A Tri-Criterion Framework for Interpretable, Adaptive, and Robust Fuzzy Decision-Making Under Uncertainty An IRA-FIS Approach with an Environmental Health Risk-Assessment Case Study

Rafael Canorea-García[1], Martin González y Santiago[2], Roger Sanz Gonzalez[3], Víctor Rodríguez González[4], Javier Diz-Casal[5], Claudio Paya Santos[*6]

[1] ESIC University, Spain
Email: rafael.canorea@esic.university
ORCID: https://orcid.org/0000-0002-6637-4369

[2] Fernando Pessoa University Canary Islands, Spain
Email: mgonzalez@ufpcanarias.es
ORCID: https://orcid.org/0000-0002-4174-4474

[3] Valencian International University, Spain
Email: rsanzg@professional.universidadviu.com
https://orcid.org/0009-0006-1419-2367

[4] Isabel I of Castile University, Spain
Email: victor.rodriguez.gonzalez@ui1.es
ORCID: https://orcid.org/0000-0002-5348-9730

[5] Valencia International University, Spain
Email: jdizca@professor.universidadviu.com
ORCID: https://orcid.org/0000-0003-1332-8905

[*6] Valencia International University, Spain
Email: claudio.paya@professor.universidadviu.com
ORCID: https://orcid.org/0000-0002-1908-9960

Abstract: Fuzzy decision-making provides a principled way to model uncertainty, imprecision, and subjectivity in complex real-world systems, yet three properties that largely determine whether a fuzzy model is actually adopted in practice — interpretability, adaptability, and robustness — are rarely evaluated jointly within a single, reproducible procedure. This paper introduces IRA-FIS, a fuzzy inference framework that couples a Mamdani-type rule base with three explicit, quantifiable criteria: a rule-complexity interpretability index, a bounded online adaptability mechanism for membership-function drift, and a robustness index derived from Monte Carlo perturbation of inputs. We formalize each criterion mathematically, define a configurable composite IRA score for comparing competing fuzzy designs, and demonstrate the framework on an environmental health risk-assessment case study that classifies exposure risk from three uncertain stressors: pollutant concentration, exposure duration, and population vulnerability. Sensitivity analysis shows that the resulting system maintains stable centroid outputs under input noise of up to 20% while preserving a compact, expert-auditable rule base of fewer than thirty rules. We situate the contribution relative to recent neuro-fuzzy, type-2 fuzzy, and explainable-AI hybrids reported between 2023 and 2026, and discuss limitations and avenues for extending IRA-FIS toward type-2 fuzzy sets and constrained data-driven rule learning. The framework is offered as a methodological contribution directly aligned with this Collection's call for algorithmic developments and practical applications of fuzzy logic in decision-making that improve interpretability, adaptability, and robustness.

A Tri-Criterion Framework for Interpretable, Adaptive, and Robust Fuzzy Decision-Making Under Uncertainty An IRA-FIS Approach with an Environmental Health Risk-Assessment Case Study

Keywords: fuzzy logic; decision-making under uncertainty; interpretability; adaptability; robustness; Mamdani inference; risk assessmentIntroduction

1. Introduction

Human decision-makers and artificial systems alike must routinely act on information that is incomplete, imprecise, or expressed in qualitative rather than numerical terms. Two broad views of this difficulty coexist in the literature. The first, formalist view treats knowledge as a set of propositions that are strictly true or false, with degrees of confidence captured probabilistically. The second view treats knowledge as inherently graded: a statement such as "exposure was moderate" does not become more or less true as more data arrive, but rather designates a concept whose boundaries are themselves fuzzy. Fuzzy set theory, introduced by Zadeh (1965), formalizes the second view by allowing partial membership in a set, and fuzzy logic extends classical bivalent reasoning to a continuum of truth values between 0 and 1.

Fuzzy decision-making methods have since matured into a broad family of tools — Mamdani (1974) and Takagi-Sugeno inference systems, adaptive neuro-fuzzy inference systems (ANFIS; Jang, 1993), and type-2 and type-3 fuzzy extensions that model uncertainty about the membership functions themselves — with documented applications in engineering control, healthcare diagnosis, environmental management, finance, and supply-chain risk. Fuzzy logic has also been applied to criminal intelligence analysis (Martino et al., 2026) and to the management of uncertain operational situations (Payá Santos et al., 2023), illustrating its versatility in high-complexity, high-ambiguity domains. A 2024 systematic review of 27 comparative studies in environmental engineering found that fuzzy inference systems matched or outperformed competing machine-learning methods in 21 of 27 cases (77.8%), attributing this advantage specifically to the transparency of fuzzy reasoning under uncertain, noisy, and heterogeneous environmental data (Bressane et al., 2024).

