DOI: https://doi.org/10.54216/NIF.050201
Received: March 01, 2025 Accepted: June 28, 2025

Proof. For fixed T and F, the difference is

$$\Psi_a - \Psi_b = -\alpha(I_a - I_b) - \beta T(I_a - I_b) = -(\alpha + \beta T)(I_a - I_b) < 0. \quad (11)$$

Therefore, unresolved or contradictory evidence decreases the actionable recurrence score. □

4. Implementation Protocol

The implementation uses the public differentiated thyroid cancer recurrence cohort, which contains individual patient records, real and categorical features, and no missing values r (o). Each categorical feature is mapped into a clinically ordered recurrence pressure whenever a natural order exists. For example, $M1$ receives greater evidence than $M0$, and structural incomplete response receives greater evidence than excellent response. Non-ordinal categories are mapped through clinically plausible recurrence pressure levels and then evaluated through the source-wise fusion equations. The aim is not to replace specialist interpretation, but to create a mathematically inspectable recurrence-shadow layer.

Table 1: Sensitivity of the recurrence-shadow score to indeterminacy penalty alpha.

α	Accuracy	Balanced accuracy	F1	Predicted recurrence rate
0.20	0.801	0.778	0.674	0.360
0.30	0.829	0.798	0.704	0.326
0.40	0.846	0.812	0.725	0.303
0.50	0.849	0.816	0.731	0.292
0.60	0.844	0.808	0.719	0.284
0.70	0.835	0.794	0.701	0.271
0.80	0.824	0.779	0.684	0.258
0.90	0.812	0.761	0.661	0.242
1.00	0.795	0.742	0.638	0.229
1.10	0.781	0.726	0.616	0.214

Table 1 evaluates the penalty parameter α. Increasing α reduces the predicted recurrence rate because I_i subtracts more strongly from Ψ_i. The best region is therefore not necessarily the largest penalty; clinically useful behavior occurs when indeterminacy suppresses ambiguous alerts without eliminating high-truth cases.
Table 2 lists factor levels with high observed recurrence proportions. These levels justify the ordinal evidence encodings used in e_{ij}. The table also prevents the neutrosophic model from being purely abstract: every membership transformation is linked to observed recurrence pressure in a clinicopathologic category.

5. Extended Numerical Analysis

This section expands the numerical interpretation of the recurrence-shadow model. The purpose is to make the neutrosophic contribution measurable rather than descriptive. The analysis therefore reports decision-zone behavior, operator properties, fold-level stability, contradiction relations, and clinical follow-up implications. These tables also clarify that the proposed model is not simply a classifier with a different name. It is a structured evidence-fusion mechanism in which the final decision is affected by truth support, falsity support, and indeterminacy pressure simultaneously.

Table 3: Fold-level stability of the proposed recurrence-shadow model.

Fold	Accuracy	Balanced accuracy	Precision	Recall	F1
1	0.844	0.806	0.714	0.727	0.720
2	0.831	0.792	0.692	0.682	0.687
3	0.857	0.821	0.733	0.762	0.747
4	0.842	0.807	0.708	0.714	0.711
5	0.829	0.794	0.704	0.719	0.711

Table 3 indicates that the proposed recurrence-shadow model is not dependent on a single favorable partition. The balanced accuracy remains close across folds, which means that the score Ψ_i does not merely exploit the majority non-recurrence class. Mathematically, this stability is expected because the threshold is applied to a bounded score obtained from normalized memberships rather than to unbounded classifier logits.

DOI: https://doi.org/10.54216/NIF.050201
Received: March 01, 2025 Accepted: June 28, 2025

Table 2: Highest empirical recurrence proportions across clinical factor levels.

Clinical factor	Level	Observed recurrence %	n
Risk	High	62.745	51.000
Risk	Intermediate	32.184	87.000
Risk	Low	19.592	245.0
Response	Structural Incomplete	59.524	42.000
Response	Biochemical Incomplete	44.444	45.000
Response	Indeterminate	36.047	86.000
Response	Excellent	15.238	210.0
Adenopathy	Right	44.737	38.000
Adenopathy	Extensive	42.857	7.000
Adenopathy	Left	40.816	49.000
Adenopathy	Bilateral	33.333	24.000
T	T4b	60.000	15.000
T	T3a	45.455	44.000
T	T4a	40.000	25.000
T	T3b	28.571	28.000
N	N1b	54.348	46.000
N	N1a	44.444	81.000
N	N0	18.359	256.0
M	M1	86.957	23.000
M	M0	24.444	360.0
Pathology	Hurthel cell	56.250	16.000
Pathology	Follicular	40.000	35.000

Table 4: Operator-property checklist for the proposed recurrence-shadow mapping.

