International Journal of Engineering and Applied Computing
Vol. 1, pp. 1-16, June 2026

RESEARCH ARTICLE

Goal-Oriented Software Requirements Elicitation under Uncertainty: A Rough-Set Approximation Approach

Mohammad Asim[1]. Rinky Ahuja[1]. Farhana Mariyam[2]. Mohammad Salman[3]

[1]School of Engineering and Technology, Sushant University, Gurugram, India
E-mail: mohammad.asim.phd@gmail.com, rinkiahuja@sushantuniversity.edu.in
[2]School of Business, Sushant University, Gurugram, India
E-mail: farhanamariyam@sushantuniversity.edu.in
[3]Department of Informatics, University of Zurich, Switzerland
E-mail: salman.mohammad@uzh.ch

Corresponding author: mohammad.asim.phd@gmail.com

Received on: 28/04/2026 | **Revised on**: 09/06/2026 | **Accepted on**: 22/06/2026

ABSTRACT

Goal-oriented methods are widely used for eliciting and analyzing software requirements. In these methods, requirements are derived from organizational goals through iterative refinement and decomposition until the responsibility of each lowest-level sub-goal is assigned to a specific agent or system. However, when stakeholders express goals in vague, contradictory, or context-dependent language, traditional goal-oriented frameworks such as KAOS and i* lack a principled mechanism for handling such ambiguity. As a result, they either require premature precision or allow uncertainty to propagate silently into the requirements specification. To address this limitation, several fuzzy logic–based approaches have been proposed to strengthen goal-oriented methods, but these approaches often depend on prior knowledge of the data set. Moreover, because decision makers may use different types of fuzzy numbers to represent ambiguity, the selection process can become highly subjective and may lead to an inappropriate requirements set, potentially causing software failure. To overcome these issues in fuzzy logic–driven goal-oriented methods, we propose a Rough-Set Approximation–based Goal-Oriented Requirements Elicitation framework. The applicability of the proposed method is demonstrated using the requirements of an Institute Examination System.

KEYWORDS

Goal-oriented requirements elicitation; Rough-set theory; Uncertainty modeling; Requirements prioritization; KAOS; i*; Approximation spaces; Software requirements engineering; Institute examination ystem

1. INTRODUCTION

The aim of requirements elicitation techniques is to identify different types of requirements such as functional, non-functional, legal requirements so that a system can be developed as per the needs of stakeholders. Goal-oriented methods play an important role in identifying the requirements from the goals and sub-goals of the clients by constructing an AND/OR graph (Dardenne *et al.*, 1983; Bukhsh *et al.*, 2020; Akram *et al.*, 2026). When stakeholders articulate goals through vague, contradictory, or context-dependent language, even the most structured elicitation process produces an inherently ambiguous input, one that downstream methods such as KAOS and i* are architecturally unprepared to handle. The above-mentioned frameworks can be considered strong decomposers for goal statements, but they take for granted one condition that cannot always be guaranteed – the precision of goal statements provided by stakeholders allowing for proper operationalization (Tassneem *et al.*, 2025; Nazim *et al.*, 2024). The issue lies in that the assumed level of precision is quite difficult to meet. Specifically, when referring to the IES project scenario, a stakeholder statement like "the system should be responsive" could easily belong and at the same time does not belong to a requirements baseline, depending on whom from among stakeholders one asks. GORE frameworks, being deterministic tools, leave no room for doubt – either the statement is included or excluded. The traditional practice of dealing with ambiguity has been using fuzzy logic approaches such as Fuzzy-AHP and Fuzzy-TOPSIS. Both techniques have received widespread recognition due to their capability of handling imprecise statements linguistically (Nazim *et al.*, 2022). However, their flexibility implies the predefinition of membership functions that introduce subjectivity into the process of decision making (Babar *et al.*, 2015; Gondal *et al.*, 2025).

Rough-set theory does not pose such a constraint. The formalism developed by Pawlak provides the boundaries for the approximation without relying on either a distribution or a membership function (Pawlak, 1982). Considering this fundamental difference, there is one question to be answered: why are there no cases when rough-set approaches are used for upstream goal categorization, whereas they are extensively utilized at the post-elicitation prioritization step only (Sadiq and Devi, 2023; Mariyam *et al.*, 2023). In order to find the solution to the problem posed by fuzzy-based techniques, RA-GORE (Rough-Set Approximation Goal Oriented Requirements Elicitation) is developed. Uncertainty is regarded as information in the proposed approach. A goal statement can be categorized into any of the three epistemic zones, namely definite, boundary, and highly uncertain depending on the extent of disagreements between stakeholders measured through the width of rough-set boundary intervals. This categorization is not only administrative but also diagnostic since goals included into the lower approximation area characterized by the agreement in stakeholders' opinion about definite inclusion are immediately referred to the KAOS and i* decompositions.

Goals that lie within the border area, where there is some form of difference of opinion, are considered and subjected to specific elicitation before decomposing them. High-uncertainty goals are deliberately not included in any elicitation efforts and postponed until such time when disagreement is settled to avoid propagating ambiguous information within the requirements specification process. Goals falling within the border region are considered for conflict resolution. Goals within the high uncertainty region, on the other hand, are raised even before the elicitation begins. Contributions made by our research are enumerated below:

1. A goal-oriented method has been developed with rough-set based approach, i.e., RA-GORE
2. The RA-GORE has been validated by using the requirements of an IES
3. The results of RA-GORE have been compared with Fuzzy-TOPSIS, Rough-TOPSIS, and AHP. the results confirm the central claim: boundary interval width predicts ranking divergence. Every ranking disagreement between fuzzy and rough methods in our study occurs precisely at requirements where Boundary Interval Width (BI) values are 3–4× the consensus mean, a finding with direct operational implications for requirements engineers working under resource constraints.

In continuation, this research paper is organized as follows. Related work is presented in Section 2. The theoretical background of the proposed work is given in Section 3. The proposed RA-GORE is elaborated in Section 4. The case study using IES is presented in Section 5 to provide an understanding of the RA-GORE method. Comparative analysis with other methods, including fuzzy TOPSIS, rough TOPSIS, and AHP, is given in Section 6. Conclusions and future work are provided in Section 7.

2. LITERATURE REVIEW

Existing Goal-Oriented Software Requirements Analysis (GOSRA) models have been improved by computational intelligence approaches, which seek to tackle uncertainty using fuzzy and rough sets theory, machine learning, clustering, and optimization (Robinson, 1989; Lamsweerde *et al.*, 2003; Amyot *et al.*, 2010; Horkoff *et al.*, 2019; Kaiya *et al.*, 2002). Intelligent GOSRA (Horkoff et al., 2019) approaches have been developed particularly for uncertainty, impreciseness, and incompleteness. As an example, Sadiq and Jain (2014) developed a model that combined Analytic Hierarchy Process (AHP) with a binary sort tree approach to determine requirements of IESs. Their focus was to strengthen the approach, and towards this, they adopted a strategy of utilizing α-level weighted F-preference relation to handle uncertainty in the decision-making process. In another work, Sadiq and Jain (2015) developed a model to solve goal selection problems by proposing Extended Attributed Goal-Oriented Requirements Analysis (AGORA), employing fuzzy logic approach. Based on PRFGOREP (Sadiq et al., 2014) and FSGGOREP (Sadiq et al., 2015), Mohammad et al (2017) introduced a Fuzzy-Attributed Goal-Oriented Software Requirements Analysis (FAGOSRA) technique for examining software requirements.

