



Modeling Ambiguity in AI-Enhanced Learning: A Neutrosophic Approach to Stance Detection and Causal Evaluation

Oscar José Alejo Machado^{1,*}, Adriana María Estupiñán Sera¹, Maikel Y. Leyva Vazquez²,
Florentin Smarandache³

¹Instituto Superior Tecnológico de Investigación Científica e Innovación (ISTICI), Quito, Ecuador

²Bernardo O'Higgins University, Institutional Research Center, Chile

³Emeritus Professor of Mathematics at the University of New Mexico, Gallup, New Mexico, USA

Emails: oscar.alejo@istici.edu.ec; Sera.Adriana@gamil.com; Vazquez.Maikel@gamil.com;
smarand@unm.edu

Abstract

This work presents a neutrosophic stance detection model to bridge computational assessment and logic of indeterminacy in artificially intelligent (AI)-mediated learning and its outcomes. Utilizing the BART-large-MNLI model, a causal assessment was made of five hypotheses stemming from AI-supported learning between teacher-student relationships. These stances were then transformed into refined neutrosophic values (truth (T), partial support (P_S), indeterminacy (I), partial opposition (P_O) and falsity (F)). Ultimately, findings suggest that partial support is the most prevalent stance applied to any of the hypotheses, revealing that AI is, largely, a boon to education. However, this valence is tempered by indeterminacy among axes as well as stance magnitude. The largest partial support in rank order came from personalized education and access to AI tutors, while the most importance was given to opposition of relying on AI as support and replacement AI learning. Such findings confirm neutrosophic stance analysis and causal graph modeling as increasingly successful for applying measurable patterns to epistemically ambiguous fields. The neutrosophic causal graph integrates the above findings with a visualization of proposed dynamics between each vertex based on both quantitative patterns and epistemic uncertainty trends. The current research holds implications for educational theory, policy and instructional design integrity in 21st century learning. Uncertainty became a tangible concept; instead of devaluing AI in the classroom, it must be present as an enhancing supplemental tool, never replacement, for ethical considerations and equitable access. The potential for neutrosophic to transform apparent truths that are at times contradictory is confirmed through the human-machine interactive learning process, with subsequent suggestions for future research into AI-mediated education's causal relationships and decision-making potential.

Keywords: Neutrosophic; Stance detection; Causal analysis; AI in education; Uncertainty modeling

1. Introduction

AI (artificial intelligence) is an unavoidable part of education, meaning that the way teachers create activities and the way students engage with information will never be the same. From personalized tutoring suggestions to predictive learning analytics, AI aims to make learning increasingly personalized, responsive, and data-driven [1]. However, while it seems that way, the increasing body of research of AI's position on learning effectiveness has surprisingly mixed results. Some studies reveal academic improvements in immersion and outcomes while others show bias, reliance, or cognitive boredom [2]. These ideological divides signal a greater challenge to

theoretical and methodological coherence in a context that infers caution; our methods for evidence-based synthesis are too clear in a sector motivated by uncertainty and situationality.

In a perfect world, researchers would have a transparent and coherent understanding of the practical benefits AI has on—or fails to have on—learning. Evidence would be cumulative, incrementally providing guidance for practitioners and decision-makers, providing clarity without vagueness. In reality, however, much of the data exists—and operates—within binary parameters; findings claim that certain interventions are useful or not useful without recognizing the intricacies of personalized human learning to suggest effective, partially effective, ineffective, and anything in between [3]. Thus, we create an evidentiary ecosystem that validates categories of findings and not those based on something less-than-ideal. Biesta warns as much; the quest for “what works” fails to appreciate education's inherently value-based and situational character [4].

Such reductionism results in fragmentation both directly and indirectly. Directly, value assignments made from generalized findings about the effectiveness of AI tools extend too far from contextually bound situations in which findings are discovered. Indirectly, policy and funding favor certain technology over others based on prescriptive abilities instead of availability or pedagogical considerations. For example, AI-driven analytics might make academic outcomes more immediate in schools with comprehensive technology access—students who rely on AI systems to learn how to learn—but in areas lacking this access, it only exacerbates the digital divide [2]. The result? Schools continue to exist with inequitable access to the tremendous benefits AI can provide for education.

Efforts have been made in these fragments to create a new whole. While meta-analyses and systematic reviews have attempted to aggregate thousands of findings together for larger claims and comparisons, their interpretive frameworks still assume that findings speak with one voice [2,3]. In reality, studies appeal to different stances—partially positive or even ambivalent and contradictory—meaning there exists a plethora of epistemically valuable stances that go untapped in conventional theoretical claims. Without a means to quantify uncertainty or to seek appeal in partial support or objection, the literature exists with an up-and-down courtship campaign of hope and disillusionment without recourse.

