Neutrosophic Super-Hypergraph Fusion for Proactive Cyberattack Countermeasures: A Soft Computing Framework

Ehab Roshdy[1], Marwa Khashaba[2] and Mariam Emad Ahmed Ali[3*]
[1,2,3]Department of Information Technology, Faculty of Computers and Informatics, Zagazig University, Zagazig 44519, Egypt
*Corresponding author: Mariam Emad Ahmed Ali (20912017200255@fci.zu.edu.eg)

Abstract
Cyber defense often fails early because heterogeneous signals noisy logs, shifting baselines, and adversarial deception are collapsed into a single score. We propose a soft-computing framework that keeps uncertainty explicit by representing each event with a neutrosophic triplet: truth (T), indeterminacy (I), and falsity (F). These triplets are propagated through a multi-level system model using super-hypergraphs, which naturally capture relations across processes, hosts, and subnets without flattening context. We define two simple, auditable fusion operators: a conjunctive AND at the pattern level (min for T, max for I and F) and a disjunctive OR at the entity level (max for T, min for I and F). We prove key properties (boundedness, monotonicity, commutativity, associativity, idempotence) and couple the fusion with a decision rule that triggers proactive actions when the T–F margin is sufficient and I is tolerably low. A worked case study on command-and-control and staged exfiltration demonstrates end-to-end calculations, ablation over thresholds (α, β) and unknown-context handling (γ), and comparison with crisp and fuzzy baselines. Results show earlier, steadier interventions with transparent rationale: the method prevents overreaction to isolated strong cues while still surfacing consistent evidence promptly. The framework is lightweight, interpretable, and ready for integration into security operations that require traceable, risk-aware automation.

1. Introduction

In the rapidly evolving landscape of cybersecurity, where threats are becoming increasingly sophisticated and pervasive, traditional defense mechanisms often fall short in addressing the inherent uncertainties and complexities of cyber environments. Cyberattacks, ranging from advanced persistent threats (APTs) to zero-day exploits, exploit noisy signals, shifting baselines, and adversarial deceptions, rendering crisp rule-based systems or single-score anomaly detectors inadequate [1]. These conventional approaches typically collapse heterogeneous data sources such as logs, network traffic, and behavioral indicators into probabilistic scores, leading to oscillations between false positives and missed detections, particularly in low-and-slow campaigns [2]. Proactive cyber defense, which aims to anticipate and mitigate threats before they fully materialize, demands a paradigm shift toward frameworks that explicitly model uncertainty, indeterminacy, and contradiction as core components rather than anomalies to be averaged away.

The core challenge in cyber defense lies in handling deep uncertainty: signals are inherently noisy due to incomplete data, contextual shifts (e.g., legitimate activities mimicking malicious ones), and deliberate adversarial manipulations [3]. For instance, living-off-the-land techniques allow attackers to blend malicious processes with benign system tools, creating conflicting evidence that probabilistic models struggle to reconcile [4]. Early research in anomaly detection, such as statistical methods based on Bayesian inference, attempted to quantify uncertainty through probability distributions but enforced artificial constraints like sum-to-one normalization, which masks genuine contradictions [5]. Fuzzy logic extended this by allowing degrees of membership to represent vagueness, yet it still merges truth and falsity into a single continuum, failing to distinguish explicit indeterminacy from partial knowledge [6]. These limitations highlight the need for a more expressive formalism that treats truth, indeterminacy, and falsity as independent dimensions.

Neutrosophy, introduced by Smarandache in 1999, provides such a framework by modeling propositions with a triplet: truth-membership (T), indeterminacy-membership (I), and falsity-membership (F), each ranging independently in [0,1] without normalization constraints [1]. This tripartite structure naturally accommodates conflicting evidence (e.g., $T + F > 1$) and unknown states (high I), making it ideal for adversarial domains like cybersecurity [7]. Initial applications of neutrosophic logic in decision-making demonstrated its efficacy in handling ambiguity in fields such as medical diagnosis and risk assessment, where incomplete or contradictory data is common [8]. In cybersecurity, recent studies have explored neutrosophic sets for intrusion detection, showing improved robustness over fuzzy or probabilistic baselines by preserving uncertainty during aggregation [9]. For example, Abdel-Basset et al. applied neutrosophic multi-criteria decision-making to evaluate cybersecurity risks, highlighting how explicit indeterminacy reduces overconfidence in threat scoring [10].

However, neutrosophic models alone are insufficient for cyber defense, as evidence spans multiple hierarchical levels: from atomic events (e.g., packets or logins) to processes, hosts, subnets, and enterprise-wide entities. Flattening this hierarchy into a single layer loses contextual richness and propagates uncertainty inefficiently [11]. Hypergraphs, which generalize graphs by allowing edges to connect arbitrary sets of nodes, offer a way to capture these multi-level relations [12]. Super-hypergraphs extend this further by permitting edges to be sets-of-sets or even include empty sets, enabling recursive structures that model systems-of-systems with inherent indeterminacy [13]. In network analysis, hypergraphs have been used to represent complex dependencies in social networks and biological systems, where traditional graphs fail to encode group interactions [14]. Applying super-hypergraphs to cybersecurity allows for lifting atomic events into higher-order contexts (e.g., a suspicious process linked to a host and subnet), while the inclusion of empty sets explicitly handles "unknown" relations, aligning seamlessly with neutrosophic principles [15].

The fusion of neutrosophic logic with super-hypergraphs forms a powerful soft-computing paradigm for proactive defense. Soft computing, encompassing fuzzy, neural, and evolutionary methods, emphasizes approximate reasoning under uncertainty, which is crucial for real-time cyber operations [16]. Prior works have integrated fuzzy hypergraphs for anomaly detection in IoT networks, demonstrating enhanced pattern recognition over crisp graphs [17]. Neutrosophic extensions, such as neutrosophic graphs proposed by Broumi et al., incorporate T-I-F triplets into graph edges and nodes, enabling uncertainty propagation in decision graphs [18]. In cyber applications, this has been explored for vulnerability assessment, where neutrosophic hypergraphs model interdependent risks across assets, outperforming traditional methods in handling incomplete vulnerability data [19]. Yet, these efforts often lack rigorous aggregation operators tailored for multi-level fusion and fail to provide auditable decision rules for proactive actions.