In FAGOSRA, both Fuzzy CV and PM were introduced to analyze goals and requirements related to them. The technique was tested using the IES example and the linguistic variable approach, which was performed with the help of the IFAPM. Another group developed FAGOSRA framework for decision making among many parties and created FAGOSRA_MS as a result. The GOSRA method evaluates functional goals (FG) and non-functional goals (NFGs) by dividing them into sub-goals until system requirements appear. Traditional methods of GOSRA analysis use crisp and fuzzy logic while fuzzy approach is based on membership functions which are known beforehand. Since membership functions are often selected in a subjective way, this process may become non-objective and unreliable for the purpose of goal ranking. In order to overcome these disadvantages, Mariyam et al. (2023) introduced a rough set theory–based method that uses a rough preference matrix to represent stakeholder judgments with greater fidelity. Existing studies, however, devote little attention to the systematic elicitation of testable requirements (TReq). Motivated by this gap, Asim et al. (2026) introduced an approach for deriving TReq early in the lifecycle, encouraging tight cooperation between requirements engineers and testers during the opening stages of requirements engineering and aiding the identification of a complete and consistent requirement set.

The critical review of the GORE literature revealed two principal gaps (Yu, 1997; Perini *et al.*, 2009; Shao *et al.*, 2017; Saxena *et al.*, 2025; Asim and Ahuja, 2026). First, none of the available frameworks employs rough-set approximation spaces as a classification mechanism during the elicitation phase itself—in other words, none partitions goal statements into a lower approximation (definite goals), a boundary region (contested goals), and a complement (excluded goals); instead, every collected goal is accepted as an equally valid input. Second, present-day GORE frameworks provide no structured procedure for dealing with boundary-region goals, that is, goals that can be neither definitively accepted nor definitively rejected given the stakeholder information available. Such situations are typically settled by the analyst's personal judgment, thereby reintroducing the very subjectivity that rough-set methods are meant to eliminate. To address these gaps, this paper develops a RA-GORE method, and the requirements of an IES are used to illustrate the steps of the proposed approach.

3. THEORETICAL BACKGROUND

This section sets out the two mathematical foundations on which RA-GORE rests and clarifies exactly how they connect—a distinction that is important for interpreting the framework's outputs correctly. Section 3.1 introduces Pawlak's classical rough-set theory, which supplies the framework's structural logic: the notion that knowledge can be divided into what is known with certainty, what is known only as a possibility, and what remains genuinely undecidable given the information available. Section 3.2 then presents the rough-number construction of Zhai *et al.* (2008), which puts that notion into practice using ordinal stakeholder judgments. Because the two formalisms play distinct roles within RA-GORE, they should not be treated as interchangeable.

3.1 Rough-Set Theory

The central insight of rough-set theory, introduced by Pawlak (1982), is deceptively simple: if two objects are indistinguishable based on available information, no classifier has grounds to treat them differently. Rather than imposing an artificial resolution, as probability theory does by assigning likelihoods, or fuzzy theory by assigning membership grades, rough-set theory makes the in distinguishability itself a first-class object. This is precisely what makes it suited to requirements elicitation, where forcing false precision is not a minor convenience but a source of downstream specification failure. In RST, imprecise or uncertain concepts are described through two precise sets known as the lower and upper approximations (Pawlak, 1982), as depicted in Fig. 1.

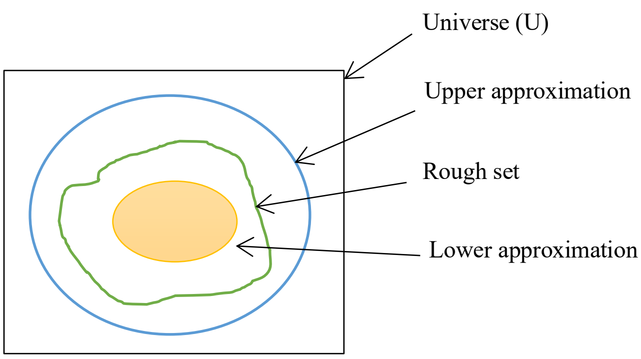

Fig. 1: Schematic demonstration of RST

Formally, a rough-set information system is defined as IS = (U, A, V, f), where U is a non-empty finite universe of objects, A is a set of condition attributes, V is the domain of attribute values, and f: U × A → V assigns attribute values to objects. For any subset of attributes B ⊆ A, an indiscernibility relation IND(B) partitions U into equivalence classes, groups of objects that are informationally identical with respect to B. The coarser the attribute set, the larger these equivalence classes, and the more uncertainty remains. For any target concept $X \subseteq U$, in RA-GORE, a candidate requirement class, rough-set theory defines three regions that together partition the universe without ambiguity:

Lower Approximation: $\underline{B}(X) = \{x \in U \mid [x]\_B \subseteq X\}$ (1)

The set of objects that certainly belong to X: every object in their equivalence class falls inside X. These are the definite requirements, no additional information can change their classification.

Upper Approximation: $\overline{B}(X) = \{x \in U \mid [x]\_B \cap X \neq \emptyset\}$ (2)

The set of objects that possibly belong to X: at least one member of their equivalence class falls inside X. These represent the outer boundary of what available information permits.

Boundary Region: $BN\_B(X) = \overline{B}(X) - \underline{B}(X)$ (3)

The set of objects that cannot be classified with certainty given current information. This is the intellectually critical region for RA-GORE. A non-empty boundary region is not a failure of the method, it is an honest representation of what the data actually supports. The goal is not to eliminate it by assumption but to understand why it exists and what additional information would resolve it. The accuracy of approximation α_B(X) = $|\underline{B}(X)| / |\overline{B}(X)|$ quantifies how much of what is possibly X is also certainly X. When $\alpha = 1$, the concept is crisp, the information system contains enough discriminating attributes to classify every object definitively. When $\alpha < 1$, the concept is rough, genuine epistemic residue remains that no reweighting or rescaling can resolve without new information.

3.2 Rough Numbers for Stakeholder Assessment Aggregation

Pawlak's formalism operates on categorical attributes and exact equivalence relations , assumptions that hold cleanly in formal information systems but break down immediately in requirements elicitation, where stakeholders provide ordinal assessments on linguistic scales that carry inherent imprecision. Rough numbers, introduced by Zhai *et al.* (2008) and applied to requirements prioritization by Sadiq and Devi (2023), bridge this gap by extending the rough-set intuition to ordinal, multi-valued data without requiring predefined membership functions.