Despite this track record, a recurring criticism of fuzzy systems is that individual studies tend to report predictive accuracy without systematically reporting whether the resulting model remains interpretable to the domain expert, whether it can be updated as new evidence arrives without an expensive full redesign, and whether its conclusions are stable under realistic levels of input noise or rule perturbation. These three properties — interpretability, adaptability, and robustness — are each individually studied in the literature, but rarely evaluated jointly within one reproducible procedure attached to a single fuzzy model. This separation matters because the three properties can trade off against one another: a highly adaptive, data-driven rule base (as in ANFIS-style learning) may sacrifice the auditability that domain experts require, while a maximally interpretable, small rule base may be brittle under distributional shift. The challenge extends to emerging domains such as global cybersecurity governance, where decision-making under legal and technical uncertainty equally demands transparent and robust tools (Paya Santos et al., 2026).

This paper addresses that gap with three contributions:

A formal tri-criterion framework, IRA-FIS, that defines an Interpretability index, an Adaptability index, and a Robustness index for any Mamdani-type fuzzy inference system, together with a configurable composite IRA score that lets stakeholders weight the three properties according to application needs.

A generalizable algorithmic procedure for computing the three indices that is independent of application domain, requiring only the rule base, the membership functions, and a perturbation budget.

A worked case study in environmental health risk assessment, with explicit rule tables, sensitivity analysis under Monte Carlo input perturbation, and a qualitative comparison against plain Mamdani, ANFIS, and interval type-2 fuzzy designs reported in the recent literature.

The remainder of the paper is organized as follows. Section 2 reviews the theoretical background on fuzzy sets, fuzzy inference, and the uncertainty taxonomy relevant to decision-making, and surveys recent (2023–2026) developments in interpretable, adaptive, and robust fuzzy systems. Section 3 presents the IRA-FIS framework and its three indices. Section 4 develops the environmental health risk-assessment case study. Section 5 discusses the results, limitations, and generalizability. Section 6 outlines future directions, and Section 7 concludes.

2. Background and Related Work

2.1. Fuzzy Sets and the Uncertainty Taxonomy

A fuzzy set A on a universe of discourse X is defined by a membership function $\mu\_A: X \rightarrow [0,1]$, where $\mu\_A(x)$ denotes the degree to which element x belongs to A, rather than a binary in/out judgment. Linguistic variables — for example, “low”, “moderate”, and “high” concentration — are represented as families of fuzzy sets that jointly partition a numerical domain, allowing qualitative expert knowledge to be expressed and combined with quantitative data (Zadeh, 1965, 1975).

It is useful to distinguish two qualitatively different sources of uncertainty. Epistemic uncertainty arises from incomplete or limited knowledge about a fact, event, or measurement; it is, in principle, reducible by gathering more evidence. Aleatory uncertainty, by contrast, reflects intrinsic variability in the phenomenon itself and cannot be eliminated by collecting more data. Classical fuzzy sets primarily address epistemic, linguistic vagueness; type-2 and type-3 fuzzy sets extend the formalism to model uncertainty about the membership grades themselves, which is closer to a hybrid epistemic-aleatory treatment and has seen renewed interest in fuzzy control of nonlinear systems.

2.2. Fuzzy Inference Systems

A fuzzy inference system (FIS) maps inputs to an output through a base of IF-THEN rules, a fuzzification stage, an inference stage, and, where a crisp output is required, a defuzzification stage. The Mamdani system (1974) combines fuzzy antecedents and consequents via fuzzy implication and aggregates the outputs before defuzzification, commonly via the centroid method. The Takagi-Sugeno system uses functional consequents, simplifying defuzzification and lending itself more naturally to gradient-based parameter tuning. Mamdani systems are typically favored when interpretability and proximity to expert reasoning are priorities, which is why we adopt a Mamdani backbone for IRA-FIS, while acknowledging that the framework’s three indices are equally definable for Takagi-Sugeno systems.

Adaptive neuro-fuzzy inference systems (ANFIS; Jang, 1993) embed a Sugeno-type FIS within a neural-network training procedure, allowing membership-function parameters to be tuned from data via gradient descent or hybrid learning rules. ANFIS and its variants have demonstrated strong predictive performance across hydrology, energy, and medical applications, but the resulting membership functions, after training, do not always retain a transparent linguistic correspondence to the original expert-defined terms — the central interpretability trade-off this paper seeks to make explicit and measurable.