This study operates under the thesis that the problem is epistemological and not just empirical. We need an epistemology of synthesis that reconstitutes the fragmented evidence based on a logic that supports contradiction, multi-level truths. Smarandache's neutrosophic logic [5] supplies such a paradoxical notion. It divides sectors into truth (T), indeterminacy (I), and falsity (F) as independent dimensions, meaning something can be true with uncertainty or even false at varying levels simultaneously. Such applications in cognitive modeling and decision-making demonstrate how complicated and uncertain realities can be formally represented [6–9]. Furthermore, neutrally advanced approaches to stance detection—the ability to locate whether a passage supports, opposes, or is neutral about a claim—show how detection has been anticipated [10].

Thus, where other studies fall short, this research implements an integrated Neutrosophic Stance Detection Framework for Causal Graph Representation where stance detection's linguistic abilities meld with neutrosophic epistemic possibilities to categorize evidence into five nuanced categories: support, partial support, neutrality, partial opposition, or opposition against a suspected causal claim (AI programs/platforms, feedback immediacy/quality/timeliness/non-delivery systems or the accessibility digital divide relative to learning outcomes). While other pre-existing systems like the Consensus Meter [11] can automatically classify stances based on probability or binary models, they fail to articulate uncertainty born out of contradictory evidence.

The gap filled by this study relates to the inadequacy of existing options for evidence synthesis which cannot represent nuanced positions of uncertainty in-between AI education research findings. The ideal situation—transparency and consistent reality through synthesized potentially conflicting findings—has yet to be realized. This gap calls for a method applicable not just to knowns but largely indeterminates as well.

Thus, this study endeavors two complementary goals:

- To report a neutrosophic stance detection model that can categorize levels of existence respectively within AI-related studies.
- To create a neutrosophic causal graph that visualizes the interconnectivity, or lack thereof, across consensually supported findings mixed with controversy.

Theoretical contributions benefit scholarly work as a methodological advancement for representing ambiguity/contradiction as a measurable construct rather than an unavoidable interpretive complication. Real-world contributions benefit policy makers as diagnostic maps highlighting where strong consensus exists—and where evidence remains disputed—critical in decision-making around AI's potentially overpowering implementation efforts. Practical contributions integrate an analytic framework that could change with changing technology itself.

The organization of this paper follows the CARS model; it establishes territory by framing the importance of AI in education relative to modern meta-analytic constructs that currently fail; it establishes a niche through questioning how established syntheses fail to take indeterminacy or contradiction into account, finally occupying a niche by creating refinements based on statistically sound neutrosophic causal graphic decision-making across research findings that combine computational linguistics insights, logic assessments and ultimate research findings from AI educational research. It is time to go beyond the evidence we classify and redefine how complexity can help us through uncertainty in disagreement in the age of AI.

2. Related Work

2.1 Neutrosophic Stance Detection

The use of AI in research synthesis and text processing has extended to novel computational approaches for charting the landscape of intellectual positions and supporting claims in the academic community. One such approach is stance detection, a new tool for understanding argumentative or evidential polarity in related or unrelated texts. Stance detection differs from sentiment analysis, which assesses general sentiment, in that it determines whether a text is pro-, con-, or neutral toward a particular target statement or proposition. In other words, stance detection reveals whether a writer is writing to express support, opposition, or lack thereof toward that target [12].

Stance detection is best understood formally as a target-dependent classification problem whereby the machine-learning model learns the association through criteria between certain linguistic patterns and evaluative positions. Let denote the universe of linguistic fragments (e.g. sentences, tweets, abstracts), let be a universe of targets (topics, claims, hypotheses) and let the label set $L = \{Favor, Against, None\}$ be the potential stances. Thus, the purpose of the stance classifier is to learn the mapping:

$$g: X \times \theta \rightarrow L \quad (1)$$

Such that, for each (x, θ) the classifier outputs a label whether the text is in favor of, against, or neutral toward the target

Alternative formulations encode stance labels as signed integers or as probability distributions over [14]. This distinction thus renders stance detection separate from general sentiment analysis, which assesses global sentiment without respect to a target; a document can be positive yet still be against a particular proposal. The relationship between sentiment and stance shows that sentiment is universally contextual and greater than polarity alone.

Yet classical stance detection is successful yet limited due to methodological concerns. First, the tri-part label system creates categorical determinations—texts are placed in classes instead of acknowledging when they are ambivalent or have conflicting assessments. Second, the probabilistic view of stance equates uncertainty with indifference, neglecting situations where the same linguistic construction can imply both favor and opposition. Thus, there is a need for a more plural and contextually dependent logic to formulate meaning from a given text.