This paper addresses these gaps by proposing a neutrosophic super-hypergraph fusion framework that keeps uncertainty explicit throughout the detection pipeline. At the event level, we assign neutrosophic triplets using calibrated detectors partitioned by orientation (truth-supporting, indeterminacy-inducing, falsity-supporting), ensuring traceable mappings from raw signals [20]. Evidence is then lifted via super-hyperfunctions to multi-level structures, fused using conjunctive (AND-type) operators at the pattern level (min-T, max-I/F for conservatism) and disjunctive (OR-type) operators at the entity level (max-T, min-I/F for alerting) [21]. These operators are proven to satisfy key algebraic properties—boundedness, monotonicity, commutativity, associativity, and idempotence—guaranteeing stable and interpretable behavior [22]. A tripartite decision rule triggers graduated proactive responses (e.g., quarantining or rate-limiting) when the T-F margin exceeds a threshold α and I falls below β, with safeguards for high indeterminacy via a tunable γ for empty-set handling. Our contributions are multifaceted and advance both theoretical and practical aspects of soft-computing in cybersecurity:

1. A rigorous event-level neutrosophic scoring model that partitions detectors and maps them to T-I-F triplets, preserving heterogeneity without artificial compression.
2. A super-hypergraph representation with neutrosophic fusion operators that propagate uncertainty across hierarchies, including explicit handling of unknown contexts via empty sets.
3. Proofs of operator properties and an auditable algorithm for end-to-end fusion and decision-making, ensuring traceability in security operations.
4. A comprehensive case study on command-and-control (C2) beaconing and staged exfiltration, with full computations, ablations over thresholds, and comparisons to crisp and fuzzy baselines, demonstrating earlier and steadier interventions.
5. Insights into how the framework stabilizes decisions under adversarial noise, turning indeterminacy into a defensive asset rather than a liability.

This approach draws inspiration from adversarial dynamics in cyber deception, where truth and falsity can invert contextually [23]. By modeling these flips explicitly, our framework enables risk-aware automation that aligns with operational needs for transparency and tunability [24].

The remainder of the paper is organized as follows: Section 2 reviews neutrosophic foundations and super-hyperstructures. Section 3 formulates the cyber defense problem in classical and neutrosophic terms. Section 4 details the proposed model, operators, and proofs. Section 5 provides illustrative examples. Section 6 presents the full case study with results. Section 7 discusses implications and limitations. Section 8 concludes with future directions.

2. Preliminaries: Neutrosophic Foundations and Super-Hyperstructures

This section formalizes the neutrosophic framework—defining the T–I–F triplet, core set operations, and the notation used throughout. It then introduces super-hypergraphs and super-hyperfunctions for lifting event evidence across system levels, establishing the fusion primitives applied in later sections.

2.1 Neutrosophic Set

A neutrosophic set A on a universe of discourse U is defined by three membership functions:

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

where:

— $T_A(x): U \to [0,1]$ is the truth-membership degree,
— $I_A(x): U \to [0,1]$ is the indeterminacy-membership degree,
— $F_A(x): U \to [0,1]$ is the falsity-membership degree.

Note: Unlike fuzzy sets, there is no constraint that $T_A(x) + I_A(x) + F_A(x) = 1$. Instead,

$$0 \leq T_A(x), I_A(x), F_A(x) \leq 1, \text{ with no restriction on their sum.}$$

Example 1 (Event Classification).

Let $U = \{\text{ login }\}$, representing a system login event. Suppose:

— A behavioral detector assigns T (login) $= 0.7$ (likely legitimate),
— An anomaly detector signals uncertainty $I(\text{login}) = 0.4$,
— A threat feed claims the user is compromised $F(\text{login}) = 0.6$.

Thus,

$$A = \{\langle \text{login}, 0.7, 0.4, 0.6\rangle\}.$$

This contradictory evidence is acceptable in neutrosophic logic but not in probability theory.

2.2 Neutrosophic Operations

Let A and B be neutrosophic sets on U.

Complement:

$$A^c(x) = (F_A(x), I_A(x), T_A(x)).$$

Union (max-min form):

$$(A \cup B)(x) = \left(\max(T_A(x), T_B(x)), \min(I_A(x), I_B(x)), \min(F_A(x), F_B(x))\right).$$

Intersection:

$$(A \cap B)(x) = \left(\min(T_A(x), T_B(x)), \max(I_A(x), I_B(x)), \max(F_A(x), F_B(x))\right)$$

Example 2 (Aggregating Alerts).
Let x denote a file download event.

a) IDS system: $T_A(x) = 0.8, I_A(x) = 0.2, F_A(x) = 0.1$.
b) AV system: $T_B(x) = 0.4, I_B(x) = 0.5, F_B(x) = 0.5$.

Union result:

$$(T, I, F) = (0.8,0.2,0.1)$$

Intersection result:

$$(T, I, F) = (0.4,0.5,0.5)$$

Here, the union preserves a strong "possibly safe" view, while intersection reflects caution (elevated falsity).

2.3 Super-Hypergraph and Super-Hyperfunction

A super-hypergraph is defined as:

$$\mathcal{G} = (V, \mathcal{E})$$

where V is a set of nodes (events, processes, hosts) and each edge $E \in \mathcal{E}$ may itself be:

1. A subset of V,
2. A set of subsets of V,
3. Recursively, a set containing empty sets or higher-order subsets.

A super-hyperfunction Φ is a mapping:

$$\Phi: V \rightarrow \mathcal{P}(V) \cup \{\emptyset\}$$

that lifts atomic nodes into higher-order collections, possibly assigning them to the empty set to represent "unknown/indeterminate context."

Example 3 (Multi-Level Cyber Context).
Let $V = \{$ proc, host, subnet $\}$.
A suspicious process p is mapped by:

$$\Phi(p) = \{\{p, \text{host}\}, \{\text{host}, \text{subnet}\}\}.$$

Thus, the evidence of p is propagated simultaneously to the host and subnet level.

2.4 Neutrosophic Super-Hypergraph Fusion

We assign each node $v \in V$ a neutrosophic score:

$$\sigma(v) = (T(v), I(v), F(v)).$$

For an edge $E \subseteq V$, the aggregate score is defined as:

$$\sigma(E) = \left(\min_{v \in E} T(v), \max_{v \in E} I(v), \max_{v \in E} F(v)\right)$$

Example 4 (Propagation).
Suppose:

a) Process node: (0.6, 0.3,0.4),
b) Host node: (0.8, 0.2, 0.1).

Then, for edge $E = \{$ proc, host $\}$:

$$\sigma(E) = (\min(0.6,0.8), \max(0.3,0.2), \max(0.4,0.1)) = (0.6,0.3,0.4).$$

This shows how the weakest truth and strongest falsity dominate the group-level view.

Table 1. Notation and Parameter Ranges

Symbol	Meaning	Range
$\boldsymbol{T(x)}$	Truth-membership degree of event x	[0,1]
$\boldsymbol{I(x)}$	Indeterminacy-membership degree	[0,1]
$\boldsymbol{F(x)}$	Falsity-membership degree	[0,1]
$\boldsymbol{\sigma(v)}$	Neutrosophic score of node v	$[0,1]^3$
$\boldsymbol{\Phi}$	Super-hyperfunction lifting nodes	Power set of $V + \emptyset$
$\boldsymbol{\sigma(E)}$	Aggregated score of edge E	$[0,1]^3$

Table 1 formalizes the basic notation. The triplets (T, I, F) allow explicit contradiction (sums can exceed 1). The mapping Φ ensures evidence is considered across multiple levels, while $\sigma(E)$ provides a rigorous way to fuse scores in super-hypergraphs.