2.3. Recent Developments (2023–2026)

Three strands of recent literature directly motivate the IRA-FIS design.

Interpretability.

Work on rule-base minimization has shown that large, automatically induced fuzzy knowledge bases can be substantially compressed without loss of decision accuracy. A rough-set-based Reducing Rules and Conditions (RRC) algorithm applied to an environmental risk-assessment knowledge base of 6,215 rules and 26,959 conditions extracted a compressed base of 2,562 rules and 9,319 conditions while preserving classification performance (Ciaramella, 2024), illustrating both the scale that real fuzzy knowledge bases can reach and the importance of an explicit complexity metric for auditability. Complementary work on fuzzy inference systems for disease diagnosis has emphasized that the semantic interpretability of fuzzy rules is a distinct property from predictive accuracy and should be reported alongside it. Fuzzy logic has also been applied to criminal intelligence analysis, where traceability of reasoning is a critical operational requirement (Martino et al., 2026; Payá Santos et al., 2023).

Adaptability.

Adaptive fuzzy inference layers have been proposed as a meta-optimization mechanism for decision-support systems operating under data uncertainty, balancing interpretability, computational feasibility, and adaptability by adjusting parameters online as new evidence becomes available. Reviews of fuzzy machine learning in environmental engineering similarly call for adaptive FIS models capable of dynamically recalibrating as IoT and remote-sensing data streams evolve (Bressane et al., 2024), while explicitly flagging the computational cost of large-scale optimization as an open challenge.

Robustness.

Robustness in the sense used here — stability of model output under input or rule perturbation — has been examined through fuzzy-measure-based approaches to robust motion planning and adaptive control, and through hybrid fuzzy-ethical risk frameworks that explicitly report sensitivity analyses as part of an interpretable, traceable risk score. The robustness challenge also has a governance dimension: in domains such as global cybersecurity, where harmonizing international legal frameworks under high normative uncertainty constitutes itself a fuzzy decision problem (Paya Santos et al., 2026), model stability under parameter variation is a policy requirement, not merely a technical one.

2.4. Research Gap

Taken together, this literature confirms that fuzzy decision-making is an active and productive area, but each strand — interpretability, adaptability, robustness — is typically advanced and evaluated in isolation, on different case studies, with different metrics. IRA-FIS responds to this gap by defining the three properties on a common formal footing, applicable to any Mamdani-type rule base, and by reporting all three for a single worked system rather than for three separate ones.

3. The IRA-FIS Framework

3.1. Overview

IRA-FIS augments a standard Mamdani fuzzy inference pipeline (fuzzification → rule evaluation → aggregation → defuzzification) with three diagnostic procedures that run alongside model construction: (i) an Interpretability index computed from the rule base and membership-function structure; (ii) an Adaptability mechanism that allows bounded, expert-constrained online adjustment of membership functions as new data arrive, together with an index quantifying how much adjustment was required; and (iii) a Robustness index computed via Monte Carlo perturbation of the inputs and, optionally, the rule confidence weights. The three indices are combined into a single composite IRA score for comparing alternative rule bases or membership-function choices for the same problem.

3.2. Interpretability Index (II)

Interpretability is operationalized as a decreasing function of rule-base complexity and membership-function overlap, following the principle that a fuzzy system is more interpretable when it uses fewer rules, fewer antecedents per rule, and membership functions that are clearly separated along the linguistic spectrum. We define:

$$II = w_1 \cdot (1 - R/R\_max) + w_2 \cdot (1 - \bar{A}/\bar{A}\_max) + w_3 \cdot (1 - \bar{O})$$

where R is the number of active rules and R_max a stakeholder-defined upper bound considered auditable for the domain (we use R_max = 50, consistent with expert-review practice reported for clinical and environmental fuzzy knowledge bases); $\bar{A}$ is the mean number of antecedents per rule and $\bar{A}$_max its bound; $\bar{O}$ is the mean pairwise overlap between adjacent membership functions, computed as the normalized intersection area between consecutive triangular or trapezoidal membership functions; and w_1, w_2, w_3 are non-negative weights summing to 1. II ranges over [0,1], with higher values indicating a more easily audited model.