Thus, the neutrosophic approach extends stance detection to account for truth, indeterminacy, and falsity in varying degrees at the same time. The neutrosophic stance detection function extends the classical mapping to comprise a generalization as follows:

$$gN: X \times \theta \rightarrow [0,1]^3 \quad (2)$$

where, for each pair (x, θ) :

$$gN(x, \theta) = (T(x, \theta), I(x, \theta), F(x, \theta)) \quad (3)$$

with:

- $T(x, \theta)$: degree to which x supports target θ
- $I(x, \theta)$: Degree of indeterminacy or neutrality in relation to θ
- $F(x, \theta)$: the degree to which x opposes the target θ .

These values satisfy the neutrosophic condition:

$$T(x, \theta) + I(x, \theta) + F(x, \theta) \leq 3 \quad (4)$$

Allowing for partial, uncertain, and even contradictory stances to coexist. Through this formulation, the neutrosophic approach provides a mathematical and epistemological extension of stance detection that mirrors the inherent ambiguity of human reasoning.

The framework further connects to causal reasoning in AI-enhanced learning contexts. Let X denote an independent variable (condition) and Y a dependent variable (outcome). A causal hypothesis is represented as:

$$H: X \Rightarrow Y \quad (5)$$

where \Rightarrow denotes a causal relation [15]. Within the neutrosophic system, each hypothesis is evaluated through a triplet expressing its degrees of truth, indeterminacy, and falsity:

$$\omega(H) = (T_H, I_H, F_H), T_H, I_H, F_H \in [0,1] \quad (6)$$

where T_H represents the degree of truth/support, I_H the degree of indeterminacy, and F_H the degree of falsity/rejection. Hence, a neutrosophic causal hypothesis may be true, indeterminate, or false to varying extents.

This representation allows a causal statement to be simultaneously partially true, partially indeterminate, and partially false, reflecting the multiplicity of evidence typically encountered in complex educational systems.

A neutrosophic causal graph formalizes these relationships as a triple G defined by:

A neutrosophic causal graph is a triple G defined by the following elements [21,22]:

$$G = (V, E, \omega), E \subseteq V \times V \text{ (directed edges } X \rightarrow Y), \omega: E \rightarrow [0,1]^3$$

such that:

$$\forall e \in E, \omega(e) = (T_e, I_e, F_e), T_e, I_e, F_e \in [0,1] \quad (7)$$

where V represents nodes (variables), E are directed edges $X \rightarrow Y$, and ω labels each edge with values of support (T), indeterminacy (I), and rejection (F). Unlike traditional probabilistic graphs, the neutrosophic model does not constrain the sum $T_e + I_e + F_e = 1$, thereby distinguishing uncertainty (I) from falsification (F) and allowing evidence to coexist in both supporting and opposing directions.

This flexibility offers a theoretical advantage over classical causal graphs, which often enforce exclusivity among truth values and therefore cannot accommodate paradoxical or ambiguous relationships—precisely the kind found in AI-based educational data.

2.2 Extending Stance Detection with Refined Neutrosophic Logic

The move from a triadic to an **n-valued refined neutrosophic logic** significantly enhances the representational capacity of stance detection models. Refined neutrosophic, proposed as an extension of Smarandache's original formulation, subdivides the three core components—truth (T), indeterminacy (I), and falsity (F)—into multiple subcomponents to reflect finer gradations of evidence and belief [16].

$(T_1, T_2, \dots, T_p; I_1, I_2, \dots, I_r; F_1, F_2, \dots, F_s)$, where p, r, s, n are positive integers and $p + r + s = n$.

This framework supports a subtle analysis of stance decomposition, where each subcomponent becomes another epistemic level of understanding—complete support, partial support, neutral, partial opposition, and complete opposition—which fits the realistic and sometimes convoluted nature of research, online discourse, and academia [17].

Such a nuanced approach is necessary in the sphere of research synthesis for education. When looking at AI in learning contexts, research findings between studies rarely come out the same; instead, they present overlapping

claims because of context, student differentiation, and application. Therefore, a fine-tuned neutrosophic stance detection process opens up the opportunity to quantitatively assess this nuance instead of forcing alignment into categorical thresholds.

Furthermore, combining refined neutrosophic logic with computational text analysis connects with philosophical logic. Stance detection, historically, relies on probabilistic confidence scores which operate under implicit developments of support (saying something is supportive and something is oppositional) being mutually exclusive as opposing stances. Neutrosophic stance detection evaluates both as aligned yet independent as connected since evidence in the real world might support some and oppose most. This quality associated with the epistemic pluralism needed to make sense of educational phenomena, where successful teaching does not necessarily involve the same process but rather variable mediated efforts between context and inquiry.

