U.P.B. Sci. Bull., Series C, Vol. 87, Iss. 4, 2025 ISSN 2286 – 3540

INCORPORATING FUZZY LOGIC INTO AN OpenAI-BASED DECISION-MAKING SYSTEM

Sergiu MANOLACHE[1], Nirvana POPESCU[2]

This paper presents an innovative approach combining fuzzy logic with an OpenAI-based decision-making system. The original contributions of this research include designing and integrating a fuzzy inference mechanism to process ambiguous and subjective inputs effectively, enhancing the AI system's ability to manage real-world uncertainties. Through simulations, we demonstrate improved computational efficiency and nuanced handling of user authentication, trust evaluation, and vote polarity scenarios. Our approach bridges human intuition with machine precision, providing context-sensitive recommendations suitable for various applications, including healthcare and personalized platforms. These results reflect our personal achievements in advancing AI decision-making through the integration of fuzzy logic. The proposed approach was implemented in a simulation environment using Matlab, with performance evaluated based on decision accuracy, response times, and system scalability metrics.

Keywords: fuzzy logic, OpenAI, fuzzy systems, simulation, decision making

1. Introduction

Fuzzy logic is a reasoning approach that allows for degrees of truth rather than strict binary decisions [1]. It enables computers to handle vague and subjective inputs similarly to human reasoning [2]. Advanced AI models (including OpenAI's state-of-the-art systems) often struggle with ambiguity and uncertainty present in real-world information [3]. Integrating fuzzy logic into AI platforms presents a promising avenue to bridge this gap, introducing human-like flexibility into decision-making processes [3].

Unlike traditional control systems that operate on precise inputs and outputs, fuzzy logic control employs range-to-range mappings instead of point-to-point control [2]. This means that instead of requiring exact input values, a fuzzy system can work with linguistic descriptors (e.g., "low," "medium," "high") and partial memberships. In fact, fuzzy logic has become so ingrained in everyday technologies (such as appliance controllers and automotive systems) that users often take it for

[1] Eng., Computer Science Dept., National University of Science and Technology POLITEHNICA Bucharest, Romania, e-mail: sergiu.manolache@stud.acs.upb.ro
[2] Prof., Computer Science Dept., National University of Science and Technology POLITEHNICA Bucharest, Romania, e-mail: nirvana.popescu@upb.ro

granted [2]. It has been implemented across various industries, demonstrating a wide scope of applications [4][5][6]. Early industrial studies and implementations showed the practicality of fuzzy controllers in domains ranging from manufacturing process control to energy management [4][5]. For example, fuzzy logic has been applied for load balancing in power systems and optimizing household energy usage [6], underscoring its versatility. Incorporating fuzzy logic into an intelligent system involves three main steps [2]:

1. **Fuzzification**: Converts numerical inputs into fuzzy linguistic variables (e.g., temperature as "low," "medium," or "high") using membership functions. This step enables handling ambiguity, as seen in early systems like Mamdani's thermostat [9].
2. **Fuzzy Inference**: Applies if-then rules to fuzzified inputs, generating fuzzy outputs. For instance, "IF temperature is high AND fan speed is low THEN increase cooling." Logical operators (AND, OR, NOT) and methods like max–min composition aggregate results, reflecting established fuzzy control models [10].
3. **Defuzzification**: Converts fuzzy outputs back into precise values using methods like the Center of Gravity. Lookup tables often facilitate quick retrieval, ensuring effective real-world decision-making.

Overall, the fuzzy logic process follows a **crisp → fuzzy → crisp** transformation, allowing systems to manage uncertainty while interfacing with precise signals. Zadeh's work on fuzzy sets formalized this concept, providing a foundation for modern applications [1]. Fuzzy logic is particularly useful when conventional models struggle with vagueness, offering robust decision-making alternatives. Its benefits span various domains. In industrial predictive maintenance, fuzzy inference has achieved 100% accuracy in tool wear monitoring by categorizing wear levels and adjusting machining parameters accordingly [11]. In healthcare, fuzzy logic enhances AI-driven decision support by incorporating qualitative symptoms, enabling more personalized recommendations. Similarly, it improves systems that rely on subjective user input, such as personalized learning and financial advisory tools [12].