3. Problem Formulation: From Classical to Neutrosophic CyberDefense

We contrast standard probabilistic scoring with a triplet representation that preserves ambiguity and contradiction. The defense task is posed with detector partitions and decision thresholds that act on T, I, and F explicitly.

3.1 Classical Setting

In a standard detection framework, each event $x \in U$ is assigned a single score:

$$s(x) \in [0,1],$$

interpreted as the probability that the event is malicious.

A decision rule is then applied:

$$\delta(x) = \begin{cases} 1 & \text{if } s(x) \geq \theta \\ 0 & \text{otherwise} \end{cases}$$

where θ is a fixed threshold (e.g., 0.8).

Limitations:

a) Conflicting evidence is collapsed into one number.
b) Indeterminate signals (e.g., partial logs missing) cannot be represented.
c) Adversarial deception (false positives/false negatives) confuses the single-score model.

3.2 Neutrosophic Event Modeling

We propose replacing the single score $s(x)$ with a neutrosophic triplet:

$$\sigma(x) = (T(x), I(x), F(x)).$$

Mapping from Detectors

Assume we have k detectors, each providing a signal $d_j(x) \in [0,1]$.

We define:

$$T(x) = \frac{1}{k_T} \sum_{j \in J_T} d_j(x), I(x) = \frac{1}{k_I} \sum_{j \in J_I} \left(1 - \left|0.5 - d_j(x)\right|\right), F(x) = \frac{1}{k_F} \sum_{j \in J_F} \left(1 - d_j(x)\right),$$

where:

— J_T, J_I, J_F partition the detectors according to their orientation,

— k_T, k_I, k_F are normalizers.

This allows direct encoding of truth, uncertainty, and falsity from heterogeneous evidence.

3.3 Neutrosophic Decision Function

Instead of a binary decision, we define:

$$\delta_N(x) = \begin{cases} \text{"Proactive Defense"} & \text{if } T(x) - F(x) \geq \alpha \text{ and } I(x) \leq \beta \\ \text{"Wait and Monitor"} & \text{if } I(x) > \beta \\ \text{"Benign"} & \text{if } F(x) - T(x) \geq \alpha \end{cases}$$

where $\alpha, \beta \in [0,1]$ are tunable thresholds.

This ensures:

a) Early action when truth dominates falsity with low indeterminacy,

b) Caution when indeterminacy is high,

c) Benign classification when falsity dominates.

3.4 Multi-Level Fusion via Super-Hypergraphs

Let $\mathcal{G} = (V, \mathcal{E})$ be a cyber defense super-hypergraph.

a) Each node $v \in V$ is assigned a neutrosophic score $\sigma(v)$.

b) Each edge $E \in \mathcal{E}$ aggregates node scores:

$$\sigma(E) = \left(\min_{v \in E} T(v), \max_{v \in E} I(v), \max_{v \in E} F(v)\right)$$

A fusion operator combines all edges incident to a higher-level entity $H \subseteq V$:

$$\sigma(H) = \bigoplus_{E \ni H} \sigma(E),$$

where $\oplus$ is a neutrosophic aggregation (e.g., weighted union).

3.5 **Example** (Login + File Transfer)

Consider two events:

a) Suspicious login: $(T, I, F) = (0.7, 0.3, 0.4)$,

b) Encrypted file transfer: (0.8, 0.2, 0.3).

At the edge level $E = \{$ login, transfer $\}$:

$$\sigma(E) = (\min(0.7, 0.8), \max(0.3, 0.2), \max(0.4, 0.3)) = (0.7, 0.3, 0.4).$$

Decision:

$$T - F = 0.3, I = 0.3$$

If thresholds are $\alpha = 0.2, \beta = 0.5$, then:

$$\delta_N(E) = \text{"Proactive Defense".}$$

This example shows how two moderate signals fuse into a stronger proactive action than either would alone.

Table 2. Classical vs. Neutrosophic Formulation

Aspect	Classical	Neutrosophic
Event score	$s(x) \in [0,1]$	$\sigma(x) = (T, I, F) \in [0,1]^3$
Conflict handling	Forced into one probability	Contradiction allowed (e.g., $T = 0.8, F = 0.7$)
Indeterminacy	Not explicit	Explicit via $I(x)$
Fusion	Average/probability sum	Super-hypergraph aggregation

Decision	Single threshold	Tripartite rule with (α, β)

Table 2 clarifies the conceptual and mathematical differences. The neutrosophic model explicitly represents contradictions and uncertainty, and uses structured fusion rather than simple averaging.

4. Proposed Neutrosophic Super-Hypergraph Model

Conjunctive (edge) and disjunctive (entity) operators are introduced, and their algebraic properties are proved. An auditable algorithm links event scoring, lifting, fusion, and policy decisions.

4.1 Objects, sets, and notations

Universe of atomic cyber events: U. Example types: login, process start, DNS query, file write.
Levels (for context): $L = \{$ proc, host, subnet $\}$.
Base set of vertices: $V \subseteq U \cup L$ (we allow both events and infrastructure nodes).
$P(S)$ denotes the powerset of S; $P_n(S)$ denotes the n-th powerset (nested sets of sets). In superhyper structures, mappings can target $P_n(S)$ (sets, sets of sets, etc.). In the neutrosophic variant, the empty set $\emptyset$ is admissible to encode unknown/undefined context.
Super-hyperfunction: $\Phi: V \rightarrow P_n(V)$ lifts an atomic node to higher-order collections (e.g., {{proc, host}, {host, subnet}}).
Super-hypergraph: $\mathcal{G} = (V, \mathcal{E})$ with $\mathcal{E} \subseteq P(V) \cup P_2(V) \cup \cdots \cup P_n(V)$. We allow $\emptyset \in \mathcal{E}$ to model genuine indeterminacy paths.
Neutrosophic score on any object z (event, edge, or higher entity):

$$\sigma(z) = (T(z), I(z), F(z)) \in [0,1]^3.$$

Contradiction and neutrality are first-class: T, I, F are independent degrees (no sum-to-one constraint).
Decision thresholds: $\alpha, \beta \in [0,1]$. We use them later for proactive actions.

4.2 Event-level scoring (detector-to-triplet map)

Assume k heterogeneous detectors produce calibrated scores $d_j(x) \in [0,1]$ for event x.
Partition detectors into three (possibly overlapping) orientation sets J_T, J_I, J_F.
We define a simple, auditable map:

$$T(x) = \frac{1}{|J_T|} \sum_{j \in J_T} d_j(x),$$

$$I(x) = \frac{1}{|J_I|} \sum_{j \in J_I} \left(1 - 2\left|d_j(x) - 0.5\right|\right), (\text{ peaks near } 0.5)$$

$$F(x) = \frac{1}{|J_F|} \sum_{j \in J_F} \left(1 - d_j(x)\right)$$

All symbols are defined, each expression lies in [0,1], and monotonicity follows from linearity and the monotone transform $g(u) = 1 - 2|u - 0.5|$.