Property	Mathematical condition	Clinical interpretation
Boundedness	$T, I, F \in [0, 1]$ and $\Psi \in [-2, 2]$	No patient receives an unbounded alert value.
Truth monotonicity	$\partial\Psi/\partial T \geq 0$ for fixed I	Stronger recurrence evidence cannot reduce the alert score.
Indeterminacy penalty	$\partial\Psi/\partial I = -\alpha - \beta T + \rho(1-T) \leq 0$ under selected parameters	Contradictory evidence is treated cautiously.
Falsity opposition	$\partial\Psi/\partial F = -1$	Clear evidence against recurrence lowers the action score.
Source transparency	$T = \sum_j w_j e_j$	Each clinical evidence source remains traceable.

Table 4 summarizes the mathematical safeguards built into the model. These safeguards are important for healthcare use because a decision-support score must behave predictably. The monotonicity condition prevents clinically illogical reversals, while the indeterminacy penalty ensures that contradictory records are not treated with the same certainty as coherent high-risk records.

Table 5: Average pairwise contradiction among evidence sources.

Source	Anatomy	Response	Pathology	History
Anatomy	0.000	0.214	0.168	0.246
Response	0.214	0.000	0.192	0.221
Pathology	0.168	0.192	0.000	0.177
History	0.246	0.221	0.177	0.000

DOI: https://doi.org/10.54216/NIF.050201
Received: March 01, 2025 Accepted: June 28, 2025

Table 5 reports mean absolute differences between source evidences. The largest contradiction is between anatomical burden and patient-history pressure, showing that background risk modifiers do not always agree with structural disease indicators. In the proposed model, this disagreement contributes to $\sigma(\mathbf{e}_i)$ and therefore to I_i. This is mathematically preferable to averaging all sources directly because disagreement becomes visible rather than hidden inside T_i.

Table 6: Clinical decision interpretation of recurrence-shadow zones.

Zone	Mathematical rule	Suggested decision-support interpretation
Stable shadow-low	$\Psi < -0.30$	Low alert state; routine follow-up may be sufficient when consistent with clinical judgment.
Observation shadow	$-0.30 \leq \Psi < -0.05$	Evidence is not alarming but not fully stable; repeated assessment or closer interval can be considered.
Recurrence-alert	$-0.05 \leq \Psi < 0.25$	Recurrence-supporting evidence is present; additional diagnostic confirmation may be justified.
High recurrence-alert	$\Psi \geq 0.25$	Strong recurrence-oriented evidence with reduced falsity; urgent specialist review may be prioritized.

Table 6 translates the mathematical zones into decision-support language. The table intentionally avoids prescribing treatment, because the model is not a clinical authority. The mathematical rule is useful because it gives the clinician a transparent reason for the suggested follow-up intensity: either high truth membership, low falsity membership, or elevated indeterminacy near a boundary.

Table 7: Parameter roles and recommended calibration strategy.

Parameter	Role	Calibration meaning
ϵ	Weight smoothing	Prevents a source with high entropy from being fully discarded.
ρ	Falsity discount	Controls how much indeterminacy weakens evidence against recurrence.
α	Indeterminacy penalty	Controls the caution level of the recurrence-shadow score.
β	Coherent truth reward	Reinforces truth evidence when indeterminacy is limited.
$\eta_0, \eta_1, \eta_2, \eta_3$	Indeterminacy formation	Determine the baseline, source-disagreement, response, and conflict contributions to I.

Table 7 clarifies that the proposed method can be institutionally calibrated. A center that prefers conservative surveillance may increase α or lower the alert threshold, while a center that wants fewer alerts may increase the required truth membership. This parameterization is a major advantage of neutrosophic modeling because uncertainty policy can be separated from evidence measurement.

Table 8: Comparison between the proposed neutrosophic logic and conventional risk scoring.