3.3. Adaptability Index (AI)

Adaptability is operationalized through a bounded online update rule for membership-function parameters rather than unconstrained data-driven retraining, so that any adjustment remains traceable to a specific evidentiary trigger and stays within expert-approved limits. For a triangular membership function parameterized by (a, b, c), IRA-FIS allows the update:

$$b' = b + \eta \cdot (\hat{b} - b), \text{ subject to } |b' - b_0| \leq \delta$$

where $\hat{b}$ is a new peak estimate suggested by incoming data, $\eta \in (0,1]$ is a damping factor, b_0 is the original expert-defined peak, and δ is an expert-approved maximum drift. This keeps the system within a bounded neighborhood of the original, expert-validated configuration — deliberately more conservative than full ANFIS-style gradient training (Jang, 1993), trading some predictive flexibility for retained auditability. The Adaptability Index reports, after a defined observation window, the mean normalized drift:

$$AI = 1 - (1/M)\, \Sigma\, |b'_m - b_{0,m}| \,/\, \delta_m$$

for m = 1,…,M membership functions in the system. AI close to 1 indicates the system required little adjustment to remain calibrated (suggesting the original design generalizes well); a low but non-zero AI indicates the bounded-adaptation mechanism is actively and usefully absorbing distributional change within the expert-approved envelope.

3.4. Robustness Index (RI)

Robustness is assessed by Monte Carlo perturbation: each crisp input x_i is resampled N times as $x_i + \varepsilon$, where $\varepsilon \sim U(-p \cdot x_i,\, p \cdot x_i)$ for a noise level p (e.g., p = 0.10, 0.20, 0.30), and the resulting defuzzified outputs are recorded. The Robustness Index at noise level p is:

$$RI(p) = 1 - \sigma_y(p) / \mathrm{range}(Y)$$

where $\sigma_y(p)$ is the standard deviation of the N perturbed outputs and range(Y) is the full output range, so that $RI(p) \in [0,1]$ is comparable across systems with different output scales. We additionally report a rule-perturbation variant, attenuating a random subset of rule firing strengths by up to 20% to emulate partial sensor dropout or rule misfire.

3.5. Composite IRA Score

The three indices are combined into a single composite score:

$$IRA = \alpha \cdot II + \beta \cdot AI + \gamma \cdot \overline{RI}, \text{ with } \alpha + \beta + \gamma = 1$$

where $\overline{RI}$ is the Robustness Index averaged across the tested noise levels, and α, β, γ are stakeholder-defined weights. We deliberately do not propose a single universal weighting: a safety-critical application may set γ high, while a policy-communication tool may set α high. What IRA-FIS standardizes is the computation of the three underlying indices, not the value judgment of how to combine them.

3.6. Procedure Summary

The end-to-end IRA-FIS procedure is summarized below:

1. Define linguistic variables and elicit an initial expert rule base; compute the baseline Interpretability Index (II).
2. Run Mamdani inference (fuzzification, rule aggregation, centroid defuzzification) on the validation dataset to obtain baseline outputs.
3. Apply the bounded adaptation mechanism over a defined observation window of new cases; recompute the Adaptability Index (AI).
4. Run Monte Carlo input perturbation (and, optionally, rule-firing perturbation) at two or more noise levels; compute the Robustness Index (RI) at each level.
5. Combine II, AI, and $\overline{RI}$ into the composite IRA score using stakeholder-defined weights; compare against alternative designs if applicable.

This procedure requires only the rule base, the membership-function definitions, and a perturbation/observation budget, and is therefore portable across application domains.

4. Case Study: Environmental Health Risk Assessment Under Multi-Stressor Uncertainty

4.1. Problem Definition

We illustrate IRA-FIS on an environmental health risk-assessment problem: classifying the exposure risk of a population to a pollutant based on three uncertain inputs commonly used in environmental risk characterization: pollutant concentration (a normalized exposure index, 0–100), exposure duration (hours per day, 0–24), and population vulnerability (a composite index of age, pre-existing conditions, and socioeconomic exposure factors, 0–10). The output is a continuous Risk Level (0–100), defuzzified by the centroid method, that downstream policy tools can threshold into categorical alerts. This problem class is structurally representative of the broader family of environmental and public-health fuzzy decision problems documented in the recent systematic literature, including coagulant-dosage control, drought-index estimation, and groundwater-quality assessment (Bressane et al., 2024).

4.2. Fuzzy Sets and Membership Functions

Each input variable is partitioned into three linguistic terms using triangular membership functions, and the output variable into five terms to provide finer-grained risk communication. Table 1 reports the membership-function parameters used.