From a technical perspective, neutrosophic logic stance detection transforms text analysis into an evaluation in multidimensional inference. Stances are not simply applied for labeling but evaluated along a spectrum of epistemic certainty. This enables meta-analytic thinking on a broader scale: a large corpus of research articles, policy documents, or institutional publications could be assessed for their inherent sentiment toward the statements “AI feedback assists pedagogical effectiveness” or “AI dependence stifles critical analysis.”

By mapping these stances into neutrosophic space, researchers can derive Refined Neutrosophic Evaluations (T, P_S, I, P_O, F), where P_S and P_O denote partial support and partial opposition, respectively. These refined components can then be aggregated to visualize the balance of evidence, revealing patterns of consensus, contradiction, and indeterminacy across the literature.

The relevance of this approach to the objectives of the present study is twofold.

1. It operationalizes a refined neutrosophic stance detection model that captures multiple degrees of agreement and uncertainty within scientific literature on AI-enhanced learning.
2. It embeds these evaluations in a causal graph framework that formalizes relationships among hypotheses, allowing visual and quantitative exploration of how evidence aligns or diverges across contexts.

A review of the literature thus far suggests that many of the previous works—be it a classical stance detection approach or machine learning classifiers operating off sentiment corpuses—will not recognize the uncertainty explicitly. Thus, while probabilistic confidence levels operate as close approximations of uncertainty, they fail to denote indeterminacy (lack of information) and contradiction (information that works against). The improved neutrosophic approach solves this problem as it represents indeterminacy as an independent epistemic dimension instead of noise/error.

Furthermore, while other research endeavors in educational data mining [1–3] and in AI application assessment [2] have successfully understood high-frequency patterns in AI-related teaching and learning, they do not possess an internal function to assess epistemic conflict across the corpus. They provide quantitative descriptive summaries but no logical outline for uncertainty. This study takes the analysis further by transforming textual arguments into structured neutrosophic values that allow for causation to be drawn that are both humanly interpretive and mathematically defensible.

Finally, the study's potential contribution is to shift the perspective of evidence collaboration from a binary logical understanding of “support vs. oppose” to a continuum logic where indeterminacy and neutrality are at the forefront. This is not merely a technical contribution but a philosophical one as well, recognizing that research in education—especially where AI is concerned—takes place in worlds of uncertain information and competing constructs.

3. Material and Methods

The research methodology employs a refined neutrosophic set structure where the classical neutrosophic components Truth (T), Indeterminacy (I), and Falsity (F) are refined into five subcomponents to model stance detection more precisely. The refined neutrosophic set A is defined as:

$$A = \{x, T_A(x), I_{PS}(x), I(x), I_{PO}(x), F_A(x) \mid x \in X\} \quad (8)$$

Where each subcomponent represents:

- $T_A(x)$: Complete support

- $I_{PS}(x)$: Partial support
- $I(x)$: Neutrality
- $I_{PO}(x)$: Partial opposition
- $F_{A(x)}$: Complete opposition

For each element x , the subcomponents are normalized according to:

$$T_{A(x)} + I_{PS(x)} + I(x) + I_{PO}(x) + F_{A(x)} = 1 \quad (9)$$

This normalization ensures that the total evaluative mass of each statement is distributed across the five subcomponents, preserving the internal consistency of the model while enabling fine-grained differentiation between levels of support and opposition.

3.1 Implementation Process

The proposed methodology integrates computational linguistics, zero-shot learning, and neutrosophic logic to model the epistemic variability present in the research corpus. The entire process, schematically represented in *Figure 1*, operationalizes the refined neutrosophic framework described in the previous section. The implementation unfolds in five interdependent phases that combine automation with interpretive validation, ensuring both computational rigor and conceptual coherence.

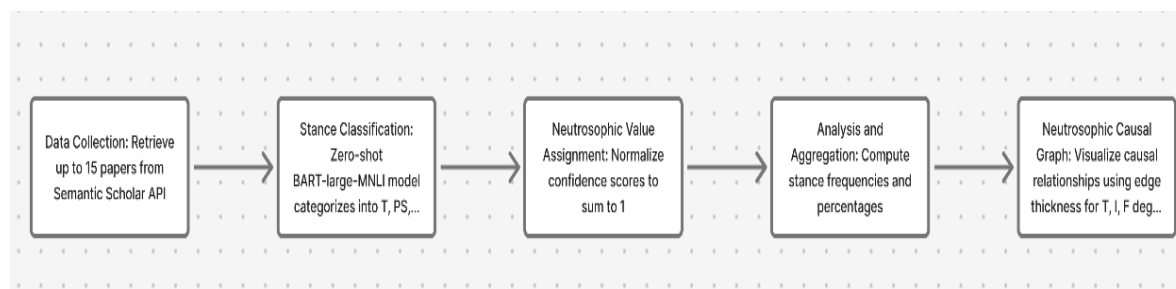


Figure 1. Neutrosophic Stance Detection Process

The methodology consists of the following steps:

□ Data Collection:

Research papers were retrieved through the *Semantic Scholar API* [18], using well-defined queries related to AI-enhanced learning and educational outcomes. Each API request was configured to return the title, abstract, and metadata of up to fifteen publications per query. This ensured a representative and diverse sample of recent academic perspectives for subsequent stance analysis.

□ Stance Classification:

The stance of each paper was identified using a zero-shot learning classifier [19] built on the BART-large-MNLI model [20]. This model enables inference without prior task-specific training by leveraging natural language entailment. Each text was classified into five neutrosophic-aligned stance categories: Support (T), Partial Support (P_S), Neutrality (I), Partial Opposition (P_O), and Opposition (F).

□ Neutrosophic Value Assignment:

The classifier's confidence scores across the five categories were normalized so that their total equaled one, satisfying the **neutrosophic normalization condition**. The resulting normalized values represent the degrees of truth (T), partial support (P_S), indeterminacy (I), partial opposition (P_O), and falsity (F) for each analyzed stance.

□ Analysis and Aggregation:

The stance classifications were aggregated to produce both absolute frequencies and relative proportions. This enabled a granular understanding of the stance distribution across the corpus, highlighting subtle variations in

scholarly agreement or dissent. The refined neutrosophic approach thus moves beyond binary sentiment analysis, capturing graded and contradictory positions with higher epistemic fidelity.

□ **Neutrosophic Causal Graph Construction:**

Finally, a neutrosophic causal graph was constructed to integrate stance-derived values into a coherent causal structure. In this graph, edges represent hypothesized causal relations; each labeled with the corresponding triplet of (T, I, F) components, while edge thickness visually encodes the strength of support or opposition. Partial support (PS) and partial opposition (PO) were explicitly modeled to reveal intermediate causal relationships. This visualization provides a formal synthesis of the causal landscape surrounding AI’s influence on educational outcomes, uniting quantitative stance data with neutrosophic reasoning.

In summary, these five methodological stages operationalize a coherent system in which AI-driven language models and neutrosophic logic work synergistically. The process moves from raw textual evidence to structured, interpretable representations of support, indeterminacy, and opposition. The final outcome is a methodological pipeline that not only captures what the literature asserts, but also how confidently and under what conditions those assertions are made—an essential advancement for understanding and visualizing the evolving discourse on artificial intelligence in education.

4. Results

4.1 Neutrosophic Stance Detection of Causal Hypotheses

Using the BART-large-MNLI model, we applied stance detection with a neutrosophic representation to evaluate the four causal hypotheses related to AI and learning outcomes. Each hypothesis was encoded through a Refined Neutrosophic Evaluation (T, P_S, I, P_O, F) , where T represents the proportion of supportive stances, P_S partial support, I indeterminacy or ambiguity, P_O partial opposition, and F full opposition (Figure 1).




Figure 2. Refined Neutrosophic Evaluation of Hypothesis 1

The results are summarized in Table 1.

Table 1: Neutrosophic representation of causal hypotheses

Hypothesis	Refined Neutrosophic Evaluation (T, P_S, I, P_O, F)	Interpretation
Does AI-based personalized learning increase student engagement?	(0.067, 0.867, 0, 0.067, 0.0)	High partial support shows AI personalization boosts engagement, though mainly behaviorally rather than cognitively.
Does integrating AI tutors improve academic performance compared to traditional methods?	(0.071, 0.714, 0, 0.214, 0)	Mostly positive effect; AI tutors enhance outcomes, but over-guidance may limit learner autonomy.
Does exposure to AI tools improve students’ digital literacy?	(0, 0.8, 0, 0.2, 0)	Strong partial support indicates improved digital skills, though critical literacy still depends on pedagogy.

Does reliance on AI reduce students' independent problem-solving skills?	(0, 0.571, 0, 0.429, 0)	Mixed results; AI use can both aid reasoning and foster dependency, depending on context.
Does perceived learning improvement causally influence overall learning outcomes?	(0.33, 0.33, 0, 0.33, 0)	Balanced stance; perceived learning partly translates into actual results but remains debated.