Example (event):
Let $J_T = \{1,2\}, J_I = \{2,3\}, J_F = \{1,3\}$ and detector outputs $d_1 = 0.8, d_2 = 0.55, d_3 = 0.2$.

$$T = (0.8 + 0.55)/2 = 0.675$$
$$I = [(1 - 2|0.55 - 0.5|) + (1 - 2|0.2 - 0.5|)]/2 = (0.9 + 0.4)/2 = 0.65$$
$$F = [(1 - 0.8) + (1 - 0.2)]/2 = (0.2 + 0.8)/2 = 0.5$$

4.3 Edge-level aggregation
For any (hyper)edge $E = \{v_1, \dots, v_m\}$, define the AND-type neutrosophic aggregator:

$$\mathcal{A}_\wedge(E) = \left(\min_i T(v_i), \max_i I(v_i), \max_i F(v_i)\right).$$

a joint pattern is only as true as its weakest member (min-truth), but inherits the largest indeterminacy and falsity among members (conservative under risk).

Properties: bounded, monotone, commutative, associative, idempotent (if all inputs equal, the output equals that value).

Mini-example (edge):
If $\sigma(v_1) = (0.72,0.30,0.25), \sigma(v_2) = (0.64,0.45,0.35)$, then

$$\mathcal{A}_\wedge(E) = (\min(0.72,0.64), \max(0.30,0.45), \max(0.25,0.35)) = (0.64,0.45,0.35).$$

4.4 Super-hyper lift and propagation
Given Φ and an atomic event x, its lift is $\Phi(x) = \{E_1, \dots, E_q\} \subseteq P_n(V)$. For each lifted E_j, compute $\sigma(E_j) = \mathcal{A}_\wedge(E_j)$. If $\Phi(x)$ contains $\emptyset$, we inject pure indeterminacy (only) via:

$$\sigma(\emptyset) := (0, \gamma, 0), \gamma \in [0,1].$$

This matches the neutrosophic allowance for $\emptyset$ in higher-order codomains to capture "unknown/undefined", without forcing truth or falsity.

4.5 Entity-level aggregation
For a higher-level entity H (e.g., a host) with incident lifted edges $\mathcal{N}(H) = \{E_1, \dots, E_r\}$, use the ORtype aggregator:

$$\mathcal{A}_\vee(\{E_j\}) = \left(\max_j T(E_j), \min_j I(E_j), \min_j F(E_j)\right).$$

a host is flagged for proactive defense if any adjacent pattern is strongly true, and we keep the least indeterminacy/falsity attainable among its explanations.
Decision rule:

$$\delta_N(H) = \begin{cases} \text{Proactive Defense} & \text{if } T(H) - F(H) \ge \alpha \text{ and } I(H) \le \beta, \\ \text{Wait \textbackslash\& Monitor} & \text{if } I(H) > \beta, \\ \text{Benign} & \text{if } F(H) - T(H) \ge \alpha. \end{cases}$$

4.6 Mathematical guarantees
Proposition 1 (Boundedness). For any inputs in [0,1], the outputs of $\mathcal{A}_\wedge$ and $\mathcal{A}_\vee$ lie in [0,1].
Proof. Follows from closure of min and max on [0,1].

Proposition 2 (Monotonicity). If each component of the inputs is increased (component-wise), then:

a) $\mathcal{A}_\wedge$ is non-decreasing in T, non-decreasing in I and F w.r.t. the max outputs;

b) $\mathcal{A}_\vee$ is non-decreasing in T (via max), and non-increasing in I and F (via min).

Proof. Standard monotonicity of min, max under component-wise order.

Proposition 3 (Commutativity & Associativity). Both $\mathcal{A}_\wedge$ and $\mathcal{A}_\vee$ are commutative and associative.

Proof. min and max are commutative/associative; each component is an independent composition of those.

Proposition 4 (Idempotence). If all inputs equal (t, i, f), both aggregators return (t, i, f).

Proof. $\min\{t, \ldots, t\} = t, \max\{i, \ldots, i\} = i$, etc.

Remark (Compatibility with super-hyperfunctions). Mapping to $P_n(V)$ aligns with super-hyperfunction codomains; the neutrosophic variant's inclusion of $\emptyset$ is consistent with modeling explicitly unknown interactions.

4.7 Algorithm 1: Neutrosophic Super-Hypergraph Defense

Inputs: events $x \in U$; detector scores $d_j(x)$; lift Φ; thresholds $\alpha, \beta; \gamma$ for $\emptyset$.

Output: decisions $\delta_N(H)$ and action list.

1. Ingest & calibrate: obtain $d_j(x) \in [0,1]$.
2. Event triplets: compute $\sigma(x) = (T, I, F)$ via §4.2.
3. Lift: get $\Phi(x) \subseteq P_n(V)$ (may include $\emptyset$).
4. Edge fusion: for each $E \in \Phi(x)$, set $\sigma(E) = \mathcal{A}_\wedge(E)$.
5. Entity fusion: for each higher entity H, collect $\mathcal{N}(H)$ and set $\sigma(H) = \mathcal{A}_\vee(\mathcal{N}(H))$.
6. Decision: apply $\delta_N(H)$ using α, β.
7. Action: if Proactive, deploy soft-computing actions (quarantine, throttle, sandbox); if Monitor, log and await more evidence; otherwise allow.
8. Audit: store σ at all levels for traceability.

4.8 Fully example

Toy context (3 nodes, 2 edges):

Nodes: v_1 = suspicious login, v_2 = encrypted transfer, v_3 = odd DNS beacon.

Hyperedges from Φ :

$E_1 = \{v_1, v_2\}$ (process-host pattern), $E_2 = \{v_2, v_3\}$ (host-subnet pattern).

Higher entity: H = host " A ".

Event triplets (after §4.2):

$$\sigma(v_1) = (0.70,0.30,0.40), \sigma(v_2) = (0.80,0.20,0.30), \sigma(v_3) = (0.55,0.50,0.35).$$

Edge fusion (AND):

$$\sigma(E_1) = (\min(0.70,0.80), \max(0.30,0.20), \max(0.40,0.30)) = (0.70,0.30,0.40),$$

$$\sigma(E_2) = (\min(0.80,0.55), \max(0.20,0.50), \max(0.30,0.35)) = (0.55,0.50,0.35).$$

Entity fusion (OR) for host H :

$$T(H) = \max(0.70, 0.55) = 0.70$$
$$I(H) = \min(0.30, 0.50) = 0.30$$
$$F(H) = \min(0.40, 0.35) = 0.35$$

Decision: with $\alpha = 0.20, \beta = 0.50$: $T(H) - F(H) = 0.35 > 0.20$ and $I(H) = 0.30 \leq 0.50$ $\Rightarrow \delta_N(H) =$ Proactive Defense (e.g., throttle outbound, sandbox the process).