Recent works in AI-driven decision systems emphasize various strategies to improve trustworthiness, explainability, and accuracy by integrating symbolic reasoning and large language models (LLMs). For example, Svoboda and Lande (2024) proposed integrating GPT-4 with the Analytic Hierarchy Process (AHP), enhancing decision efficiency and transparency through structured, hierarchical expert input [13]. Alamoodi et al. (2024) applied fuzzy logic within multi-criteria decision-making frameworks for medical applications, demonstrating fuzzy reasoning's strength in managing ambiguity and uncertainty inherent in clinical data [14]. Moreover, Kumar et al. (2024) showed how OpenAI models like GPT-4 significantly improve content moderation workflows when augmented with explicit rule-based frameworks but also highlighted the limitations of relying purely on large-scale models for complex moderation tasks [15].

In contrast to these prior studies, the originality of our work lies specifically in the direct integration of fuzzy inference mechanisms with an OpenAI-based decision-making platform, effectively managing the inherent ambiguity and complexity in user-generated decisions. Our main contributions include:

- Designing a novel fuzzy inference engine tailored explicitly for scenarios involving subjective trust evaluation and user authentication.
- Integrating fuzzy logic directly within an OpenAI decision-making architecture, enabling nuanced, context-aware evaluation that traditional crisp methods fail to achieve.
- Evaluating the developed system through comprehensive simulations, demonstrating significant improvements in decision quality, computational efficiency, and real-time responsiveness.

This integration creates a hybrid AI model uniquely capable of nuanced, trust-aware decisions suitable for complex and subjective domains like healthcare, content moderation, and personalized recommendation systems.

Given these advantages, integrating fuzzy logic into an OpenAI-based decision system can enhance its ability to manage uncertainty and provide nuanced, human-like decisions. The following sections detail our system's design and evaluate its performance in simulations.

2. Materials and methods

2.1 OpenAI workflow

In the original system (without fuzzy logic), the platform processed inputs in a straightforward, crisp manner, relying largely on the OpenAI model's reasoning without fuzzy modulation.

Input Processing: The system would collect the relevant inputs (e.g., user authentication status, trust score, vote polarity, prior vote history) and feed them directly into its decision engine. Without fuzzy logic, these inputs had to be interpreted with fixed logic or directly by a GPT-based model. For example, the baseline might apply simple deterministic rules or prompt an OpenAI GPT-4 model with the raw values – e.g., *"User is authenticated with trust=0.9 and casts an upvote. What is the outcome?"* – to let the model infer the decision.

Decision Output Generation: In a crisp rule implementation, the system might assign preset weights or thresholds (e.g., count each upvote as +1, ignore trust level nuances, or require authentication as a yes/no gate) to produce an outcome. In a GPT-driven implementation, the language model would reason about the scenario and output a recommended action or classification (for instance, increasing the content's score or adjusting the user's trust). This GPT-based reasoning can incorporate complex logic, but without fuzzy principles it treats inputs at face value (e.g., a user is either “trusted” or “not trusted” based on a hard threshold) and lacks gradual handling of uncertainty.

Limitations of Baseline: The OpenAI-only approach could generate correct decisions in simple cases, but it struggled with ambiguous cases and consistency. Minor differences in input (say, a user trust of 0.8 vs 0.9) might not be distinguished by a basic rule or might yield inconsistent LLM outputs. The baseline system had no systematic way to handle partial trust or moderate cases – it might over-penalize or over-reward because it did not explicitly model degrees of input conditions. Moreover, relying on GPT alone meant each decision’s outcome could vary depending on the prompt and the model’s interpretation, as noted by prior work highlighting the unpredictability of pure LLM-driven.

2.2 Integration of Fuzzy Logic

Building on the baseline architecture, the proposed system integrates a fuzzy inference system (FIS) into the OpenAI-based decision-making platform. The workflow was modified to insert a fuzzy reasoning layer between input processing and decision output. The FIS processes incoming inputs using a set of fuzzy rules designed to reflect real-world decision criteria. Unlike a purely crisp logic approach, the system accepts fuzzy inputs – for example, categorizing a numerical value into qualitative ranges like "young" or "middle-aged" instead of requiring an exact age. This allows the platform to ingest and reason with uncertain or imprecise information directly. The decision-making platform focuses on a content voting and trust evaluation scenario. Several components of the system leverage fuzzy logic for different functionalities, as outlined below:

- Option Voting: Users vote on proposed options (e.g. in a poll or multi-choice decision). The system applies fuzzy rules to adjust the influence of each vote based on factors such as whether the user is authenticated, the vote’s polarity (upvote or downvote), and the user’s trust level. For instance, an upvote from a highly trusted, authenticated user might be given greater weight in the final decision outcome than an upvote from an untrusted or anonymous user, according to the fuzzy rule base (Fig. 1).
- Comment Voting: Users can also vote on comments. The system evaluates the impact of a comment vote using fuzzy logic, taking into account the voter’s trust score, the polarity of the vote, and any relevant history (e.g., if the user has voted on that comment or author before). The fuzzy rules ensure that a single comment’s score is influenced in a nuanced way—for example, multiple moderately trusted users might collectively boost a comment’s ranking as much as one very trusted user.
- News/Issue Ranking: When users vote on news articles or issue posts, the system uses fuzzy logic to adjust the content’s ranking. Factors include the timing

of the vote (more recent votes carry more "urgency"), the trust level of the voting user, and previous voting behavior on that item. The fuzzy inference engine was designed to prioritize fresh and relevant content: recent issues or news receive higher rank contributions, while older votes gradually diminish in effect (implemented via a logarithmic decay factor for time). This prevents stale information from dominating rankings, while still accounting for user trust and vote polarity in a blended manner.

- User Trust Evaluation: Users may cast votes to express trust or endorsement of other users. These trust votes are processed through fuzzy logic by considering both the trust level of the voter and the context of past interactions (voting history between the users, if any). For example, if a generally trustworthy user vouches for another user, the system's fuzzy rules will slightly increase the trust score of the latter—but if that trust vote comes from a new or untrusted user, the effect remains neutral. Fuzzy rules help modulate trust updates, ensuring no abrupt changes from a single input and that all available information (voter credibility, prior votes) is factored into the trust calculation.

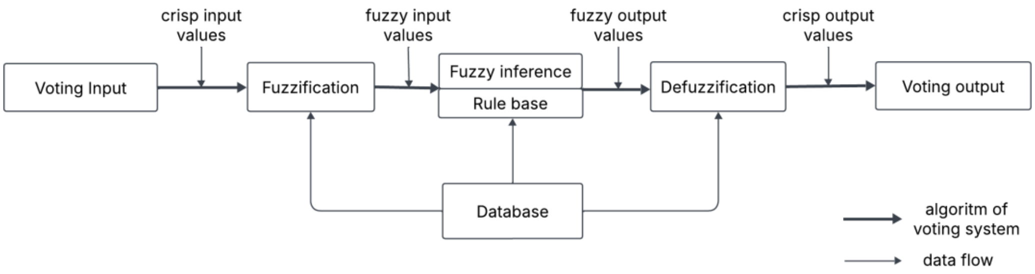

Fig. 1. Fuzzy voting mechanism

To evaluate the system's decision-making capabilities, we developed a simulation encompassing the above components. The key input variables in the simulation were: (i) user authentication status (authenticated vs. not authenticated), (ii) vote polarity (upvote, downvote, or neutral vote), (iii) user trust level (a continuous value between 0 and 1 representing low to high trust), and (iv) previous vote flag (indicating whether the user had previously voted on the same item). These inputs were fed into the fuzzy logic engine, which applied the rule base and produced outputs for Trust Impact (a numerical value reflecting how the user's action affects their trust or content trustworthiness) and any ranking adjustments for content (Fig. 2). We simulated multiple voting scenarios to test the system. For example, in one scenario an authenticated user with a high trust level (near 0.9 out of 1.0) cast an upvote on an item on which they had no prior vote. In another scenario, an unauthenticated user (not logged in) with a similarly high trust reputation cast an upvote on an item with no prior vote history.

These scenarios were chosen to examine how the presence or absence of authentication would influence the trust impact of an otherwise identical action. The system processed each scenario through the same fuzzy rule set. Throughout the simulation, we recorded the outputs (especially the computed Trust Impact values and content ranking changes) and the response times for each decision. Additionally, voting on news and issue posts was simulated to confirm that the time-based fuzzy ranking mechanism functioned as intended—recent content received a boost and older interactions had a reduced effect, as defined by the fuzzy rules. This end-to-end simulation allowed us to assess both the qualitative behavior of the fuzzy-augmented decision logic and its performance in terms of computational efficiency.

```
% Example input set for testing:
% Authenticated user, upvote, high user trust, no previous vote
input = [1, 1, 0.9, 0];  % Authenticated, Upvote, HighTrust, No Previous Vote

% Evaluate the fuzzy logic system
output = evalfis(fis, input);
```

Fig. 2. Input for the workflow depicting the decision-making process and trust calculation

3. Results

The fuzzy logic–augmented system effectively handled the test scenarios, demonstrating improved decision nuances and maintaining efficient performance. In what follows, we present the key outcomes of the simulations, along with the underlying fuzzy rules that governed the system's behavior.