Variable	Linguistic Term	Membership Function	Parameters (a, b, c)
Concentration [0–100]	Low	Triangular	(0, 0, 40)
	Moderate	Triangular	(20, 50, 80)
	High	Triangular	(60, 100, 100)
Duration [0–24 h]	Short	Triangular	(0, 0, 8)
	Medium	Triangular	(4, 12, 20)
	Long	Triangular	(16, 24, 24)
Vulnerability [0–10]	Low	Triangular	(0, 0, 4)
	Moderate	Triangular	(2, 5, 8)
	High	Triangular	(6, 10, 10)
Risk [0–100]	Very Low	Triangular	(0, 0, 20)
	Low	Triangular	(10, 25, 40)
	Moderate	Triangular	(30, 50, 70)
	High	Triangular	(60, 75, 90)
	Very High	Triangular	(80, 100, 100)

Table 1. Membership-function parameters for the environmental health risk-assessment case study.

4.3. Rule Base

The rule base implements the full factorial combination of the three input variables' linguistic terms (3 × 3 × 3 = 27 rules), with consequents assigned by an expert-defined weighting scheme in which vulnerability is given the highest influence on risk (weight 0.45), followed by concentration (0.35) and duration (0.20), reflecting the established public-health principle that population susceptibility is frequently the dominant driver of realized harm at a given exposure level (Payá Santos et al., 2023). Table 2 reports a representative subset of the 27 rules; the complete rule base is available from the corresponding author upon request.

IF Concentration is	AND Duration is	AND Vulnerability is	THEN Risk is
Low	Short	Low	Very Low
Low	Medium	High	Moderate
Moderate	Short	Moderate	Low
Moderate	Long	High	High
High	Short	Low	Low
High	Medium	High	High
High	Long	Moderate	Moderate
High	Long	High	Very High

Table 2. Representative subset of the 27-rule IRA-FIS rule base (3 antecedents each).

4.4. Mamdani Inference and Defuzzification

Inference follows the standard Mamdani (1974) pipeline: each rule's firing strength is computed via the minimum (AND) t-norm over the three fuzzified antecedents; rule outputs are aggregated via the maximum (OR) s-norm; and the aggregated output fuzzy set is defuzzified via the centroid method. We implemented the system in Python using scikit-fuzzy and verified it against five representative test scenarios spanning the input space (Table 3).

Concentration	Duration (h)	Vulnerability	Defuzzified Risk
20	4	2.0	7.78
50	12	5.0	25.00
85	20	8.5	75.00
70	6	9.0	60.30
30	18	3.0	33.77

Table 3. Defuzzified Risk Level for five representative test scenarios under the baseline IRA-FIS configuration.

The results follow the expected monotonic pattern. The fourth scenario (high concentration, short duration, very high vulnerability: 60.30) scores above what the third scenario's intermediate-duration analogue would suggest, illustrating how the vulnerability-weighted rule base correctly captures that a highly susceptible population can reach elevated risk even under comparatively brief exposure — a pattern difficult to encode in a purely linear scoring rule but that emerges naturally from the rule base's explicit conditioning on vulnerability, as discussed in the context of uncertain operational situations by Payá Santos et al. (2023).

4.5. IRA Evaluation of the Case Study

4.5.1. Interpretability Index

Applying the formula in Section 3.2 with $R = 27$ rules (against $R\_max = 50$), $\bar{A} = 3$ antecedents per rule (against $\bar{A}\_max = 5$), and a mean pairwise membership-function overlap of $\bar{O} = 0.055$, with weights $w_1 = 0.4$, $w_2 = 0.3$, $w_3 = 0.3$, yields:

II = 0.4(1 − 27/50) + 0.3(1 − 3/5) + 0.3(1 − 0.055) = 0.588

This places the system in the upper-middle range of the [0,1] interpretability scale: the rule base is compact relative to the auditability bound, and the membership functions show low overlap, indicating that adjacent linguistic terms remain visually and semantically distinguishable — a property that directly supports stakeholder trust, consistent with the transparency advantage reported for fuzzy inference systems by Bressane et al. (2024).

4.5.2. Adaptability Index

We illustrate the bounded adaptation mechanism with a realistic scenario: new cohort data suggest that the empirical peak for "high" vulnerability should shift slightly ($\hat{b}$ = 9.3 vs. original b_0 = 10.0). Applying the bounded update with η = 0.5 and δ = 1.0:

b′ = 10.0 + 0.5(9.3 − 10.0) = 9.65, with drift used = 0.35 (35% of the allowed envelope)

AI = 1 − 0.35/1.0 = 0.65

An AI of 0.65 indicates the system absorbed a meaningful but not extreme calibration shift while remaining within its expert-approved envelope.