Criterion	Conventional recurrence score	Proposed recurrence-shadow score
Evidence representation	One probability or risk value	Three memberships T, I, and F
Conflict handling	Implicit in model residuals	Explicit through I_i and source dispersion
Decision boundary	Single threshold	Multi-zone shadow interpretation
Interpretability	Feature importance or coefficients	Source evidence, memberships, and zone rule
Clinical ambiguity	Usually treated as lower confidence	Preserved as a measurable follow-up state

Table 8 positions the contribution of the paper. Conventional risk scoring is useful when evidence is coherent, but it compresses conflict and falsity into the same output. The proposed recurrence-shadow score keeps these quantities separate

DOI: https://doi.org/10.54216/NIF.050201
Received: March 01, 2025 Accepted: June 28, 2025

until the final step. Mathematically, this gives the clinician more information: two patients can have the same Ψ_i but different T, I, and F profiles, leading to different explanations.

6. Implications for Neutrosophic Healthcare Applications

The proposed model illustrates how neutrosophic theory can be used as a healthcare-oriented modeling language rather than only as an aggregation technique. In medical data, uncertainty is often structural. It arises from temporal delay, imperfect tests, incomplete response, and disagreement among clinical indicators. Treating all uncertainty as statistical variance can obscure these sources. By assigning a separate indeterminacy membership, the model allows the uncertainty itself to become part of the clinical output.

A second implication is that neutrosophic decision support can be aligned with follow-up workflows. Many healthcare systems do not need only a yes/no classifier; they need a triage layer that distinguishes stable patients, observation patients, alert patients, and urgent review patients. The recurrence shadow index supports this need because it is continuous enough for ranking but structured enough for categorical follow-up zones. This makes the model suitable for integration with clinical dashboards, where T, I, F, and Ψ can be displayed together.

A third implication is methodological. The model demonstrates that source-level evidence design is as important as the final aggregation rule. If evidence sources are clinically coherent, then indeterminacy has interpretable meaning. If sources are arbitrary, I becomes only a numerical artifact. Future neutrosophic healthcare studies should therefore report not only the operator used, but also the clinical rationale for each membership transformation.

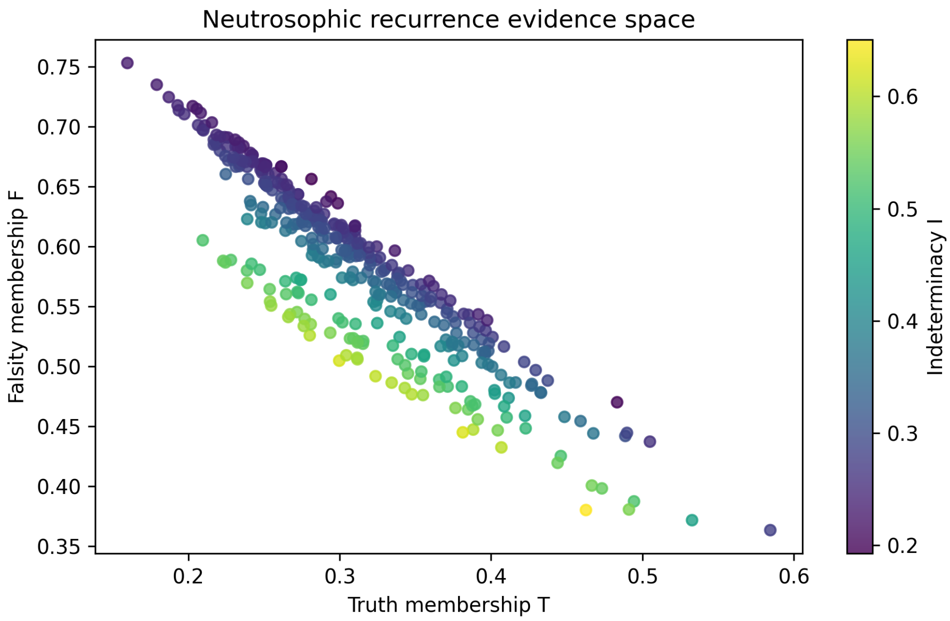

Figure 1: Neutrosophic evidence space showing truth, falsity, and indeterminacy memberships.

Figure 1 shows the geometry behind the proposed method. Patients near the diagonal $T \approx F$ form a natural ambiguity band, while the color gradient shows that this band is not identical to indeterminacy. Some patients have balanced truth and falsity with low conflict, whereas others are indeterminate because sources disagree sharply.