The key idea is this: rather than representing a stakeholder's assessment of a goal as a single point value, a rough number encodes it as an interval whose width reflects how much disagreement exists among all assessors about objects ranked at or near that value. An assessor who assigns a goal "Good" on a 7-point scale receives a narrow rough number if most other assessors agree, and a wide rough number if assessors are spread across the scale. The interval is not assumed, it is computed from the assessment distribution itself. Formally, for a set of stakeholder assessments $\{x_1, x_2, \ldots, x_k\}$ of a goal attribute drawn from universe U, the rough number of assessment x_i as defined in (Mariyam *et al.*, 2023; Sadiq and Devi, 2022). The boundary interval BI (x_i) = $\overline{Lim\ (x_i)} - \underline{Lim\ (x_i)}$ is the operationally critical output. A large BI does not make the assessor wrong; it indicates that the stakeholder community lacks convergence on an understanding of that particular goal attribute. It is the very signal RA-GORE capitalizes on – broad boundary intervals indicate, prior to ranking the goal statement in question, that this requirement requires additional elicitation rather than analysis.

4. THE PROPOSED RA-GORE FRAMEWORK

Unlike most requirements engineering frameworks, the RA-GORE framework starts from an implicit assumption, namely, that regardless of how chaotic the process of elicitation is, there always comes a moment when a list of goals emerges, which may be considered stable and consistent input data. RA-GORE framework explicitly challenges this approach. With its five-stage framework, it does not operate with goal statements; instead, it operates with their qualification, to determine, even prior to any decomposition or prioritization, which goals are supported by the available stakeholder data. As such, it has a certain advantage over other frameworks. Unlike traditional approaches to requirements engineering, RA-GORE deals with uncertainty at the stage of goal elicitation itself. Of course, this approach implies additional procedural work at stages 3 and 4, but every goal qualified at stage 5 is a clear requirement rather than just a hypothetical one rushed through lack of time. Fig. 2 depicts the complete process flow. The five phases are detailed below, with emphasis on the design decisions that set RA-GORE apart from conventional sequential elicitation approaches.

Phase 1: Goal Elicitation

Stakeholder goals are collected through structured interviews, requirements workshops, and document analysis. The defining constraint of Phase 1 is deliberate: goal statements are recorded without enforced precision. Analysts do not resolve ambiguities on the spot, paraphrase vague statements into formal goal language, or adjudicate between conflicting stakeholder expressions. This restraint is not passivity, it is a methodological commitment. Premature precision at the elicitation stage does not eliminate uncertainty; it conceals it, pushing the problem downstream where it is harder to detect and more expensive to correct. Phase 1 of RA-GORE creates the raw material that the rough-set formalism in Phases 2–3 is specifically designed to handle.

Phase 2: Information System Construction

Raw goal statements are encoded as objects in a Pawlak information system IS = (U, A, V, f), where each goal becomes an object in U and its characteristics become attribute values in A. Five elicitation attributes are defined: Stakeholder Priority , aggregated ordinal importance rating across stakeholders; Conflict Score , degree of inter-stakeholder disagreement on the goal; Goal Type , functional (FR) or non-functional (NFR); Decomposition Depth , number of refinement levels required to reach operationalizable sub-goals; and Operationalizability , whether the goal can be directly translated into verifiable system behavior. The selection of these five attributes is not arbitrary. Together they cover the three dimensions that determine whether a goal is elicitation-ready: what stakeholders want (priority, goal type), whether they agree on it (conflict score), and whether it can be acted on (decomposition depth, operationalizability). Attributes that cannot be assessed from stakeholder input alone, implementation feasibility, cost estimation, technical architecture, are deliberately excluded from Phase 2, as they belong to design rather than elicitation.