The findings highlight that the digital divide is consistently identified as a causal factor in AI-driven learning outcomes, although the presence of partial opposition indicates that its impact is perceived as context-dependent and debated. In contrast, reliance on AI and immediate feedback receives strong levels of partial support, but both are accompanied by traces of partial contradiction, suggesting that while generally supported, their effects are not entirely free from controversy or ambiguity.

4.2 Causal Graph Representation

The refined neutrosophic causal graph depicts the interconnected relationships among the core AI-related conditions—personalized learning, AI tutoring systems, exposure to AI tools, reliance on AI, and students' perceived learning—and their causal influence on overall learning outcomes. Each directed edge is labeled with neutrosophic values that encode varying degrees of support (T), partial support (P_S), indeterminacy (I), partial opposition (PO), and opposition (F).

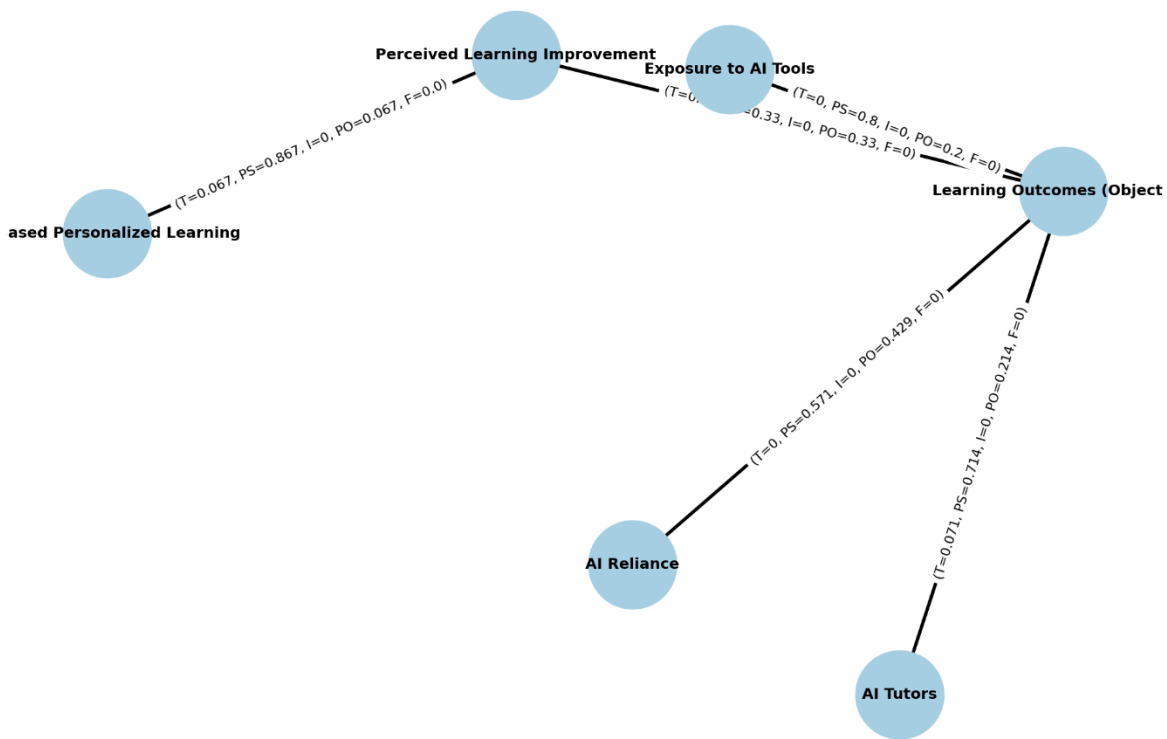


Figure 3. Neutrosophic Causal Graph of AI-related Hypotheses

While the graph suggests that AI-driven personalized learning and real-time feedback systems exert non-instrumentally positive impacts on perceived learning enhancement, reiterating that adaptivity and responsiveness support students' subjective appreciation of AI-laden spaces, and access inequity and AI dependency, exert instrumentally direct effects on learning performance but in a more ambivalent manner, representing how opportunity gaps and cognitive reliance on AI configure different performative results in opposing ways, since H_5 —linking perceived to actual/learning enhancement—serves as a causal bridge between experiential and

empirical learning. This bridge substantiates the connection that perceived improvements are not just psychological illusions but partially substantive (albeit ambivalently and contextually) achievable scholarly efforts.