4.5.3. Robustness Index

Monte Carlo perturbation was applied to the five test scenarios of Table 3, resampling each input 300 times at three noise levels. Table 4 reports the results.

Noise Level	Mean Output Std. Dev.	Robustness Index RI(p)	Trials per Scenario
10%	1.99	0.980	300
20%	4.08	0.959	300
30%	6.13	0.939	300

Table 4. Robustness Index under Monte Carlo input perturbation at three noise levels.

Across all three noise levels the Robustness Index remains above 0.93. Rule-firing perturbation (random attenuation up to 20%) yielded a mean output standard deviation of 2.57 and RI = 0.974, confirming graceful rather than catastrophic degradation under partial information loss.

$\overline{RI}$ (averaged across the three input-noise levels) = 0.959

4.5.4. Composite IRA Score

Using equal weights (α = β = γ = 1/3), appropriate for a general-purpose policy-communication deployment:

IRA = (1/3)(0.588) + (1/3)(0.650) + (1/3)(0.959) = 0.732

A safety-critical deployment with (α, β, γ) = (0.2, 0.2, 0.6) yields IRA = 0.812; a public-facing transparency-first tool with (0.5, 0.2, 0.3) yields IRA = 0.711. The composite score is intended as a comparison device across competing designs rather than an absolute benchmark.

4.6. Qualitative Comparison with Alternative Fuzzy Architectures

Table 5 situates IRA-FIS against three alternative architectures reported in the recent literature, based on the qualitative emphasis each design places on the three criteria; a head-to-head numerical benchmark on an identical dataset is left as future work (Section 6).

Architecture	Interpretability	Adaptability	Robustness	Typical Use Case
Plain Mamdani FIS	High	Low (manual redesign)	Moderate–High	Expert-elicited rules, stable domain
ANFIS (Sugeno + NN)	Low–Moderate	High (gradient-based)	Moderate	Large labeled datasets, accuracy-critical
Interval Type-2 FIS	Moderate	Moderate	High	Second-order uncertainty
IRA-FIS (this work)	Moderate–High	Moderate (bounded)	High	Expert-auditable, evolving evidence

Table 5. Qualitative positioning of IRA-FIS relative to representative alternative fuzzy architectures.

5. Discussion

The case study results support the central premise of this paper: it is feasible to report interpretability, adaptability, and robustness for a single fuzzy decision-making system using a common, reproducible procedure. The Interpretability Index (0.588) reflects a deliberate design choice to retain three antecedents per rule for semantic completeness — a requirement recognized in high-stakes domains such as criminal intelligence analysis, where reasoning traceability is a critical operational constraint (Martino et al., 2026) — ; a design team prioritizing maximal auditability could trade some expressiveness for a higher II. The bounded Adaptability mechanism (AI = 0.65) is intentionally more conservative than full data-driven retraining: it absorbs a meaningful calibration shift while remaining traceable to a specific evidentiary trigger and within an expert-approved envelope, which we view as a desirable property for regulatory or public-health contexts where unexplained model drift would itself be a governance risk. The Robustness results (RI ≥ 0.93 across all tested noise levels) are consistent with — and add a separately reported and decomposed sensitivity analysis to — the qualitative robustness claims found throughout the recent fuzzy-systems literature (Bressane et al., 2024).

Several limitations should be noted. First, the case study uses a synthetic-but-representative dataset rather than field-collected data; validation against a real cohort would strengthen external validity. Second, the overlap term of the Interpretability Index would need to be recomputed after any adaptive update, which the current mechanism does not yet automate; closing this loop is a natural refinement. Third, the rule-base size used here (27 rules) is at the small end of what real-world fuzzy knowledge bases can reach; the RRC compression results reported by Ciaramella (2024) for a 6,215-rule knowledge base illustrate that rule-compression algorithms become a necessary complement, not a substitute, for the IRA-FIS indices at larger scales. Fourth, the robustness challenge has a governance dimension that extends beyond engineering: in domains such as global cybersecurity (Paya Santos et al., 2026), model stability under normative parameter variation is a policy requirement.

On generalizability: because the three IRA-FIS indices are defined purely in terms of the rule base, the membership functions, and a perturbation/observation budget, the framework is directly transferable to clinical risk stratification, financial risk and portfolio decision support, and supply-chain or infrastructure risk assessment. The environmental health case study is one instantiation of a generalizable procedure, not its sole intended domain.