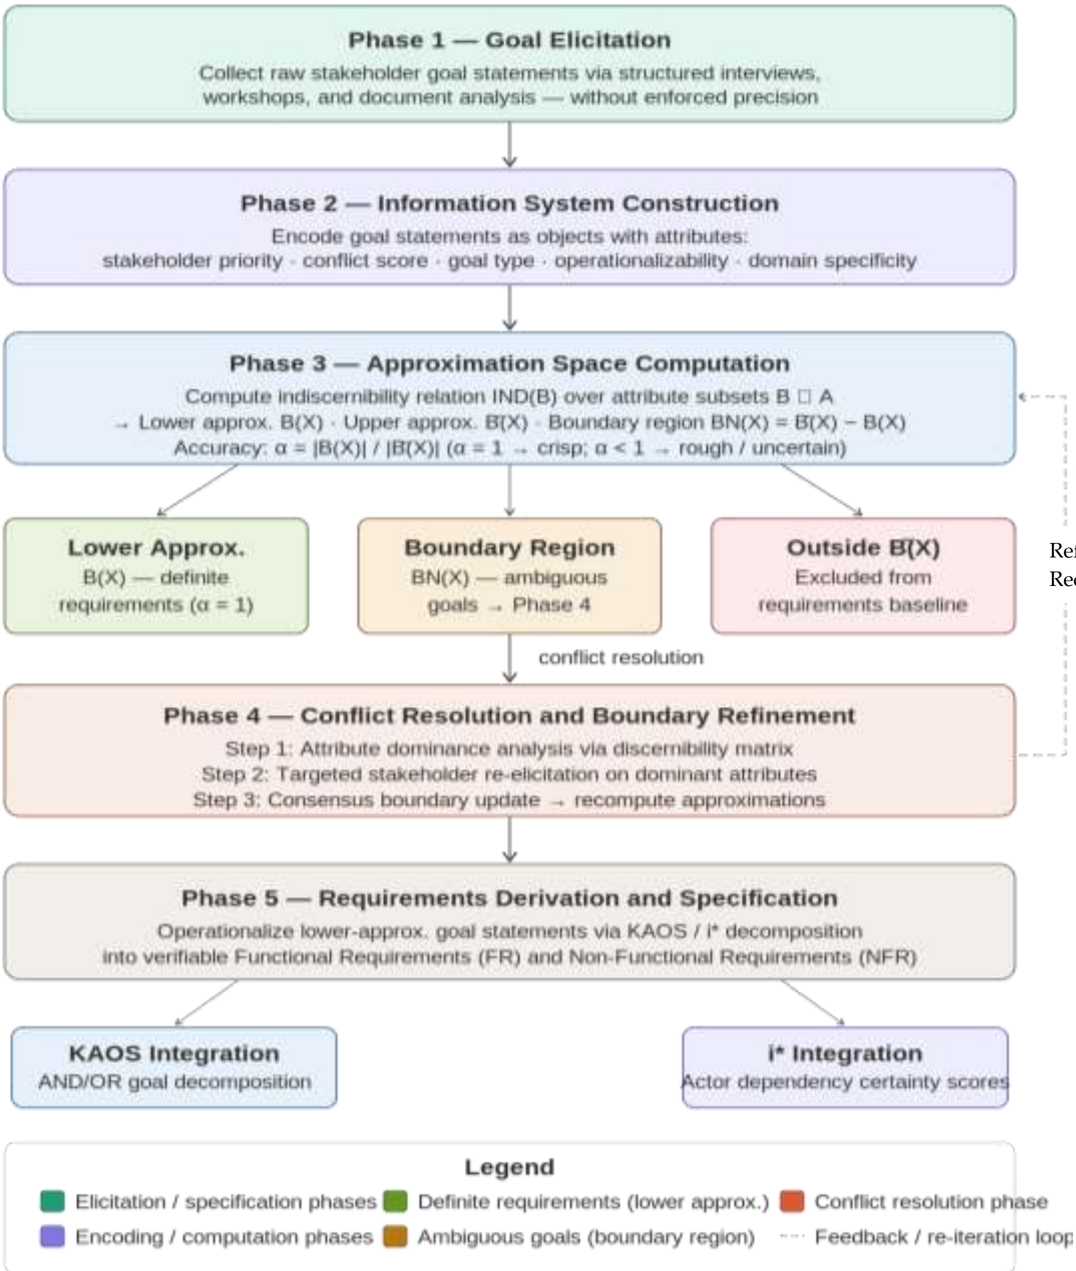

Fig. 2: Block diagram of the proposed method

Phase 3: Approximation Space Computation

For each candidate requirement concept X, the indiscernibility relation IND(B) is computed over selected attribute subsets B ∈ A, and lower approximation $\underline{B}(X)$, upper approximation $\overline{B}(X)$, and boundary region BN_B(X) are derived as defined in Section 3.1. Approximation accuracy α_B(X) is computed for each concept. The output of Phase 3 is a tripartite classification: Lower Approximation (α = 1), goal statements that all available stakeholder information confirms as definite requirements, proceeding directly to Phase 5; Boundary Region (α < 1, above exclusion threshold), goal statements that available information neither confirms nor excludes, routed to Phase 4; and Complement of Upper Approximation, goal statements that available information definitively excludes from the requirements baseline. The boundary region is the critical output, not a failure of the computation but its most informative result. A wide boundary region means the information system genuinely cannot resolve the classification, which is exactly the condition that warrants further stakeholder engagement rather than analyst judgment.

Phase 4: Conflict Resolution and Boundary Refinement

Phase 4 exists because boundary-region goals cannot be resolved by computation alone, they require structured human intervention, conducted in a way that minimizes the reintroduction of subjectivity. The protocol operates in three steps, ordered deliberately:

Step 4.1: Attribute Dominance Analysis

The discernibility matrix is computed to identify which attributes are responsible for the indiscernibility, i.e., which attribute values, if clarified, would most reduce the boundary region. This step determines where to focus re-elicitation rather than returning to stakeholders with open-ended questions that would generate more noise.

Step 4.2: Targeted Re-elicitation

Structured clarification prompts targeting the dominant attributes identified in Step 1 are presented to stakeholders. The goal is attribute-level precision on the dimensions that matter, not full goal restating. Constraining the scope of re-elicitation is a deliberate design choice: unconstrained re-elicitation tends to introduce new ambiguities faster than it resolves existing ones.

Step 4.3: Consensus Boundary Update

Revised attribute values are incorporated into the information system and approximations are recomputed. If the goal statement moves into the lower approximation, it is cleared for Phase 5. If it remains in the boundary region after two re-elicitation cycles, it is escalated, with full boundary region documentation, to a requirement review committee for organizational-level resolution. Two cycles is the operational threshold; further iteration without new stakeholder information does not reduce the boundary interval and risks analyst rationalization replacing genuine consensus.

Phase 5: Requirements Derivation and Specification

Lower-approximation goals, those that passed through Phase 3 directly or cleared the boundary through Phase 4 resolution, are operationalized into formal functional and non-functional requirements using KAOS AND/OR goal decomposition patterns. The epistemic status of the goals is transferred to requirements: baseline requirements are tagged for definite goals while conditional requirements will be tagged for resolved goals, with their respective BI values retained as an indicator of risk.

5. CASE STUDY

The Institute Examination System (IES) was selected as the validation environment for RA-GORE because it closely reflects the characteristics for which the framework was designed. First, IES involves multiple stakeholder groups, including examination controllers, faculty members, students, and administrative staff, each possessing distinct and often conflicting priorities. For example, examination controllers emphasize scheduling accuracy and auditability, students focus on result transparency and grievance handling, while administrators prioritize system reliability and performance. Such differences lead to varying evaluations of the same requirements, making rough-set boundary approximations particularly effective for capturing uncertainty and disagreement. Second, IES encompasses a balanced mix of functional and non-functional requirements, representing the complexity commonly observed in real-world information systems. This enables the assessment of RA-GORE's ability to analyze and prioritize diverse requirement types within a unified framework. Third, IES is an operational system actively used in a multi-department university rather than a synthetic benchmark designed for experimental convenience. Consequently, stakeholder assessments reflect genuine institutional priorities and conflicts, including situations where consensus is absent. These conditions provide a realistic and rigorous setting for evaluating RA-GORE's capability to support requirement analysis under uncertainty and stakeholder divergence. The evaluation took into account ten proposed requirements (R1-R10), which were the most important functional and non-functional requirements of the IES. The selection process was conducted using five participants chosen from among stakeholders who had been identified previously as examination controllers, faculty members, students, and administrative staff.

The panel size was determined by the institutional context of IES and aligns with the stakeholder samples employed in previous rough-set-based requirements engineering studies (Sadiq and Devi, 2023). While the selected panel provided sufficient diversity to capture differing stakeholder perspectives, a larger and more heterogeneous stakeholder group would further enhance the external validity and generalizability of the findings. Expanding stakeholder participation therefore remains an important direction for future validation studies. Stakeholders evaluated each requirement using a seven-point linguistic scale comprising Very Poor (VP = 1), Poor (P = 2), Medium Poor (MP = 3), Fair (F = 4), Medium Good (MG = 5), Good (G = 6), and Very Good (VG = 7). The seven-point scale was selected because it offers a suitable balance between discrimination and usability. Compared with three-point or five-point scales, it provides greater granularity for capturing variations in stakeholder perceptions and identifying meaningful differences in requirement evaluations. At the same time, the scale remains sufficiently intuitive for participants from diverse, non-technical backgrounds. Furthermore, the odd-numbered structure incorporates a neutral midpoint (Fair = 4), enabling stakeholders to express balanced judgments without being compelled toward either a positive or negative assessment.

Each requirement was evaluated against five quality criteria: Functionality (Fu), Reliability (Re), Usability (Us), Cost (Co), and Strategic Importance (SI). These criteria were selected because they represent the most frequently reported dimensions in requirements prioritization research and align closely with the concerns expressed by stakeholders during the goal elicitation process for the Institute Examination System (IES). Collectively, they capture both technical and organizational aspects that influence requirement significance and implementation decisions. For the purpose of the baseline analysis, all criteria were modeled as benefit attributes, where higher values indicate greater desirability.

To facilitate a rigorous and unbiased evaluation, three prioritization methods were applied to the same stakeholder assessment dataset. This ensured that any observed differences in outcomes could be attributed to the underlying decision-making approaches rather than variations in input data. Fuzzy TOPSIS was selected as the primary benchmark because it is one of the most widely adopted techniques for handling linguistic assessments in requirements prioritization. The method employs triangular fuzzy numbers (TFNs) to represent and aggregate stakeholder judgments and has been extensively validated in requirements engineering and multi-criteria decision-making applications (Saxena *et al.*, 2025; Nazim *et al.*, 2022).