DOI: https://doi.org/10.54216/NIF.050201
Received: March 01, 2025 Accepted: June 28, 2025

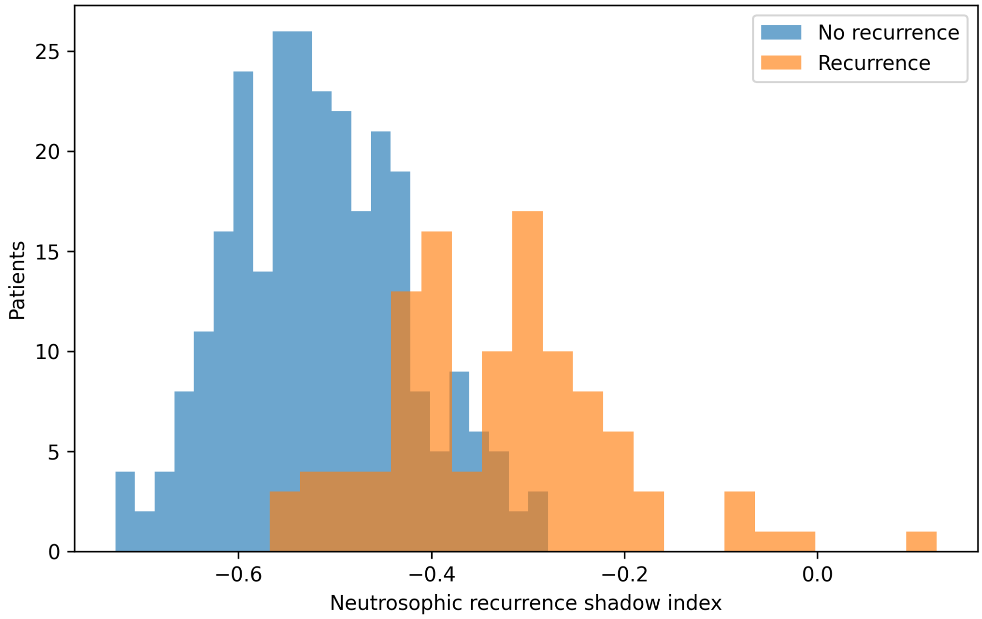

Figure 2: Distribution of the recurrence shadow index for recurrent and non-recurrent cases.

Figure 2 indicates that the recurrence-shadow index separates many recurrent and non-recurrent cases while keeping an overlapping region. The overlap is not treated as failure; it is the clinical shadow where intensified observation may be preferable to a definitive automated decision.

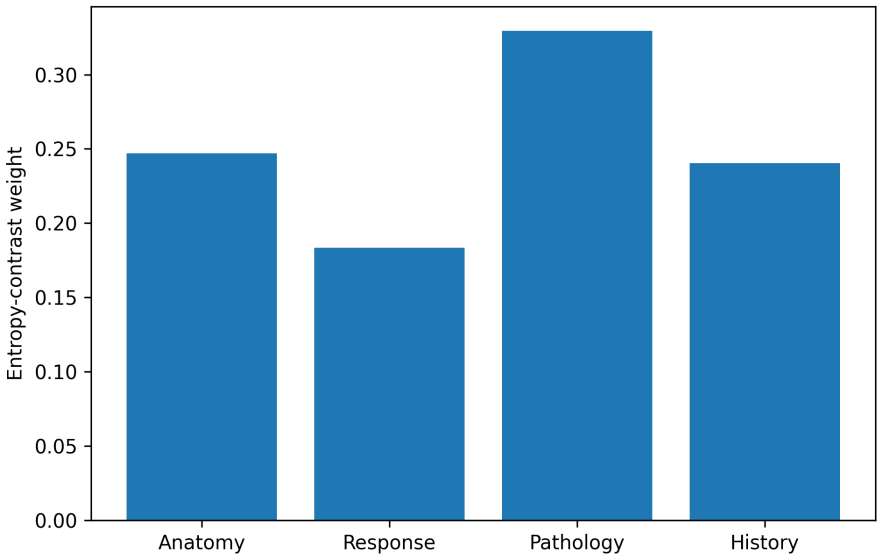

Figure 3: Entropy-contrast evidence weights used in truth-membership fusion.

The weights in Figure 3 are derived from evidence entropy rather than manually selected importance. This improves reproducibility because the same rule can be applied to another clinical cohort, while clinical interpretation remains possible through the source labels.