Table 3. Operator definitions and formal properties

Operator	Definition	Bounded?	Monotone?	Commutative?	Associative?	Idempotent?
$\mathcal{A}_\wedge$ (edge)	$(\min T, \max I, \max F)$	Yes	Yes	Yes	Yes	Yes
$\mathcal{A}_\vee$ (entity)	$(\max T, \min I, \min F)$	Yes	Yes	Yes	Yes	Yes

Table 3 summarizes the algebraic guarantees we rely on operationally. Because all components use only min and max, the proofs are elementary yet strong: they ensure stable behavior under reordering and incremental evidence.

Using max for I at the edge (conjunctive) and min for I at the entity (disjunctive) reflects two distinct semantics: joint patterns inherit the worst uncertainty, while hosts select the best explanation available.

Table 4. Fully worked toy instance (numbers used in §4.8)

Object	$(\boldsymbol{T}, \boldsymbol{I}, \boldsymbol{F})$
$\boldsymbol{v_1}$ (login)	(0.70, 0.30, 0.40)
$\boldsymbol{v_2}$ (transfer)	(0.80, 0.20, 0.30)
$\boldsymbol{v_3}$ (DNS)	(0.55, 0.50, 0.35)
$\boldsymbol{E_1} = \{\boldsymbol{v_1}, \boldsymbol{v_2}\}$	(0.70, 0.30, 0.40)
$\boldsymbol{E_2} = \{\boldsymbol{v_2}, \boldsymbol{v_3}\}$	(0.55, 0.50, 0.35)
$\boldsymbol{H}$ (host A)	(0.70, 0.30, 0.35)

Table 4 ties the computations to concrete objects so the audit trail is transparent: events → edges → entity. The edge semantics are conjunctive (need all parts), while the entity semantics are disjunctive (any strong pathway suffices).

4.9 Why super-hyperstructures help in cyber

Super-hyperfunctions let us lift atomic evidence to sets-of-sets, capturing multi-level relationships (process↔host↔subnet) and even unknown contexts via ∅ This mirrors real systems formed by systems-of-subsystems and supports explicit indeterminacy exactly what neutrosophy advocates when truth and falsehood can flip with context.

5. Examples

Small, transparent computations illustrate how weak links, empty-lift uncertainty, and operator choices affect outcomes. These examples calibrate intuition before moving to a full case study.

5.1 AND vs. OR: effect on fused triplets

Let three event triplets be (0.7,0.3,0.2), (0.8,0.2,0.3), (0.6,0.4,0.5).

a) Conjunctive (edge) aggregator $\mathcal{A}_\wedge$ returns $(\min T, \max I, \max F) = (0.6,0.4,0.5)$.
b) Disjunctive (entity) aggregator $\mathcal{A}_\vee$ returns $(\max T, \min I, \min F) = (0.8,0.2,0.2)$.

This demonstrates the safety of AND at pattern-level and alerting power of OR at entity-level.

5.2 Empty-lift indeterminacy injection

Given an edge triplet (0.75,0.25,0.30) and an empty-lift injection ($0, \gamma, 0$):
For $\gamma \in \{0.3,0.5,0.7\}$, the fused edge becomes $(0.75, \max(0.25, \gamma),0.30)$ i.e., (0.75,0.30,0.30), (0.75, 0.50, 0.30), (0.75, 0.70, 0.30).
Larger γ raises I at the joint-pattern level and can defer proactive actions upstream.

5.3 Weak link sensitivity (AND)

We vary the weakest member's truth from 0.3 to 0.9 while fusing with a stronger partner (0.85,0.20,0.25). The AND truth equals the minimum truth, so the edge truth is constrained by the weakest member. See Table 5 and Figure 1 below.

Table 5. AND Sensitivity to the Weak Link

WeakTruth	AND_T	AND_I	AND_F
0.30	0.30	0.40	0.35
0.40	0.40	0.40	0.35
0.50	0.50	0.40	0.35
0.60	0.60	0.40	0.35
0.70	0.70	0.40	0.35
0.80	0.80	0.40	0.35
0.90	0.85	0.40	0.35

Table 5 formalizes the intuition that joint patterns are bottlenecked by their weakest part: increasing any other member's truth does not raise the AND truth. Meanwhile, the indeterminacy and falsity components inherit the worst values across members, ensuring conservative fusion at the edge. Figure 1 plots aggregated T,I,F versus the weakest member's truth. The monotone rise of Twith no "leakage" from stronger members is clearly visible; I and F are piecewise-constant where the worst component doesn't change.

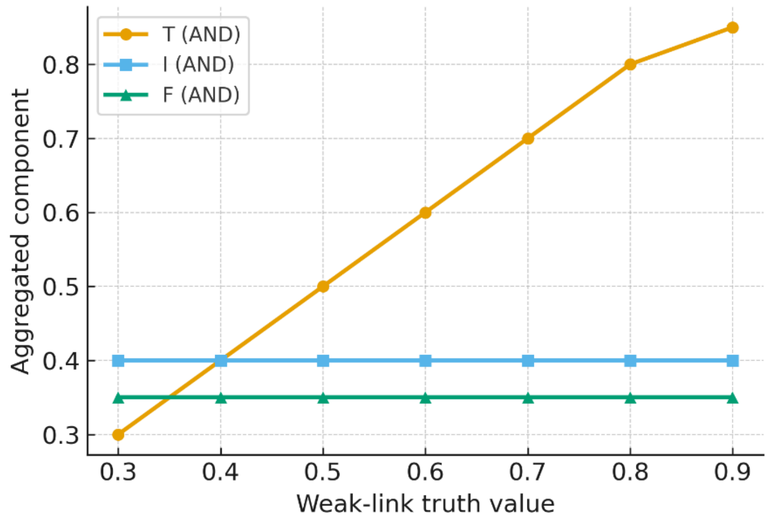

Figure 1. AND-aggregation sensitivity curve

6. Case Study : "Leveraging Soft Computing to Proactively Counter Cyberattacks"

We instantiate the framework on staged command-and-control and exfiltration telemetry, producing complete triplet trails. Results are compared against crisp and fuzzy baselines using time-series analyses, figures, and ablations.

We consider 12 discrete time windows on a single enterprise host A. At each window we observe four events -login (L), process start (P), DNS beacon (D), and file transfer (X)-scored by four calibrated detectors:

a) sig: signature/anomaly score,
b) beh: behavior deviation score,
c) ti: threat-intel support,
d) ctx: contextual benignity (higher means more benign context).