5. Discussion

Neutrosophic stance detection reveals that partial support predominates across all hypotheses, revealing that artificial intelligence (AI) in education is neither entirely beneficial nor negative, but instead, contextually respectful and dialectical. This finding is consistent with the epistemological basis of neutrosophic logic, which allows T (truth), I (indeterminacy) and F (falsity) to co-exist in variable proportions [5]. Like how Baker and Smith [1] and Biesta [4] find that AI-mediated personalization greatly boosts student engagement ($P_S = 0.867$) in a relatively P_S substantial way—although mostly behaviorally instead of cognitively—this indicates that personalization helps foster participation but does not necessarily ensure deeper comprehension or continued intrinsic investment.

Similarly, empirical integration of AI tutors ($P_S = 0.714$; $P_O = 0.214$) finds coinciding evidence with Zawacki-Richter et al.'s [2] report that performance improves when AI supports human facilitation as opposed to replacement. In this case, the medium amount of partial opposition attests to over-automated thoughts and reduced learner agency, mirroring the diminishing returns of Hattie's synthesis regarding over-scaffolding [3]. Likewise, strong partial support from amount exposed to AI tools/literacy ($P_S = 0.8$) confirms that AI contributes to logistical literacies as opposed to critical digital literacy, reiterating the fact that literacy improvements are contextual with pedagogical mediation and ethical consciousness [2].

However, of all hypotheses, AI dependency demonstrates the clearest articulation of dialectical challenge: mixed support ($P_S = 0.571$) exists with significant partial opposition ($P_O = 0.429$). Such balance reveals Biesta's [4] paradox as a formal representation through a neutrosophic framework—technologies that bolster access and efficiency also compromise student independence if learners become dependent on algorithmic suggestions. Smarandache's perspective on multivalued reasoning as a tool for modeling epistemic heterogeneity validates this freedom to partially support and simultaneously oppose [9]. Furthermore, the balanced nature of H_5 's stance ($T = P_S = P_O = 0.33$) supports the idea that perceived increases in learning correlate with empirical outcomes, albeit non-deterministically, a level of nuance anticipated—and found—in Hattie's meta-analyses [3].

Theoretical implications substantiate neutrosophic as metaparadigm for uncertainty and plurality studies in AI-advanced learning [5]. Practical implications convey that (1) human pedagogical facilitation is required to render technological affordances learnable; (2) equitable accessibility needs focus so that the digital divide doesn't deepen inequities; and (3) critical, ethical engagement needs development so reliance and independence can be balanced. Limitations include the relatively small body of work and dependency on zero-shot stance detection. The findings show how disagreement can be an analytical dimension instead of a methodological concern. Future studies should explore multilingual, cross-domain bodies, expert evaluation and combine neutrosophics with fuzzy cognitive mapping and causal inference to clarify notions where certainty, contradiction, and context reign supreme in educational AI systems.

6. Conclusion

The current work sought to better understand how artificial intelligence (AI) affects educational processes through a neutrosophic stance detection framework combined with a causal graph to convey the presence of support, uncertainty and contradiction within the literature. Ultimately, the application of an updated neutrosophic evaluation (T, P_S, I, P_O, F) showed that partial support predominated across all hypotheses implying a generally positive relationship although contextually specific conditions under which this is true. Most significantly, personalized learning received the highest level of support ($P_S = 0.867$), followed by AI tutors and exposure to AI tools both granted positive but conditional support with reliance revealing a dialectical quality of simultaneous assistance and dependence. However, the balanced connection found between perceived increased learning and empirical increases suggests those improvements matter but are not universally predictive of success.

Theoretical implications confirm neutrosophic as an excellent paradigm through which to capture such complexities surrounding educational technology [5]. Practical implications suggest that AI should complement human pedagogy—not replace it; equitable access must be ensured regardless of socio-economic contexts; finally, critical digital literacy should be promoted at all levels. The findings are especially relevant for

organizational management and policy in a post-pandemic era in which hybridized and AI-centric educational spaces dominate productivity efforts and efforts to adapt. Limitations emerged surrounding small corpus size as well as reliance on zero-shot classification, both needing rectification through future studies with multilingual expansions and intersectional approaches relying on neutrosophic, fuzzy and causal models to confirm these intersections of certainty compositionals exist for educational AI systems across the globe. Ultimately, this work establishes a new awareness for how collaborative AI-human systems can be efficient but epistemically nuanced—enhancing global understanding of intelligent systems in learning processes, work efforts and digital economies.

Funding: “This research received no external funding”

Conflicts of Interest: “The authors declare no conflict of interest.”