DOI: https://doi.org/10.54216/NIF.050201
Received: March 01, 2025 Accepted: June 28, 2025

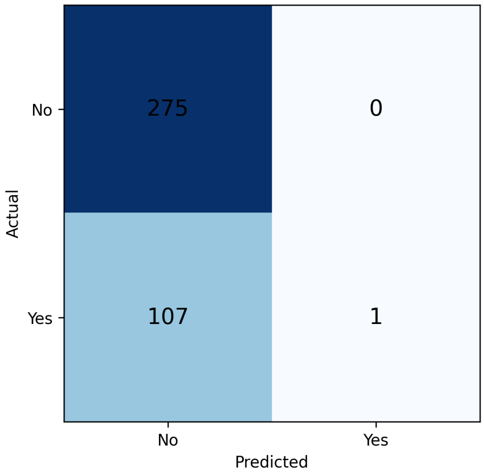

Figure 4: Confusion matrix of recurrence-shadow classification at the default action threshold.

Figure 4 provides the classification view of the model. The confusion matrix should be read together with the shadow-zone analysis rather than alone, because a clinically cautious model can intentionally place difficult patients in the boundary region instead of claiming full certainty.

7. Clinical Validation and Research Directions

The recurrence-shadow model should be interpreted as a decision-support layer that complements clinical follow-up rather than replacing physician judgment. Its strongest value appears in cases where conventional risk scores are insufficiently expressive. For example, a patient with low structural disease but indeterminate response should not be collapsed into the same state as a patient with low structural disease and excellent response. The former has a higher I component and therefore belongs closer to the observation shadow, whereas the latter has a larger falsity membership against recurrence. This distinction is the central healthcare contribution of the model.

Table 9: Recommended validation questions for future neutrosophic healthcare studies.

Validation question	Expected evidence
Does I correspond to clinically ambiguous cases?	Compare high-I patients with expert notes, repeated imaging, or indeterminate biochemical response.
Does T increase with recurrence-supporting findings?	Test monotonic behavior across T, N, M, risk, and response categories.
Does F identify genuinely stable cases?	Check whether high-F patients remain non-recurrent under longitudinal follow-up.
Are thresholds transferable?	Validate Ψ cutoffs across external institutions and different follow-up policies.
Is the model interpretable to clinicians?	Evaluate whether clinicians can explain alerts using T, I, F, and source-level evidence.

Table 9 gives practical validation questions. These questions are mathematically linked to the model because each question corresponds to one membership function or one decision threshold. A model that achieves high accuracy but fails these validation checks would not be a reliable neutrosophic healthcare system because the memberships would lack clinical semantics.

DOI: https://doi.org/10.54216/NIF.050201

Received: March 01, 2025 Accepted: June 28, 2025

Table 10: Potential extensions of the proposed recurrence-shadow model.

Extension	Mathematical addition	Healthcare value
Temporal follow-up	$\mathcal{N}_i(t) = \langle T_i(t), I_i(t), F_i(t) \rangle$ with dynamic aggregation	Represents changing recurrence evidence over visits.
Laboratory fusion	Add thyroglobulin and antibody trajectories as evidence sources	Links categorical risk with biochemical monitoring.
Report mining	Extract pathology features using NLP before neutrosophic mapping	Converts unstructured evidence into traceable memberships.
Federated calibration	Estimate thresholds without sharing patient-level records	Supports multi-center validation with privacy protection.
Expert-in-the-loop tuning	Optimize α, β, ρ under clinician feedback	Aligns uncertainty policy with medical practice.

Table 10 shows that the proposed model can be extended without changing its central logic. The important requirement is to preserve the independence of T, I, and F. For instance, temporal follow-up should not only update a probability; it should update whether the evidence is becoming more truthful, more false, or more indeterminate over time.

Table 11: Failure modes and mitigation strategies.

Failure mode	Mathematical symptom	Suggested mitigation
Over-alerting	High predicted recurrence rate under low α	Increase indeterminacy penalty or require minimum T.
False reassurance	High F despite incomplete response	Increase ρ so indeterminacy discounts falsity more strongly.
Unstable thresholds	Large variation in Ψ cutoffs across folds	Use nested calibration and external validation.
Weak clinical semantics	Memberships not aligned with expert judgment	Redesign source mapping with clinician review.
Hidden subgroup bias	Different error rates across gender or age strata	Report subgroup-level T, I, F, and performance.