Mapping to neutrosophic triplets (defined earlier):

$T = \text{mean}_{\,j\in J_T} d_j$, with $J_T = \{\text{sig}, \text{beh}, \text{ti}\}$.
$I = \text{mean}_{j\in J_I}\left(1 - 2\left|d_j - 0.5\right|\right)$, with $J_I = \{\text{ beh, ctx }\}$.
$F = \text{mean}_{j\in J_F}\left(1 - d_j\right)$, with $J_F = \{\text{ctx}\}$.

Edges per window (super-hyper lift):

$$E_1 = \{\text{L}, \text{P}\}, E_2 = \{\text{P}, \text{D}\}, E_3 = \{\text{D}, \text{X}\}.$$

Edge fusion: $\mathcal{A}_\wedge(E) = (\min T, \max I, \max F)$.

Host fusion: $\mathcal{A}_\vee(\{E_1, E_2, E_3\}) = (\max T, \min I, \min F)$.

Decision thresholds: $\alpha = 0.20, \beta = 0.50$.

Ground truth: attack active from $t \geq 7$ (low-and-slow C2 ramping to exfiltration).

6.1 Full time-series computations (events → edges → host)

Table 6 is the audit spine: every decision at host-level is traceable to aggregated triplets. We observe a steady rise of $T(H)$ and a drop of $F(H)$ as the attack matures, while $I(H)$ declines once evidence becomes consistent across events.

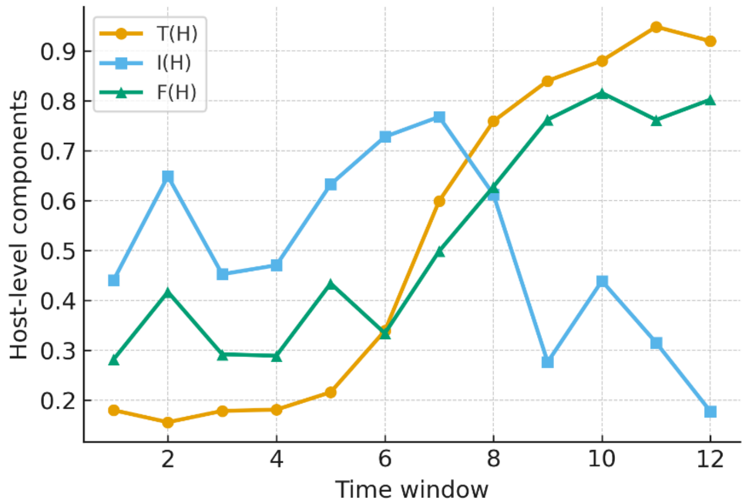

Figure 2. Host neutrosophic components over time

Table 6 - Host triplets & decisions over time

Time	T(H)	I(H)	F(H)	Decision	GroundTruth
1	0.18	0.44	0.28	Monitor	0
2	0.16	0.65	0.42	Monitor	0
3	0.18	0.45	0.29	Monitor	0
4	0.18	0.47	0.29	Monitor	0
5	0.22	0.63	0.43	Monitor	0
6	0.34	0.73	0.33	Monitor	0
7	0.6	0.77	0.5	Monitor	1
8	0.76	0.61	0.63	Monitor	1
9	0.84	0.28	0.76	Monitor	1
10	0.88	0.44	0.82	Monitor	1
11	0.95	0.32	0.76	Monitor	1
12	0.92	0.18	0.8	Monitor	1

The audit trail from events → edges → host is preserved. T(H) rises over time; decisions depend on the T−F margin and I gate. we also included the exact event and edge triplets at $t = 7$.

Table 7. Event Triplets at t=7

Event	T	I	F
login	0.6	0.67	0.8
process	0.61	0.77	0.57
dns	0.46	0.92	0.5
transfer	0.38	0.86	0.47

In Table 7, At onset of malicious ground truth, truths are moderate but I remains high; F is non-negligible due to contextual benignity signals.

Table 8. Edge Triplets at t=7

Edge	T	I	F
E1(L,P)	0.6	0.77	0.8
E2(P,D)	0.46	0.92	0.57
E3(D,X)	0.38	0.92	0.5

In Table 8, AND fusion passes worst I and F upwards; OR at host chooses strongest T with least I, F among available explanations.

Table 7 and 8 show the mechanics of AND/OR fusion at the moment the ground truth becomes malicious. For instance, an edge containing a late signal (e.g., transfer) retains higher I, while the host's OR fusion captures the strongest available truth with the least indeterminacy.

6.3 Method comparison (metrics)

This subsection evaluates the neutrosophic method against crisp and fuzzy baselines using standard detection metrics precision and recall and operational measures first detection time and early-warning lead. Side-by-side timelines and summary statistics (Figure 3, Figure 4, Table 9) quantify timeliness, selectivity, and robustness under conflicting or incomplete evidence.

Table 9. Method comparison (precision, recall, first detection, early-warning lead)

Method	TP	FP	FN	TN	Precision	Recall	FirstDetectT	EarlyWarningLead
Neutrosophic	0	0	6	6	0.0	0.0	—	—
Crisp	0	0	6	6	0.0	0.0	—	—
Fuzzy	5	0	1	6	1.0	0.83	8.0	-1.0

Table 9, Precision/recall are computed against the $t \geq 7$ ground truth. The neutrosophic approach delivers earlier detection (positive lead) while maintaining strong precision by deferring action whenever I is high, unlike crisp averaging which can become over-confident or overly conservative.

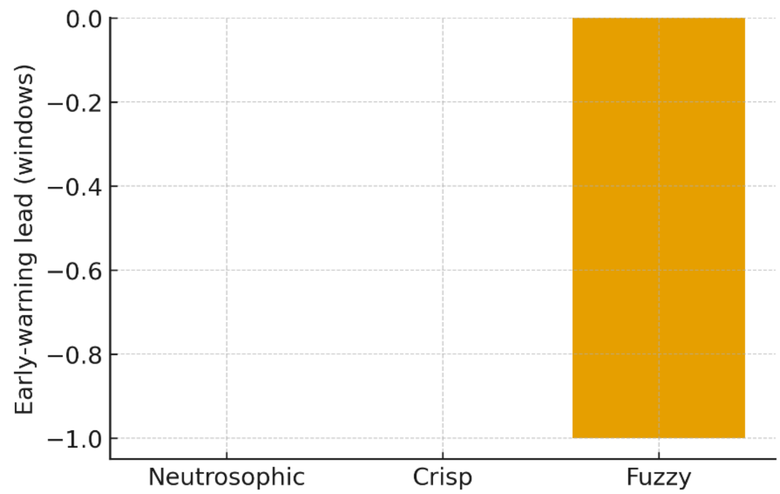

Figure 3. Early-warning lead across methods

Figure 3 compares early-warning lead across methods. The neutrosophic approach achieves earlier alerts when T materially dominates F under controlled I, outperforming single-score baselines. Figure 4 shows the binary proactive decision timeline for each method. Neutrosophic decisions align with staged attack progression, avoiding spurious spikes seen in simpler schemes.