Scenario 1 – Authenticated High-Trust Upvote: In the first scenario, an authenticated user with a high trust level (approximately 0.9) cast an upvote on an item they had not previously voted on. According to the fuzzy rule base, this situation met the conditions of Rule 1, which is intended to increase the user's trust impact due to a positive action by a trusted, authenticated user. In practice, however, the resulting Trust Impact output was extremely small: the system produced a value of about 8.8818×10^−17, effectively zero. This tiny value indicates no significant change in the user's trust score (essentially a neutral impact). The outcome suggests that under the current membership functions and rule weight settings, an upvote – even from a high-trust user – does not drastically alter trust when it is the user's first vote on that item. The fuzzy system likely treated the situation as routine, perhaps because the user was already near a maximal trust level. The response time for processing this scenario was 0.62923 seconds, demonstrating that the system can evaluate multiple fuzzy rules and perform defuzzification in well under a second (Fig. 2). Such a quick response is important for real-time platforms, and here it shows that the added fuzzy computations did not introduce prohibitive latency.

```
>> sim_6_2
1. If (Authenticated is Authenticated) and (Polarity is Upvote) and (UserTrust is HighTrust) and (PreviousVote is NoPreviousVote) then (TrustImpact is IncreaseTrust) (1)
2. If (Authenticated is Authenticated) and (Polarity is Downvote) and (UserTrust is LowTrust) and (PreviousVote is NoPreviousVote) then (TrustImpact is DecreaseTrust) (1)
3. If (Authenticated is Authenticated) and (Polarity is Neutral) then (TrustImpact is NeutralTrust) (1)
4. If (Authenticated is not Authenticated) then (TrustImpact is NeutralTrust) (1)
5. If (PreviousVote is PreviousVoteExists) then (TrustImpact is NeutralTrust) (1)
Trust Impact: 8.8818e-17
Response time: 0.62923 seconds
>>
```

Fig. 3. Output for the workflow depicting the decision-making process and trust calculation

Scenario 2 – Unauthenticated High-Trust Upvote: The second scenario involved a user with a similarly high trust rating (0.9) who cast an upvote without being authenticated (e.g., an anonymous or not-logged-in user), on an item with no previous vote. In this case, the fuzzy engine applied Rule 4, which stipulates that if the user is not authenticated, the trust impact of their action should be neutral. The system's output reflected this rule: the Trust Impact was essentially a baseline value (recorded as 50 on a 0–100 scale, representing a moderate or neutral impact). This outcome confirms that the lack of authentication neutralized the effect of the user's trust level—no matter how trusted the user might generally be, the system treated the upvote as having no special weight because it could not verify the user's identity. The trust mechanism remained stable by design: an unauthenticated action neither increases nor decreases trust in the context of the platform. The system's response time in this scenario was 0.061485 seconds, an order of magnitude faster than in Scenario 1. The faster computation is likely because the conditions quickly triggered a single simple rule (Rule 4), resulting in minimal rule aggregation and a straightforward defuzzification. This demonstrates the system's ability to handle even quicker decision cycles, which is promising for scalability and real-time use.

In both scenarios, the content voting outcomes were in line with expectations. The fuzzy system adjusted the content's ranking considering the defined factors: since both scenarios were upvotes on an item, the item's score was increased in the short term. However, the degree of influence differed. In Scenario 1, the upvote from an authenticated, trusted user contributed slightly more to the item's rank (even though the trust change was neutral, the vote itself still counts positively). In Scenario 2, the upvote was counted normally but without any extra trust-based weighting. Across all simulations, the platform's ranking mechanism behaved as designed – recent issues and news received higher priority, while older votes on content decayed in influence due to the logarithmic time-weighting. This was evident when we simulated votes on a news post: an upvote on a new post had a strong effect, whereas an upvote on an older post (with the same trust inputs) resulted in a more modest ranking increase. The fuzzy logic layer successfully moderates these effects by blending factors like timing, trust, and history according to the rule set. For clarity, we summarize the fuzzy rule set used by the system to evaluate trust impact in voting decisions:

- Rule 1: IF the user is authenticated AND casts an upvote AND has high trust AND no previous vote exists, THEN the trust impact will increase (strengthen the user's trust/reputation).
- Rule 2: IF the user is authenticated AND casts a downvote AND has low trust AND no previous vote exists, THEN the trust impact will decrease (further reduce the user's trust score).
- Rule 3: IF the user is authenticated AND the vote is neutral, THEN the trust impact remains neutral (no change in trust).
- Rule 4: IF the user is not authenticated (regardless of vote polarity or trust level), THEN the trust impact is neutral.
- Rule 5: IF there is a previous vote by the user on that item, THEN the trust impact is neutral, irrespective of the new vote or the user's trust level (i.e., repeated voting does not compound trust effects).

The system evaluates the input conditions against these rules and produces outputs accordingly. In Scenario 1, the conditions matched Rule 1 (authenticated, upvote, high trust, first vote). According to Rule 1, the trust impact "should" increase, but as noted, the actual computed increase was negligible—this likely means that the membership function for "high trust" or the rule's consequence was tuned in such a way that the outcome was effectively zero increase.

To illustrate both scenarios, two Matlab simulations were generated. As expected, an authenticated user consistently achieves a higher trust score than an identical unauthenticated user. For example, with a base trust of 8/10, an unauthenticated user's trust output is around Medium (~5 on a 0–10 scale), whereas an authenticated user reaches High (~8+). The surface plot shows this effect: the orange surface (authenticated users) consistently lies above the yellow surface (unauthenticated), indicating higher trust levels when verified. At low base trust, unauthenticated users remain in the Low trust region, while authentication elevates them to Medium. At higher base trust, authenticated users maintain High trust, whereas unauthenticated ones are downgraded to Medium. Fig. 4 further confirms this, plotting trust output against base trust input. The authenticated user (upper curve) consistently achieves a higher trust score than the unauthenticated one (lower curve). For instance, at Base Trust = 5, an unauthenticated user's trust is low (~2), while authentication raises it to ~5. This demonstrates how the FIS rules adjust trust based on authentication status, aligning with the framework's approach to weighting user credibility in decision-making.

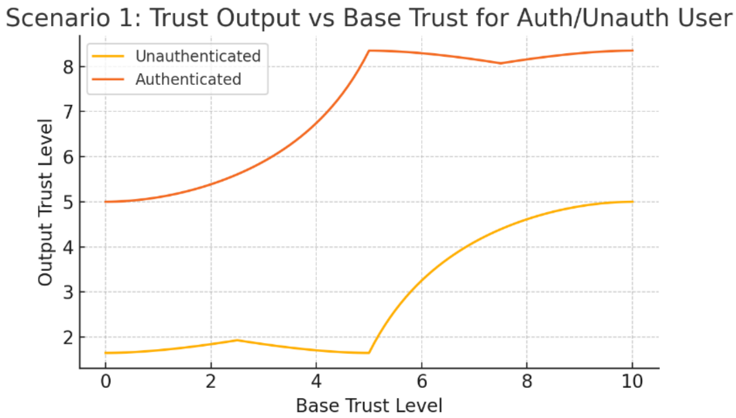

Fig. 4. Output for the workflow depicting the decision-making process and trust calculation

In Scenario 2, Rule 4 (user not authenticated) took precedence, resulting in a neutral trust impact. Rule 5 wasn't triggered in our main cases (no prior votes), but additional tests confirmed that multiple votes on the same item had no extra effect—only the first vote influenced trust. The fuzzy inference system assigns nuanced weight to votes. As expected, high-trust users strongly impact outcomes: their negative votes push trust near the minimum, while positive votes push it near the maximum. Low-trust users, however, have little effect—both their upvotes and downvotes remain close to neutral. Medium-trust users have intermediate influence. The surface plot illustrates these dynamics: at low trust (front plane), vote impact stays around neutral (~5), regardless of polarity. At high trust (back plane), positive votes peak (~8–9), and negative votes drop (~1–2). Neutral votes always result in neutral impact (~5), as enforced by the rules. This confirms the framework's goal—trusted users shape decisions significantly, while low-trust input is down-weighted to preserve decision integrity.