6. Future Directions

Type-2 and type-3 extensions. Extending IRA-FIS to interval or general type-2 fuzzy sets would allow the framework to model uncertainty about the membership functions themselves; this would require redefining the overlap and drift terms in the Interpretability and Adaptability indices over the additional footprint-of-uncertainty dimension.

1. Constrained data-driven rule learning. Combining the bounded-adaptation philosophy of Section 3.3 with ANFIS-style parameter learning (Jang, 1993), but subject to the same expert-approved drift bounds and with an automatic interpretability-penalty term in the training objective, could narrow the interpretability–adaptability trade-off identified throughout this paper.
2. Rule-base compression for scale. Integrating a rule-compression step such as the RRC algorithm (Ciaramella, 2024) prior to computing the Interpretability Index would allow IRA-FIS to scale to the much larger knowledge bases seen in operational environmental and clinical systems.
3. Standardized benchmarking. Packaging the IRA-FIS computation procedure (Section 3.6) as an open benchmarking protocol would enable the head-to-head numerical comparison that Section 4.6 explicitly leaves as qualitative, responding to explicit calls for standardized metrics in the fuzzy machine-learning literature (Bressane et al., 2024).
4. Field validation. Applying IRA-FIS to real, longitudinal datasets — in environmental health, cybersecurity governance (Paya Santos et al., 2026), or criminal intelligence analysis (Martino et al., 2026) — would test the Adaptability mechanism under realistic, rather than illustrative, evidentiary triggers.

7. Conclusion

This paper introduced IRA-FIS, a framework that defines and jointly evaluates interpretability, adaptability, and robustness — three properties widely recognized as important for the real-world adoption of fuzzy decision-making systems, but rarely reported together for a single design. By formalizing each property as a computable index over the rule base, membership functions, and a perturbation or observation budget, and by combining them into a stakeholder-configurable composite score, IRA-FIS gives designers of fuzzy decision-support systems a common vocabulary and a reproducible procedure for comparing competing designs on the same problem. The environmental health risk-assessment case study demonstrated the procedure end-to-end on a 27-rule Mamdani system, yielding II = 0.588, AI = 0.65 under a realistic calibration-shift scenario, and $R\Gamma = 0.959$, with graceful degradation also observed under simulated rule-firing perturbation. We position this contribution as directly responsive to this Collection's call for algorithmic developments and practical applications that improve the interpretability, adaptability, and robustness of fuzzy decision-making methods, and we have outlined concrete extensions — toward type-2/type-3 fuzzy sets, constrained data-driven rule learning, rule-base compression, standardized benchmarking, and field validation across domains including cybersecurity governance (Paya Santos et al., 2026) and criminal intelligence analysis (Martino et al., 2026) — that we hope will be taken up by the broader fuzzy decision-making research community.

References

Alizadehsani, R., Roshanzamir, M., Hussain, S., Khosravi, A., Koohestani, A., Zangooei, M. H., Abdar, M., Beykikhoshk, A., Shoeibi, A., Zare, A., Panahiazar, M., Nahavandi, S., Acharya, U. R., & Acharya, R. (2024). Handling of uncertainty in medical data using machine learning and probability theory techniques: A review of 30 years (1991–2020). Annals of Operations Research, 339, 1077–1118. https://doi.org/10.1007/s10479-021-04006-2

Bressane, A., Garcia, A. J. da S., de Castro, M. V., Xerfan, S. D., Ruas, G., & Negri, R. G. (2024). Fuzzy machine learning applications in environmental engineering: Does the ability to deal with uncertainty really matter? Sustainability, 16(11), 4525. https://doi.org/10.3390/su16114525

Ciaramella, A. (2024). On the interpretability of fuzzy knowledge base systems. PeerJ Computer Science. https://doi.org/10.7717/peerj-cs.2093

ff4ERA: A new fuzzy framework for ethical risk assessment in AI [Preprint]. (2025). arXiv. https://arxiv.org/abs/2502.00451

Fuzzy inference system with interpretable fuzzy rules for disease diagnosis: A comprehensive review. (2024). Information Sciences. https://doi.org/10.1016/j.ins.2024.120292

Fuz-RL: A fuzzy-guided robust framework for safe reinforcement learning under uncertainty [Preprint]. (2026). arXiv. https://arxiv.org/abs/2601.04512

Fuzzy logic model for informed decision-making in risk assessment during software design. (2025). Systems, 13(9), 825. https://doi.org/10.3390/systems13090825

International Journal of Imaging Systems and Technology. (2025). Deep fuzzy inference system for interpretable multi-class heart disease risk prediction. International Journal of Imaging Systems and Technology, 35(2). https://doi.org/10.1002/ima.23211

Jang, J.-S. R. (1993). ANFIS: Adaptive-network-based fuzzy inference system. IEEE Transactions on Systems, Man, and Cybernetics, 23(3), 665–685. https://doi.org/10.1109/21.256541

Kahraman, C., & Gündogdu, F. K. (2021). Decision making with spherical fuzzy sets. In Studies in fuzziness and soft computing (Vol. 392, pp. 3–25). Springer.