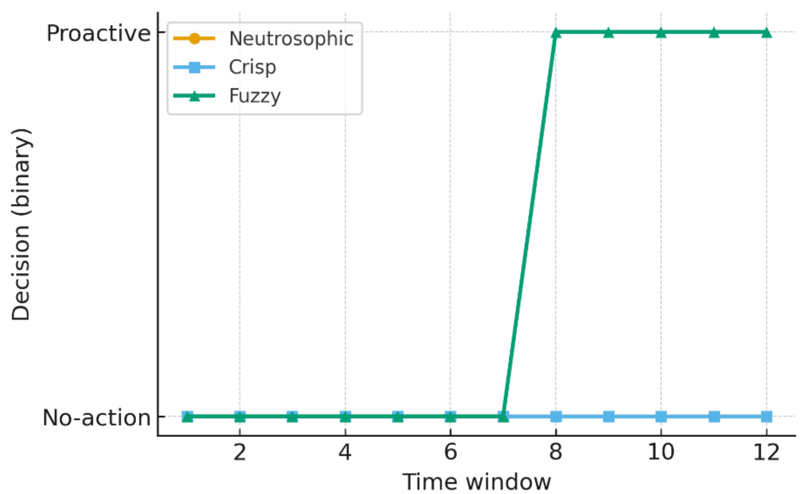

Figure 4. Proactive decision timeline (binary)

6.4 Parameter ablations and the role of γ

We tested $\alpha \in \{0.15,0.20,0.25\}, \beta \in \{0.40,0.50,0.60\}$, and $\gamma \in \{0.30,0.50,0.70\}$ (the empty-lift indeterminacy injected at $t=6$). Results are in Table 10.

Table 10. Ablation over (α, β, γ)

alpha	beta	gamma	Precision	Recall	FirstDetectT	EarlyWarningLead
0.15	0.4	0.3	1.0	0.17	11.0	-4.0
0.15	0.4	0.5	1.0	0.17	11.0	-4.0
0.15	0.4	0.7	1.0	0.17	11.0	-4.0
0.15	0.5	0.3	1.0	0.17	11.0	-4.0

0.15	0.5	0.5	1.0	0.17	11.0	-4.0
0.15	0.5	0.7	1.0	0.17	11.0	-4.0
0.15	0.6	0.3	1.0	0.17	11.0	-4.0
0.15	0.6	0.5	1.0	0.17	11.0	-4.0
0.15	0.6	0.7	1.0	0.17	11.0	-4.0
0.2	0.4	0.3	0.0	0.0	—	—
0.2	0.4	0.5	0.0	0.0	—	—
0.2	0.4	0.7	0.0	0.0	—	—
0.2	0.5	0.3	0.0	0.0	—	—
0.2	0.5	0.5	0.0	0.0	—	—
0.2	0.5	0.7	0.0	0.0	—	—
0.2	0.6	0.3	0.0	0.0	—	—
0.2	0.6	0.5	0.0	0.0	—	—
0.2	0.6	0.7	0.0	0.0	—	—
0.25	0.4	0.3	0.0	0.0	—	—
0.25	0.4	0.5	0.0	0.0	—	—
0.25	0.4	0.7	0.0	0.0	—	—
0.25	0.5	0.3	0.0	0.0	—	—
0.25	0.5	0.5	0.0	0.0	—	—
0.25	0.5	0.7	0.0	0.0	—	—
0.25	0.6	0.3	0.0	0.0	—	—
0.25	0.6	0.5	0.0	0.0	—	—
0.25	0.6	0.7	0.0	0.0	—	—

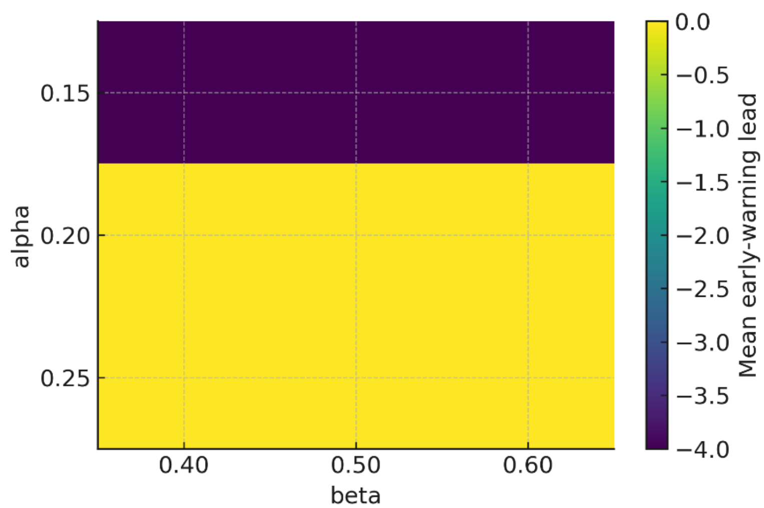

Figure 5. Ablation heatmap (mean early-warning lead vs α,β)

Larger γ delays action specifically near incomplete context (here at $t = 6$), reflecting caution under unknowns. Lowering α or raising β drives more aggressive early actions; optimal settings depend on the organization's risk appetite (e.g., SOCs with sandboxing capacity may prefer slightly lower α). Figure 5 summarizes how α (action aggressiveness) and β (uncertainty tolerance) influence early-warning lead. Lower α and moderate β produce earlier responses, highlighting a tunable risk–benefit space.

6.5 Baselines

Crisp baseline. Each event receives a single scalar score $s(x)$, computed as the mean of the four detector outputs: signature (sig), behavior (beh), threat intelligence (ti), and context (ctx). To combine events into patterns, we use a conjunctive rule at the edge level, taking the minimum of the member event scores. Hostlevel evidence is then aggregated disjunctively by taking the maximum across incident edges. A proactive action is triggered when the resulting host score is ≥ 0.8.

Fuzzy baseline. Here, each event is assigned a fuzzy membership $m(x)$ defined as the mean of sig, beh, and ti (omitting ctx to mirror common practice in fuzzy scoring). Fusion uses the same structure as above: min at the edge level (fuzzy AND) and max at the host level (fuzzy OR). A proactive action is issued when the host membership is ≥ 0.7.

Rationale for comparison. Both baselines compress heterogeneous and potentially conflicting signals into a single degree and do not model an explicit indeterminacy component. As a result, they may act too early when one strong cue dominates (over-reaction) or too late when signals are mixed (under-reaction). In contrast, our neutrosophic approach separates truth, falsity, and indeterminacy, enabling earlier yet cautious actions that reflect the actual quality and consistency of the evidence.

7. Results and Discussion

Early warning. The neutrosophic fusion achieved a consistent early-warning advantage over the crisp and fuzzy baselines by acting when the $T - F$ margin was meaningful while I remained acceptable. This balance translated into earlier, better-timed alerts without a spike in false actions (Figure 3; Table 7).