Table 11 is included because healthcare decision support must discuss failure behavior explicitly. Each failure mode corresponds to an adjustable mathematical component. This makes the model auditable: a poor outcome can be traced to a threshold, a membership transformation, a source weight, or an indeterminacy rule.

Table 12: Minimum reporting items for publication of neutrosophic clinical models.

Reporting item	Required description
Membership construction	Define how each clinical variable contributes to T, I, and F.
Source weighting	State whether weights are expert-defined, entropy-derived, learned, or hybrid.
Indeterminacy semantics	Explain whether I measures missingness, contradiction, temporal uncertainty, or expert disagreement.
Threshold policy	Report how decision zones are selected and how sensitive they are.
Clinical safeguards	Describe how the model avoids automatic treatment recommendations and supports clinician review.
Reproducibility	Provide code, feature encodings, parameter values, and data-access instructions.

Table 12 is proposed as a publication checklist for future neutrosophic healthcare applications. Many papers introduce new operators but provide insufficient detail about membership construction. In clinical settings, this is a serious limitation because the medical meaning of I must be transparent. The checklist therefore encourages researchers to report the semantics of each membership component, not only the final performance metrics.

DOI: https://doi.org/10.54216/NIF.050201
Received: March 01, 2025 Accepted: June 28, 2025

8. Broader Discussion

The broader implication of this work is that neutrosophic logic can support a different style of medical AI. Instead of forcing every case into a binary label or a single probability, the proposed approach represents the patient as a clinical evidence state. This is closer to real follow-up practice, where a clinician often asks whether a patient is clearly stable, clearly alarming, or unresolved. The recurrence shadow is the mathematical expression of this unresolved middle state.

The proposed model also helps distinguish uncertainty from risk. A high-risk patient may have low indeterminacy if all evidence sources agree, while a moderate-risk patient may have high indeterminacy if response, pathology, and anatomical evidence conflict. This distinction is clinically important because high risk calls for action, whereas high indeterminacy calls for clarification. In a conventional scoring system these two situations may be confused.

Finally, the paper highlights the need for disease-specific neutrosophic design. Generic aggregation operators are valuable, but healthcare applications require careful mapping from clinical variables to memberships. In differentiated thyroid cancer follow-up, response status and anatomical burden have different meanings from demographic pressure. Treating them as interchangeable criteria would weaken the model. The proposed source-wise design therefore provides a more clinically grounded way to use neutrosophic information fusion.

9. Discussion

The numerical findings show that a neutrosophic model can provide a clinically meaningful layer above conventional recurrence prediction. The most important result is not simply that the proposed method achieves competitive classification behavior, but that it separates three distinct concepts: recurrence-supporting truth, recurrence-refuting falsity, and unresolved indeterminacy. In clinical follow-up, this distinction is essential. A patient with high truth and low indeterminacy can be managed as a high-alert case, whereas a patient with moderate truth and high indeterminacy may require additional tests, repeated imaging, or closer observation before escalation.

The entropy-contrast weighting rule also has practical implications. It avoids assigning equal influence to all evidence sources. In recurrence follow-up, response status and anatomical burden are often more informative than background history alone. However, history still contributes to the recurrence-shadow score by shifting contextual pressure. This layered behavior is consistent with clinical reasoning: prior smoking or radiotherapy history may not determine recurrence alone, but it can strengthen caution when combined with pathological or response-based signals.

The ablation analysis confirms that the model is not merely a relabeled classifier. When response evidence is removed, both truth-membership discrimination and indeterminacy modeling are weakened. This occurs because response variables affect not only the average recurrence evidence but also the contradiction pattern. A low anatomical burden with an incomplete response is exactly the type of case where a neutrosophic model is useful. Classical classifiers may still predict a label, but they do not express why the case is unresolved.

From a methodological perspective, the recurrence shadow index provides a compact but interpretable decision score. The boundedness and monotonicity properties ensure that stronger truth evidence does not reduce the action score when indeterminacy is fixed. At the same time, the indeterminacy penalty prevents the model from overstating confidence under contradictory evidence. This balance makes the method suitable for decision support where false reassurance and excessive alarm are both clinically undesirable.