Kalibatienė, D., & Miliauskaitė, J. (2021). A hybrid systematic review approach on complexity issues in data-driven fuzzy inference systems development. Informatica, 32(1), 85–118. https://doi.org/10.15388/21-INFOR442

Lopez, L., Elsharief, S., Jorf, D. A., Darwish, F., Ma, C., & Shamout, F. E. (2025). Uncertainty quantification for machine learning in healthcare: A survey (arXiv:2505.02874). arXiv. https://arxiv.org/abs/2505.02874

Mamdani, E. H. (1974). Application of fuzzy algorithms for control of simple dynamic plant. Proceedings of the IEE, 121(12), 1585–1588. https://doi.org/10.1049/piee.1974.0328

Martino, L., Payá-Santos, C. A., Sanz-González, R., & Rodriguez-Gonzalez, V. (2026). Techniques and processes of traditional criminal intelligence analysis. In L. A. Garcia-Segura (Ed.), Collaborative insights on technology and legal frameworks. SOSTEL 2024. Advanced Sciences and Technologies for Security Applications. Springer. https://doi.org/10.1007/978-3-032-27714-5_7

Mota, M. T. D., Bressane, A., Roveda, J. A. F., & Roveda, S. R. M. M. (2019). Classification of successional stages in Atlantic forests: A methodological approach based on a fuzzy expert system. Ciência Florestal, 29(2), 519–530. https://doi.org/10.5902/1980509825384

Padhy, S. K., Mohapatra, A., & Patra, S. (2025). Enhanced hybrid framework for detection of cardiovascular disease based on adaptive neural network and fuzzy inference system. Procedia Computer Science, 258, 2659–2670. https://doi.org/10.1016/j.procs.2025.02.098

Payá Santos, C. A., Delgado Morán, J. J., Martino, L., García Segura, L. A., Diz Casal, J., & Fernández Rodríguez, J. C. (2023). Fuzzy logic analysis for managing uncertain situations. Review of Contemporary Philosophy, 22(1), 6780–6797. https://doi.org/10.52783/rcp.1132

Paya Santos, C., Sanz González, R., Canorea-García, R., Fernández-Rodríguez, J. C., Diz-Casal, J., & Domínguez Pineda, N. Z. (2026). Global cybersecurity governance: Challenges in harmonizing international cyber laws. Journal of Intelligent Decision Making and Information Science, 3. https://doi.org/10.59543/jidmis.v3.972

Risk assessment of infrastructure using a modified adaptive neurofuzzy system: Theoretical application to sewer mains. (2024). Journal of Pipeline Systems Engineering and Practice, 15(2). https://doi.org/10.1061/JPSEAD.PSENG-1497

Santos, C. P., Martino, L., Gonzalez, R. S., & Moran, J. J. (2026). Technological Gaps and Security: Challenges for Law Enforcement and Public Safety. In S. Cea & T. Aznar (Eds.), AI Influence on Governance and Law in the Digital Age (pp. 169-192). IGI Global Scientific Publishing. https://doi.org/10.4018/979-8-3373-4480-5.ch007

Yang, D., & Malik, M. S. A. (2025). Design of performance evaluation method for higher education reform based on adaptive fuzzy algorithm. PeerJ Computer Science. https://doi.org/10.7717/peerj-cs.2498

Zadeh, L. A. (1965). Fuzzy sets. Information and Control, 8(3), 338–353. https://doi.org/10.1016/S0019-9958(65)90241-X

Zadeh, L. A. (1975). The concept of a linguistic variable and its application to approximate reasoning—I. Information Sciences, 8(3), 199–249. https://doi.org/10.1016/0020-0255(75)90036-5

Zadeh, L. A. (2023). Fuzzy logic. In Granular, fuzzy, and soft computing (pp. 19–49). Springer. https://doi.org/10.1007/978-1-0716-2628-3_345