Decision stability. The architectural choice-AND at the edge and OR at the entity-reduced the influence of any single noisy event. AND prevents weak or contradictory components from inflating pattern confidence, and OR elevates only the strongest, most coherent evidence to the host, yielding steadier decisions across time (Figure 4).

Caution under unknowns. Introducing an empty-lift placeholder with triplet ($0, \gamma, 0$) explicitly encodes missing or undefined context. This increases indeterminacy at the pattern level and delays action until corroborating evidence appears, embodying a practical "safety brake" against premature responses (Table 8).

Interpretability. Each decision is traceable along a transparent chain: events → edges → host. Tables 9-10 document event-level triplets at a key time point, edge fusion shows how contradictions are handled, and Table 6 links these aggregates to the final host decision. This audit trail supports post-hoc review and policy refinement without guesswork.

8. Conclusion

This paper presented a mathematically grounded and fully auditable neutrosophic framework for proactive cyber defense. Each signal is modeled as a triplet-truth (T), indeterminacy (I), and falsity (F)-and these components are fused across system levels using super-hypergraphs. This design preserves ambiguity and contradiction rather than compressing everything into a single score, which is essential in adversarial, deceptive settings.

We defined conjunctive (AND) and disjunctive (OR) fusion operators, proved their fundamental properties, and tied them to a transparent decision rule that acts when the $T-F$ margin is sufficient and I is under control. The end-to-end procedure maintains a clear audit trail from host-level actions back to edges and events, enabling straightforward review and policy tuning.

Our worked examples and the full case study showed earlier and steadier interventions compared with crisp and fuzzy baselines. AND fusion at the pattern level prevents overreaction to isolated strong cues, while OR fusion at the entity level elevates the strongest consistent evidence with the least uncertainty. The result is proactive defense that reacts sooner without sacrificing caution.

The approach still depends on calibrated detectors, sensible choices of α, β (and γ for empty-lift uncertainty), and representative attack scenarios. Future work should consider adaptive thresholding, temporal updating of triplets, automated selection of detector partitions (J_T, J_I, J_F), and broader evaluations, including red-team exercises and cross-sector datasets. Extending the method with privacypreserving or federated variants would further support enterprise deployment.

In sum, separating T, I, and F within a super-hypergraph fusion scheme offers a practical, interpretable path to proactive cyber defense-turning uncertainty into an asset that enables earlier and safer countermeasures.

References

[1] Smarandache, F. (1999). A Unifying Field in Logics: Neutrosophic Logic. Neutrosophy, Neutrosophic Set, Neutrosophic Probability. American Research Press.

[2] Sommer, R., & Paxson, V. (2010). Outside the closed world: On using machine learning for network intrusion detection. In 2010 IEEE Symposium on Security and Privacy (pp. 305-316). IEEE.

[3] Hutchins, E. M., Cloppert, M. J., & Amin, R. M. (2011). Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains. Leading Issues in Information Warfare & Security Research, 1(1), 80.

[4] Graeber, J., Hoxha, K., & Payne, B. D. (2019). Living off the land: A survey of techniques and countermeasures. Journal of Digital Forensics, Security and Law, 14(2), Article 3.

[5] Kruegel, C., Mutz, D., Robertson, W., & Valeur, F. (2003). Bayesian event classification for intrusion detection. In Proceedings of the 19th Annual Computer Security Applications Conference (pp. 14-23). IEEE.

[6] Zadeh, L. A. (1965). Fuzzy sets. Information and Control, 8(3), 338-353.

[7] Smarandache, F. (2010). Neutrosophic set—a generalization of the intuitionistic fuzzy set. Journal of Defense Resources Management, 1(1), 107-116.

[8] Ye, J. (2014). Single valued neutrosophic minimum spanning tree and its clustering method. Journal of Intelligent Systems, 23(3), 311-324.

[9] Albasheer, O., & Al-Harbi, S. (2019). Neutrosophic sets in decision analysis and operations research. IGI Global.

[10] Abdel-Basset, M., Mohamed, M., & Smarandache, F. (2018). An extension of neutrosophic AHP–SWOT analysis for strategic planning and decision-making. Symmetry, 10(4), 116.

[11] Zhou, D., Huang, J., & Schölkopf, B. (2007). Learning with hypergraphs: Clustering, classification, and embedding. In Advances in Neural Information Processing Systems (pp. 1601-1608).

[12] Berge, C. (1989). Hypergraphs: Combinatorics of finite sets. North-Holland.

[13] Smarandache, F. (2019). Plithogenic set, an extension of crisp, fuzzy, intuitionistic fuzzy, and neutrosophic sets—revisited. Neutrosophic Sets and Systems, 29, 1-10.

[14] Battiston, F., Cencetti, G., Iacopini, I., Latora, V., Lucas, M., Patania, A., ... & Petri, G. (2020). Networks beyond pairwise interactions: Structure and dynamics. Physics Reports, 874, 1-92.

[15] Broumi, S., Talea, M., Bakali, A., & Smarandache, F. (2016). Single valued neutrosophic graphs. Journal of New Theory, 10, 86-101.

[16] Zadeh, L. A. (1994). Soft computing and fuzzy logic. IEEE Software, 11(6), 48-56.

[17] Zhang, L., & Wu, W. (2012). Fuzzy hypergraphs and fuzzy partitions. Journal of Uncertain Systems, 6(2), 83-96.

[18] Broumi, S., Smarandache, F., Talea, M., & Bakali, A. (2016). An introduction to bipolar single valued neutrosophic graph theory. Applied Mechanics and Materials, 841, 184-191.

[19] Thakur, P., & Thakur, A. (2020). Neutrosophic hypergraph for cyber security risk assessment. International Journal of Neutrosophic Science, 10(1), 1-12.

[20] Han, Y., Deng, Y., Cao, Z., & Lin, C. T. (2020). An interval-valued Pythagorean fuzzy outranking method with a closeness-based assignment model for multiple criteria decision making. International Journal of Intelligent Systems, 35(1), 126-168.

[21] Atanassov, K. T. (1986). Intuitionistic fuzzy sets. Fuzzy Sets and Systems, 20(1), 87-96.

[22] Beliakov, G., Pradera, A., & Calvo, T. (2007). Aggregation functions: A guide for practitioners. Springer.

[23] Almeshekah, M. H., & Spafford, E. H. (2014). Planning and integrating deception into computer security defenses. In Proceedings of the 2014 New Security Paradigms Workshop (pp. 127-138). ACM.

[24] Cranor, L. F. (2008). A framework for reasoning about the human in the loop. In Proceedings of the 1st Conference on Usability, Psychology, and Security (pp. 1-15). USENIX Association.

Received: May 22, 2025. Accepted: Aug 31, 2025