Rough TOPSIS, based on the rough-number framework proposed by Zhai *et al.* (2008), was included as a rough-set-based alternative. Unlike Fuzzy TOPSIS, which relies on fuzzy membership functions, Rough TOPSIS captures uncertainty through lower and upper approximation intervals derived directly from stakeholder assessments. This makes it an appropriate methodological counterpart for evaluating the uncertainty-handling mechanisms incorporated within RA-GORE. AHP Weighted Scoring was used as a conventional multi-criteria decision-making baseline. Its inclusion provides a reference point for assessing whether the rankings generated by rough-set and fuzzy approaches differ substantially from those obtained through a simpler and well-established prioritization technique. The comparison among these methods was designed to address two key research questions. First, it examines whether the zone classifications generated by RA-GORE can effectively identify situations in which fuzzy and rough-set approaches are likely to produce different prioritization outcomes. This directly evaluates the framework's diagnostic capability. Second, it investigates whether fuzzy and rough-set methods yield comparable rankings when stakeholder consensus is relatively high, while producing noticeable differences under conditions of significant disagreement. Demonstrating such behavior would provide empirical justification for the additional analytical effort associated with rough-set-based elicitation and prioritization.

Table 1 summarizes the ten candidate requirements considered in the IES case study. The requirements were intentionally selected to represent both functional and non-functional aspects of the system. Functional requirements (R1–R5 and R10) capture the core operational processes of the examination system, whereas non-functional requirements (R6–R9) address quality-related concerns such as reliability, usability, and performance, where stakeholder perspectives often vary more substantially.

Table 1: IES Candidate Requirements

Req. ID	Description	Type
R1	Exam scheduling and timetable generation	FR
R2	Student result entry and processing	FR
R3	Automated invigilation assignment	FR
R4	Grievance management portal	FR
R5	Legacy system data migration	FR
R6	Role-based access control	NFR
R7	System performance under peak load	NFR
R8	Audit trail and logging	NFR
R9	Multi-language interface support	NFR
R10	Third-party integration for SMS alerts	FR

The requirement descriptions were intentionally maintained at the goal level and were not translated into detailed formal specifications. This reflects the typical state of requirements during the early stages of elicitation, where stakeholders communicate needs, expectations, and objectives in natural language before they are refined into precise system specifications. Consequently, the dataset represents the form of input processed by Phase 2 of RA-GORE, which operates on stakeholder-expressed goals prior to formalization. The distribution of six functional requirements (FRs) and four non-functional requirements (NFRs) was not artificially imposed but emerged from the actual requirements profile of the IES. This composition provides a realistic basis for investigating whether requirement type influences the degree of uncertainty captured by rough-set boundary intervals. The relationship between requirement category and boundary interval width is examined in detail in Section 6.

6. COMPARATIVE STUDY

Four evaluation metrics were employed to examine different dimensions of agreement, uncertainty, and ranking behavior among the compared methods.

- ***Spearman's Rank Correlation Coefficient*** (ρ) was used to assess the overall similarity between ranking outcomes generated by different methods. This metric evaluates the extent to which two methods preserve the same relative ordering of requirements across all ranking positions and is particularly sensitive to changes in rank placement throughout the entire list.

- *Kendall's Tau Coefficient (τ)* was used to measure pairwise ranking consistency. Unlike Spearman's ρ, which focuses on the overall ordering pattern, Kendall's τ quantifies the proportion of requirement pairs that are ranked in the same order by two methods. The combined use of these two metrics provides a more comprehensive assessment of ranking agreement. For example, a high Spearman's ρ accompanied by a comparatively lower Kendall's τ may indicate that disagreements are concentrated among specific requirements rather than being evenly distributed across the ranking. Reporting both measures therefore enables a more detailed analysis of ranking similarities and differences.

- ***Boundary Interval Width (BI)*** represents the average width of rough-number intervals associated with each requirement, aggregated across all stakeholders and evaluation criteria. This metric serves as the primary uncertainty indicator within RA-GORE. The central hypothesis of the framework is that wider boundary intervals correspond to greater stakeholder disagreement and, consequently, a higher likelihood of ranking differences between prioritization methods. While Spearman's ρ and Kendall's τ quantify agreement, BI is intended to explain the sources of disagreement.

- ***Ranking Divergence*** (|F-Rank − R-Rank|) measures the absolute difference between the positions assigned to a requirement by Fuzzy TOPSIS and Rough TOPSIS. This metric functions as the dependent variable in the boundary interval analysis. According to the diagnostic rationale of RA-GORE, requirements characterized by larger boundary intervals should exhibit greater ranking divergence, whereas requirements located within the consensus zone should display little or no variation between the two ranking approaches.

Application of the RA-GORE framework to the ten requirements of the Institute Examination System (IES) resulted in a three-zone classification based on the average Boundary Interval (BI) width. To categorize the requirements according to the degree of stakeholder uncertainty, two threshold values were established. Requirements with BI values below 0.60 were assigned to the consensus zone, representing cases with minimal uncertainty and strong stakeholder agreement. Requirements with BI values ranging from 0.60 to 1.50 were classified into the boundary zone, indicating moderate levels of disagreement. Requirements exceeding a BI value of 1.50 were placed in the high-uncertainty zone, reflecting substantial divergence in stakeholder assessments. These threshold values were derived empirically from the observed BI distribution within the IES dataset, where distinct clusters emerged naturally. Nevertheless, the selection of threshold values may vary across domains and should be further examined through sensitivity analysis in future studies.

The resulting classification reveals a meaningful pattern of stakeholder consensus and disagreement. Three requirements—R3 (BI = 0.48), R6 (BI = 0.52), and R1 (BI = 0.57)—were categorized within the consensus zone. Their relatively narrow boundary intervals indicate a high degree of stakeholder agreement and limited classification uncertainty. In accordance with the RA-GORE process, these requirements can advance directly to the specification stage in Phase 5 without requiring additional conflict-resolution activities. Five requirements—R8, R2, R4, R10, and R5—fell within the boundary zone, with BI values ranging from 0.61 to 0.89. These requirements exhibit moderate uncertainty, suggesting the presence of differing stakeholder viewpoints that warrant further examination. Rather than being accepted or rejected outright, they are subjected to the structured conflict-resolution procedures defined in Phase 4 of the framework.

The highest levels of uncertainty were observed for R9 (BI = 1.90) and R7 (BI = 1.99). The boundary intervals associated with these requirements are approximately three to four times larger than the average BI values observed in the consensus zone, indicating substantial stakeholder disagreement. Notably, both requirements belong to the non-functional category, supporting the widely recognized observation that non-functional requirements often generate greater ambiguity due to their subjective interpretation and less clearly defined acceptance criteria. Consequently, these requirements were identified as candidates for targeted re-elicitation and stakeholder negotiation before any prioritization or specification decisions were finalized. Table 2 presents the complete zone classification ordered by increasing BI values. The distribution itself provides additional insight into the uncertainty structure of the requirements. The consensus-zone requirements form a compact cluster between 0.48 and 0.57, while the boundary-zone requirements occupy an intermediate range between 0.61 and 0.89. In contrast, the BI values for R9 and R7 increase sharply to 1.90 and 1.99, respectively. This pronounced separation between the boundary and high-uncertainty groups suggests the presence of a distinct discontinuity rather than a gradual transition, thereby supporting the use of a three-zone classification scheme based on two threshold values.