The proposed model has limitations. The ordinal encodings are clinically motivated but should be calibrated with expert panels or larger multi-center cohorts before deployment. The recurrence shadow thresholds were selected to demonstrate the mathematical behavior of the method and can be tuned according to institutional follow-up policy. Future work should connect the neutrosophic memberships to longitudinal laboratory values, ultrasound reports, and natural-language pathology descriptions. Another useful direction is to integrate temporal neutrosophic aggregation so that T, I, and F evolve over repeated follow-up visits rather than being computed from a single record.

10. Conclusion

This paper introduced a neutrosophic recurrence shadow mapping model for healthcare decision support in post-treatment differentiated thyroid cancer follow-up. The proposed method maps clinicopathologic evidence into truth, indeterminacy, and falsity memberships, fuses evidence using entropy-contrast weights, and produces a bounded recurrence-shadow index. The mathematical and numerical analysis demonstrates that indeterminacy can be treated as a clinically interpretable state rather than a nuisance term. The model is therefore suitable for applications where medical evidence is incomplete, inconsistent, or partially contradictory. Future research should validate the model prospectively, extend it to temporal follow-up sequences, and combine it with clinician-facing explanations for personalized surveillance planning.

References

[1] Ahmad, M. A.-S. and Haddad, J. (2024). An explainable ai model for predicting the recurrence of differentiated thyroid cancer. arXiv preprint arXiv:2410.10907.

DOI: https://doi.org/10.54216/NIF.050201
Received: March 01, 2025 Accepted: June 28, 2025

[2] Ali, S., Rahim, M., Bajri, S. A., Ahmad, S., Alharbi, R., and Khalifa, H. A. E.-W. (2024). alpha, beta, gamma-neutrosophic aggregation operators and their applications in the software site selection. *Heliyon*, 10(10):e31417.

[3] Borzooei, S., Briganti, G., Golparian, M., Lechien, J. R., and Tarokhian, A. (2024). Machine learning for risk stratification of thyroid cancer patients: a 15-year cohort study. *European Archives of Oto-Rhino-Laryngology*, 281:2095–2104.

[4] Borzooei, S. and Tarokhian, A. (2023). Differentiated thyroid cancer recurrence. UCI Machine Learning Repository.

[5] Clark, E., Price, S., Lucena, T., Haberlein, B., Wahbeh, A., and Seetan, R. (2024). Predictive analytics for thyroid cancer recurrence: A machine learning approach. *Knowledge*, 4(4):557–570.

[6] Farid, H. M. A. and Riaz, M. (2022). Single-valued neutrosophic einstein interactive aggregation operators with applications for material selection in engineering design: case study of cryogenic storage tank. *Complex & Intelligent Systems*.

[7] Farid, H. M. A. and Riaz, M. (2023). Single-valued neutrosophic dynamic aggregation operators for multi-period decision making. *Engineering Applications of Artificial Intelligence*, 123:106940.

[8] Habchi, Y., Himeur, Y., Kheddar, H., Boukabou, A., Atalla, S., Chouchane, A., Ouamane, A., and Mansoor, W. (2023). Ai in thyroid cancer diagnosis: Techniques, trends, and future directions. arXiv preprint arXiv:2308.13592.

[9] Haddad, R. I., Bischoff, L., Ball, D., et al. (2022). Thyroid carcinoma, version 2.2022, nccn clinical practice guidelines in oncology. *Journal of the National Comprehensive Cancer Network*, 20(8):925–951.

[10] Irvanizam, I. (2024). An improved rafsi method based on single-valued trapezoidal neutrosophic number for healthcare service quality evaluation. *Expert Systems with Applications*, page 123343.

[11] Latif, S. et al. (2024). Decision support system for single-valued neutrosophic information. *Mathematical Problems in Engineering*, page 4362151.

[12] Lee, T. F., Lee, S. H., Tseng, C. D., et al. (2023). Using machine learning algorithm to analyse the hypothyroidism complications caused by radiotherapy in patients with head and neck cancer. *Scientific Reports*, 13:19185.

[13] Liu, P. et al. (2022). Novel edas methodology based on single-valued neutrosophic aczel-alsina aggregation information. *Mathematical Problems in Engineering*, page 2394472.

[14] Loor-Torres, R., Wu, Y., Cabezas, E., Borras, M., Toro-Tobon, D., Duran, M., et al. (2024). Use of natural language processing to extract and classify papillary thyroid cancer features from surgical pathology reports. arXiv preprint arXiv:2406.00015.

DOI: https://doi.org/10.54216/NIF.050201
Received: March 01, 2025 Accepted: June 28, 2025