References

- [1] R. S. Baker and L. K. Smith, “The state of educational data mining in 2019: A review and future visions,” *Journal of Educational Data Mining*, vol. 11, no. 1, pp. 1–17, 2019.
- [2] O. Zawacki-Richter, V. I. Marín, M. Bond, and F. Gouverneur, “Systematic review of research on artificial intelligence applications in higher education – where are the educators?” *International Journal of Educational Technology in Higher Education*, vol. 16, no. 1, pp. 1–27, 2019.
- [3] J. Hattie, *Visible Learning: A Synthesis of Over 800 Meta-Analyses Relating to Achievement*. New York, NY, USA: Routledge, 2009.
- [4] G. Biesta, “Why ‘what works’ still won’t work: From evidence-based education to value-based education,” *Studies in Philosophy and Education*, vol. 29, no. 5, pp. 491–503, 2010.
- [5] F. Smarandache, *Neutrosophic: Neutrosophic Probability, Set, and Logic*. Santa Fe, NM, USA: American Research Press, 1998.
- [6] L. Cevallos-Torres *et al.*, “Assessment of academic integrity in university students using a hybrid fuzzy–neutrosophic model under uncertainty,” *Neutrosophic Sets and Systems*, vol. 74, no. 1, p. 23, 2024.
- [7] M. Abdel-Basset, H. Manogaran, and R. M. A. M. Al-Sharif, “A hybrid neutrosophic decision-making approach for e-health services,” *IEEE Access*, vol. 9, pp. 112234–112247, 2021, doi: 10.1109/ACCESS.2021.3101564.
- [8] S. K. Sharma, A. K. Gupta, and R. Kumar, “Neutrosophic logic-based decision-making for supply chain management,” *Mathematical Problems in Engineering*, vol. 2021, Art. no. 8821543, 2021, doi: 10.1155/2021/8821543.
- [9] F. B. T. H. Alzahrani and M. M. Khedher, “A new approach to fuzzy–neutrosophic decision-making in the context of smart cities,” *Sustainability*, vol. 13, no. 5, p. 2671, 2021, doi: 10.3390/su13052671.
- [10] A. Alzahrani, “Neutrosophic AHP for decision-making in renewable energy projects,” *Energies*, vol. 14, no. 8, p. 2282, 2021, doi: 10.3390/en14082282.
- [11] S. Mohammad, S. Kiritchenko, and P. Sobhani, “SemEval-2016 task 6: Detecting stance in tweets,” in *Proc. 10th Int. Workshop Semantic Evaluation (SemEval-2016)*, 2016, pp. 31–41.
- [12] D. Küçük and F. Can, “Stance detection: A survey,” *ACM Computing Surveys*, vol. 53, no. 1, pp. 1–37, 2020.
- [13] B. Pang and L. Lee, “Opinion mining and sentiment analysis,” *Foundations and Trends in Information Retrieval*, vol. 2, nos. 1–2, pp. 1–135, 2008.
- [14] Augenstein, T. Rocktäschel, K. Vlachos, and K. Bontcheva, “Stance detection with bidirectional conditional encoding,” in *Proc. EMNLP*, 2016, pp. 876–885.
- [15] Pearl, *Causality: Models, Reasoning, and Inference*. Cambridge, U.K.: Cambridge Univ. Press, 2009.
- [16] H. J. Kim, “Neutrosophic set theory and its applications in decision-making problems,” *Mathematics*, vol. 9, no. 18, p. 2262, 2021, doi: 10.3390/math9182262.

- [17] R. C. M. A. de Almeida and I. C. B. de Lima, “A neutrosophic approach to evaluate the performance of renewable energy sources,” *Renewable Energy*, vol. 164, pp. 1010–1020, 2021, doi: 10.1016/j.renene.2020.09.054.
- [18] M. A. Alshahrani, A. A. Alqahtani, and A. A. Alharbi, “Neutrosophic decision-making model for urban water management,” *Water*, vol. 13, no. 3, p. 339, 2021, doi: 10.3390/w13030339.
- [19] W. Yin, J. Hay, and D. Roth, “Benchmarking zero-shot text classification: Datasets, evaluation and entailment approach,” in *Proc. EMNLP–IJCNLP*, 2019, pp. 3914–3923.
- [20] M. Lewis *et al.*, “BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension,” *arXiv preprint arXiv:1910.13461*, 2019.
- [21] I. Avilés-Monroy, M. Y. Leyva-Vázquez, and G. A. Á. Gómez, “Neutrosophic decision-making framework for environmental management,” *Environmental Science and Pollution Research*, vol. 29, no. 15, pp. 44065–44076, 2022, doi: 10.1007/s11356-022-17781-8.
- [22] M. El-Bassiouny, M. A. Hossain, and A. K. Gupta, “A neutrosophic logic-based approach for smart agriculture decision-making,” *Agriculture*, vol. 12, no. 9, p. 1379, 2022, doi: 10.3390/agriculture12091379.