Table 2: RA-GORE Approximation Classification of IES Requirements

Req.	Avg BI	Zone	Action
R3	0.4800	Lower Approx.	Definite , pass to specification
R6	0.5240	Lower Approx.	Definite , pass to specification
R1	0.5680	Lower Approx.	Definite , pass to specification
R8	0.6120	Boundary Region	Re-elicit stakeholder input
R2	0.6893	Boundary Region	Re-elicit stakeholder input
R4	0.7333	Boundary Region	Re-elicit stakeholder input
R10	0.8160	Boundary Region	Re-elicit stakeholder input
R5	0.8933	Boundary Region	Re-elicit stakeholder input
R9	1.9047	High Uncertainty	Escalate to review committee
R7	1.9940	High Uncertainty	Escalate to review committee

The operational implication is direct: a requirements engineer using RA-GORE on this dataset would invest no further elicitation effort on R3, R6, and R1, concentrate structured re-elicitation on the five boundary-region requirements, and escalate R9 and R7 to an organizational review committee before committing them to any specification baseline. This is a qualitatively different workflow from treating all ten requirements as equally elicitation-ready, which is what conventional GORE frameworks implicitly assume. Table 3 presents the closeness coefficients (for Fuzzy and Rough TOPSIS) and weighted scores (for AHP) alongside the resulting rank positions for all ten requirements. The reader should focus on two features: first, whether the top and bottom positions are stable across methods; and second, whether mid-range positions, the operationally contested territory where release decisions are actually made, show method-dependent variation.

Table 3: Comprehensive Ranking Comparison, Fuzzy TOPSIS, Rough TOPSIS, and AHP

Req.	Fuzzy CC	F-Rank	Rough CC	R-Rank	AHP Score	A-Rank
R1	0.7099	2	0.8839	2	6.0754	2
R2	0.5886	5	0.6558	6	4.9934	6
R3	0.7542	1	1.0000	1	6.5669	1
R4	0.4728	8	0.4763	8	4.1536	8
R5	0.2870	9	0.2071	9	2.8795	9
R6	0.6659	3	0.7886	3	5.6265	3
R7	0.5845	6	0.6743	4	5.1205	4
R8	0.5925	4	0.6646	5	5.0132	5
R9	0.5684	7	0.6452	7	4.9563	7
R10	0.1657	10	0.0000	10	1.9497	10

The pattern is immediately visible. Positions 1–3 (R3, R1, R6) and positions 8–10 (R4, R5, R10) are identical across all three methods. The extremes are settled, no method, regardless of its uncertainty formalism, changes what is clearly most or least important. The action is in positions 4–7, where R7 and R8 swap ranks depending on whether Fuzzy TOPSIS or Rough TOPSIS is used. R7 ranks 6th under Fuzzy TOPSIS but 4th under Rough TOPSIS and AHP, a two-position shift that, in a resource-constrained release, could determine whether the requirement is implemented or deferred. This is precisely the mid-range territory where RA-GORE's diagnostic is designed to provide advance warning. Table 4 should be read with the Table 3 rankings in mind. The high correlation values confirm that the methods produce broadly similar orderings, but the question that matters for RA-GORE is not the overall agreement, it is the structure of the disagreement.

Table 4: Rank Correlation Analysis across Method Pairs

Comparison	Spearman ρ	Kendall τ	p-value (ρ)
Fuzzy vs. Rough TOPSIS	0.9636	0.9111	< 0.001 ($t = 10.19$, $df = 8$)
Fuzzy TOPSIS vs. AHP	0.9636	0.9111	< 0.001
Rough TOPSIS vs. AHP	1.0000	1.0000	, (perfect agreement)

Two findings stand out. First, Rough TOPSIS and AHP produce perfectly identical rankings ($\rho = 1.0$, $\tau = 1.0$) on this dataset. While this supports the internal consistency of the rough-set approach, the small sample size ($n = 10$) means this should be treated as an encouraging signal rather than a generalizable claim , replication on larger datasets is needed before drawing conclusions about rough-AHP convergence. Second, the Spearman ρ of 0.9636 between Fuzzy and Rough TOPSIS is statistically significant ($p < 0.001$), but the gap from 1.0 is itself the story: that residual 0.036 concentrates entirely in the mid-range positions identified in Table 3, not randomly across the ranking.

The previous analyses demonstrated three key findings. First, RA-GORE categorizes requirements into three distinct uncertainty zones based on boundary interval width. Second, the compared prioritization methods exhibit strong agreement for requirements located at the extremes of the uncertainty spectrum, while notable differences emerge for requirements in the intermediate region. Third, although the overall ranking outcomes produced by the methods are highly correlated, the agreement is not complete. Building on these observations, this section investigates whether boundary interval width can serve as a predictive indicator of ranking divergence. A central proposition of RA-GORE is that stakeholder uncertainty, as captured through boundary intervals, can identify requirements that are likely to produce differing prioritization outcomes before ranking calculations are performed. In other words, boundary interval width is expected not only to be associated with ranking differences but also to provide an early diagnostic signal of where such differences are most likely to occur.

Table 5 provides the evidence used to evaluate this proposition. For each requirement, the table reports the average boundary interval width, the corresponding stakeholder agreement category, the absolute ranking difference between Fuzzy TOPSIS and Rough TOPSIS, and the RA-GORE zone assignment. The primary objective is to determine whether ranking divergences are concentrated among requirements characterized by larger boundary intervals, while requirements located in the consensus zone exhibit little or no variation between ranking methods. The results reveal a clear relationship between boundary interval width and ranking divergence. All three requirements classified within the consensus zone (R3, R6, and R1) exhibit identical rankings under both Fuzzy TOPSIS and Rough TOPSIS, resulting in zero ranking divergence. In contrast, the largest discrepancy was observed for R7, which recorded a rank difference of two positions and also possessed the widest boundary interval in the dataset (BI = 1.994). Requirements R8 and R2, located near the lower boundary of the intermediate uncertainty zone, displayed minor divergences of one rank position each. Meanwhile, R4, R10, and R5 showed no ranking differences despite having moderate BI values. This observation suggests that intermediate boundary intervals indicate the potential for divergence rather than guaranteeing it.

Table 5: Boundary Interval Width vs. Ranking Divergence, Diagnostic Analysis

Req.	Avg. BI	Agreement	|F-Rank − R-Rank|	RA-GORE Zone
R3	0.4800	High	0	Lower Approx.
R6	0.5240	High	0	Lower Approx.
R1	0.5680	High	0	Lower Approx.
R8	0.6120	Moderate	1	Boundary Region
R2	0.6893	Moderate	1	Boundary Region
R4	0.7333	Moderate	0	Boundary Region
R10	0.8160	Moderate	0	Boundary Region
R5	0.8933	Moderate	0	Boundary Region
R9	1.9047	Low	0	High Uncertainty
R7	1.9940	Low	2	High Uncertainty

The observed differences can be explained by the distinct ways in which the two methods process stakeholder assessments. Fuzzy TOPSIS combines individual evaluations through the aggregation of triangular fuzzy numbers, producing a single representative value that tends to moderate extreme opinions. As a result, substantial disagreements among stakeholders may be partially masked within the aggregated score. For example, in the case of R7, stakeholder evaluations of functionality ranged from relatively low to very high assessments. The aggregation process therefore generated a moderate overall value that did not fully reflect the extent of disagreement. Rough TOPSIS, on the other hand, retains this variability through lower and upper approximation boundaries. Instead of compressing divergent opinions into a single estimate, it explicitly represents the range of stakeholder assessments through boundary intervals. Consequently, the uncertainty associated with highly contested requirements remains visible throughout the prioritization process. Because the two approaches capture stakeholder disagreement differently, ranking variations are most likely to emerge for requirements characterized by large boundary intervals.

These findings support the diagnostic objective of RA-GORE. By calculating boundary intervals before applying any prioritization technique, requirements engineers can identify requirements that are likely to yield method-dependent rankings. In the present study, R7 and R9 were flagged as high-uncertainty requirements prior to ranking analysis. The purpose of the zone classification is not to determine which ranking outcome is correct, but rather to highlight requirements whose prioritization results should be interpreted with caution until the underlying stakeholder disagreements have been further explored and resolved.

7. CONCLUSION AND FUTURE SCOPE

This study presented RA-GORE, a rough-set-based framework for goal-oriented requirements elicitation that identifies and manages stakeholder uncertainty before requirements decomposition begins. The framework treats uncertainty as valuable information rather than noise, using boundary interval widths to classify goals into consensus, boundary, and high-uncertainty zones. Goals with strong stakeholder agreement proceed directly to further analysis, while those exhibiting substantial disagreement are routed for re-elicitation or conflict resolution. The case study on the IES demonstrated that boundary interval width is a useful indicator of ranking divergence between Fuzzy TOPSIS and Rough TOPSIS. Requirements with large boundary intervals exhibited the greatest ranking differences, whereas consensus-zone requirements maintained stable rankings across methods. These findings suggest that rough-set-based uncertainty measures can provide early diagnostic insights into stakeholder disagreement and prioritization risk.

Despite these promising results, the study is limited by its single-case evaluation involving ten requirements and five stakeholders. Consequently, the proposed thresholds and quantitative findings require validation across larger and more diverse datasets. Future research should focus on industrial-scale evaluations, integration with hybrid fuzzy–rough approaches, incorporation of requirement interdependencies, exploration of intuitionistic fuzzy models for handling indeterminacy, and development of automated tool support. Overall, RA-GORE contributes a practical mechanism for exposing, classifying, and managing uncertainty in requirements engineering, enabling decision-makers to address stakeholder disagreements explicitly rather than concealing them through aggregation-based techniques.

CREDIT AUTHOR STATEMENT

1. **Mohammad Asim:** Conceptualization, Methodology, Writing – Original Draft
2. **Rinky Ahuja:** Supervision, Validation, Writing – Review & Editing
3. **Farhana Mariyam:** Visualization, Formal analysis, Validation, Writing – Review & Editing
4. **Mohammad Salman:** Investigation, Data Curation, Project administration

REFERENCES

Akram F., Ahmad T., and Sadiq M. (2026). A rough-set theory based approach for the detection of requirements discordances among stakeholders, Journal of Systems Science and Systems Engineering, Springer. https://doi.org/10.1007/s11518-026-5727-7

Amyot D., Ghanavati S., Horkoff J. and Mussbacher G. (2010). Evaluating goal models within the goal-oriented requirement language, International Journal of Intelligent Systems, Vol. 25, Issue 8, pp. 841-877. https://doi.org/10.1002/int.20433

Asim M. and Ahuja R. (2026). Intelligent goal-oriented requirements elicitation process: a bibliometric review, 7th World Conference on Artificial Intelligence: Advances and Applications, Springer, National Forensic Sciences University Goa Campus, India, January 30-31, 2026.

Asim M., Ahuja R., and Sadiq M. (2026). Elicitation of testing requirements from the functional requirements in goal-oriented software requirements engineering: a fuzzy-based approach, 4th International Conference on Power Engineering and Intelligent Systems, Springer, Organized by National Institute of Technology Uttarakhand, India, March 14-15, 2026

Babar M. I., Ghazali M., Dayang N.A. Jawawi, Shamsuddin S. M., and Ibrahim N. (2015). PHandler: An expert system for a scalable software requirements prioritization process, Knowledge-Based Systems, Volume 84, pp. 179-202. https://doi.org/10.1016/j.knosys.2015.04.010

Bukhsh F. A., Bukhsh Z. A., and Daneva M (2020). A systematic literature review on requirements prioritization techniques and their empirical evaluation, Computer Standards & Interfaces, 69. https://doi.org/10.1016/j.csi.2019.103389

Dardenne, A., Lamsweerde, A.V., and Fickas S. (1983). Goal-directed requirements acquisition, Science of Computer Programming, 20 (1-2), 3-50. https://doi.org/10.1016/0167-6423(93)90021-G

Gondal, A.N., Kausar, M., Mazhar, N., and Amjad M. (2025). Prioritization of agile requirement elicitation techniques using fuzzy logic, International Journal of Fuzzy System. Vol. 27, pp. 2637–2657. https://doi.org/10.1007/s40815-024-01920-y

J. Horkoff et al. (2016). Goal-Oriented Requirements Engineering: A Systematic Literature Map, IEEE 24th International Requirements Engineering Conference, Beijing, China, pp. 106-115. https://doi.org/10.1109/RE.2016.41

Kaiya H, Horai H. and Saeki M. (2002). AGORA: attributed goal-oriented requirements analysis method, IEEE Joint International Conference on Requirements Engineering, Essen, Germany, pp. 13-22. https://doi.org/10.1109/ICRE.2002.1048501

Lamsweerde, A. V. and Letier, E. (2003). From object orientation to goal orientation: a paradigm shift for requirements engineering. In: Wirsing, M., Knapp, A., Balsamo, S. (eds) Radical Innovations of Software and Systems Engineering in the Future. RISSEF 2002. Lecture Notes in Computer Science, vol. 2941, pp. 325-340, Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-540-24626-8_23

Mariyam F., Mehfuz S., and Sadiq M. (2023). RAGOSRA: Rough attributed goal-oriented software requirements analysis method, Journal of Intelligent & Fuzzy Systems, Vol. 44, Issue 5, pp. 7833–7843. https://doi.org/10.3233/JIFS-221300

Mohammad C. W., Shahid M., and Hussain S. Z. (2021). Fuzzy attributed goal-oriented software requirements analysis with multiple stakeholders, International Journal of Information Technology, (13)1–9. https://doi.org/10.1007/s41870-017-0073-0

Nazim M., Ahmad S., Arif M., Mishra V., and Ahmad J. (2024). Optimizing software requirements selection: a hybrid k-means and intuitionistic fuzzy TOPSIS approach, 13th International Conference on System Modeling & Advancement in Research Trends, Moradabad, India, pp. 750-755. https://doi.org/10.1109/SMART63812.2024.10882593

Nazim M., Mohammad C. W., and Sadiq M. (2022). A comparison between fuzzy AHP and fuzzy TOPSIS methods to software requirements selection, Alexandria Engineering Journal, Vol. 61, Issue 12, pp. 10851-10870. https://doi.org/10.1016/j.aej.2022.04.005

Pawlak Z. (1982). Rough sets, International Journal of Computer and Information Sciences, Vol. 11, Issue 5, pp. 341–356. https://doi.org/10.1007/BF01001956

Perini, A., Susi, A., and Ahluwalia, P. (2009). Tool-supported requirements prioritization: Comparing the AHP and CBRank methods, Information and Software Technology, Vol. 51, Issue 6, pp. 1021-1032. https://doi.org/10.1016/j.infsof.2008.12.001

Robinson, W. N. (1989). Integrating multiple specifications using domains goals, 5th International Workshop on Software Specification Design, pp. 219-226. https://doi.org/10.1145/75199.75232

Sadiq M. and Devi V. S. (2023). Prioritization and selection of the software requirements using rough-set theory, IETE Journal of Research, Vol. 69, Issue 8, pp. 5169–5186. https://doi.org/10.1080/03772063.2021.1973593

Sadiq, M. and Jain, S.K. (2014). Applying fuzzy preference relation for requirements prioritization in goal-oriented requirements elicitation process. International Journal of Systems Assurance Engineering and Management, Vol. 5, pp. 711–723. https://doi.org/10.1007/s13198-014-0236-3

Sadiq, M. and Jain, S.K. (2015). A fuzzy based approach for the selection of goals in goal oriented requirements elicitation process, International Journal of Systems Assurance Engineering and Management, Vol. 6, pp. 157–164. https://doi.org/10.1007/s13198-014-0260-3

Saxena, V.R., Kumar, B., Goel, N., and Sharma, Y. K. (2025). Software requirements prioritization using fuzzy based TOPSIS methods. In: Singh, R., Gehlot, A. (eds) Business Data Analytics. ICBDA 2023. Communications in Computer and Information Science, vol 2358. Springer, Cham. https://doi.org/10.1007/978-3-031-80778-7_9

Shao F., Peng R., Lai H., and Wang B. (2017). Drank: A semi-automated requirements prioritization method based on preferences and dependencies, Journal of Systems and Software, Volume 126, pp. 141-156, ISSN 0164-1212, https://doi.org/10.1016/j.jss.2016.09.043

Tasneem N., Zulzalil H. B., and Hassan S. (2025). Enhancing agile software development: a systematic literature review of requirement prioritization and reprioritization techniques, IEEE Access, Vol. 13 pp. 32993-33034. https://doi.org/10.1109/ACCESS.2025.3539357

Yu E. (1997). Towards modelling and reasoning support for early-phase requirements engineering, Proceedings of the 3rd IEEE International Symposium on Requirements Engineering (RE'97), pp. 226–235. https://doi.org/10.1109/ISRE.1997.566873

Zhai L., Khoo T., and Zhong S. (2008). A rough set enhanced fuzzy approach to quality function deployment, International Journal of Advanced Manufacturing Technology, 37, 613–624. https://doi.org/10.1007/s00170-007-0989-9

AUTHOR'S BIOGRAPHY

Mohammad Asim is a Ph.D. scholar in the School of Engineering and Technology, Sushant University, Gurugram, India. He received the M. Tech. degree in Computer Engineering from Maharshi Dayanand University, Rohtak, India, in 2011, and the B. Tech. degree in Civil Engineering from the Faculty of Engineering and Technology, Jamia Millia Islamia, New Delhi, India, in 1996. He also completed a Post-Graduate Diploma in Computer Applications from Jamia Millia Islamia, New Delhi, India, in 1999. He has published his research in Systems and Soft Computing (Elsevier, 2025), the International Journal of Modeling and Optimization, the IEEE International Conference on Data Storage and Data Engineering (2010), the IEEE International Conference on Computer Research and Development (2010), and the IEEE International Conference on Computer and Automation Engineering (2010). As part of his Ph.D. work, he has presented his research at the 7th World Conference on Artificial Intelligence: Advances and Applications (WCAIAA 2026), organized by the National Forensic Sciences University, Goa Campus, India, and at the 4th International Conference on Power Engineering and Intelligent Systems (PEIS 2026), organized by the National Institute of Technology Uttarakhand, India. He holds a registered design titled "IoT Controlled Self-

Adjusting Computer Webcam" (Design No. 380130-001, India, February 2023), and a second design application titled "Tactile Sensor-Based GPS Tracking and Unidirectional Vibrator Mobility Device for Visually Impaired Users" (Application No. 504311-001, filed June 2026) is currently under examination at the Indian Designs Office. His research interests include information security, cyber security, software engineering, and applied mathematics.

Rinky Ahuja is working as an Associate Professor in the School of Engineering and Technology, Sushant University, Gurugram, India. She received Ph.D. degree in Computer Science in 2022 from the Department of Computer Science, Sushant University, Gurugram, India. She completed both M. Tech. degree and B. Tech. degree in Computer Science in 2011 and 2006, respectively, from Kurukshetra University, Kurukshetra, India. She has published her research work in the International Journal of Recent Technology and Engineering, Journal of Data Acquisition and Processing, International Journal of Innovative Technology and Exploring Engineering. Her area of interest includes Machine Learning, Artificial Intelligence, Neural Networks, and Network Security.

Farhana Mariyam is working as an Assistant Professor and Head, School of Business, Sushant University, Gurugram, India. She received her Ph.D. in Computer Technology in 2026 from the Department of Electrical Engineering, Jamia Millia Islamia, New Delhi, India. She was awarded an M. Tech. degree in Computer Science from Jamia Hamdard, New Delhi, India, in 2016, and completed her B. Tech. degree in Computer Science and Engineering from Maharshi Dayanand University, Rohtak, India, in 2012. Her research interests encompass Software Engineering, Requirements Engineering, Soft Computing, and Intelligent Decision-Making Systems. She has published her research in prestigious journals and conference proceedings, including the Journal of Intelligent & Fuzzy Systems (SCI Expanded, 2022), the International Conference on Computing in Engineering & Technology (Springer, 2022), and the International Conference on Applied Soft Computing and Communication Networks (Springer, 2021). Her research focuses on applying computational intelligence techniques to address complex challenges in software engineering and decision-support systems.

Mohammad Salman received the B.Sc. Eng. degree in Electrical Engineering from Bilkent University, Ankara, Turkey, in 2018, and the M.Sc. Eng. degree in Computer Engineering from Bilkent University, Ankara, Turkey, in 2022. He is currently pursuing the Ph.D. degree at the University of Zurich, Zurich, Switzerland, where he has been a doctoral researcher since August 2022. His research focuses on trustworthy and fair artificial intelligence for healthcare applications, particularly medical image analysis. His doctoral work investigates fairness in medical image classification through the development of bias-aware loss functions for long-tailed and imbalanced datasets. He develops bias-mitigation frameworks for deep learning systems, with emphasis on fairness, transparency, explainability, and robust classification in medical decision-support systems. He is the author of the paper 'Robust One-Class Classification Using Deep Kernel Spectral Regression,' published in Neurocomputing (2024). His research interests include fair AI, medical image analysis, explainable AI, trustworthy machine learning, and deep learning for healthcare.