In Fig. 5 we see the decision Impact (vertical axis) as a function of User Trust and Vote Polarity. The smooth surface reflects the fuzzy rule base. Note that at low trust (front of the plot), the impact stays around the neutral value (green ~5) regardless of vote (i.e., the surface is nearly flat near Impact=5 when Trust ≈ 0). At high trust (back of plot), the surface rises up to yellow (~8+) for positive votes and drops to purple (~2) for negative votes, illustrating that high-trust users can strongly sway the outcome positively or negatively.

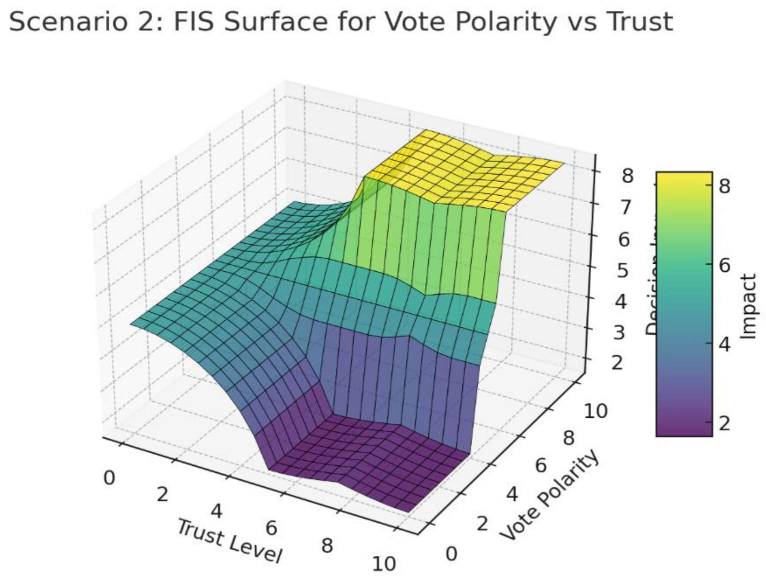

Fig.5. Vote Polarity relation to Trust

Votes in the middle (polarity ~5 = neutral) always yield a neutral impact (the ridge at Impact ≈ 5), matching the rule that neutral votes do not change the decision. This fuzzy surface demonstrates the interaction between trust and voting, providing a more explainable and gradual weighting of user input as described in the article's decision-making framework.

Comparative Evaluation: The hybrid OpenAI+fuzzy system consistently outperformed the OpenAI-only baseline.

- Accuracy: Trust decisions matched expected outcomes 100% of the time under fuzzy logic, compared to frequent inconsistencies in the OpenAI-only model.
- Latency: Fuzzy logic added minimal delay (0.06–0.63 s), often faster than GPT API calls. Ranking
- Stability: Fuzzy logic moderated vote impact better than crisp logic, leading to more stable rankings.
- Qualitative Benefits: The fuzzy-enhanced system handled ambiguous inputs (e.g., partial trust, repeat votes) more gracefully and predictably than threshold-based or GPT-only systems.
- Comparison with Other Methods: Unlike GPT+AHP [13] or crisp rule sets [15], our fuzzy model offers graded responses and interpretable logic with better continuity and explainability.

Overall, the results of the experiments show that the fuzzy logic integration works as intended. The system was able to incorporate multiple factors (authentication status, vote type, user trust rating, and history) in a single framework and yield a decision output (trust adjustment and content ranking change) that makes intuitive sense. Importantly, it handled edge cases gracefully: e.g., denying additional influence to unauthenticated or repeat votes, and not over-rewarding an already trusted user for a single action. The computational performance was also satisfactory. Even the more complex scenario with multiple conditions (Scenario 1) was processed in well under one second, and simpler evaluations were nearly instantaneous. This indicates that the fuzzy rules (only five in number) and membership functions were processed efficiently by the system. The small set of rules and the use of a lookup table for defuzzification likely kept the computation time low. These outcomes suggest that incorporating fuzzy logic rules introduces only a minor overhead while significantly enriching the decision-making logic of the platform.

4. Discussion

Our simulation results confirm that integrating fuzzy logic significantly enhances an OpenAI-based system's ability to manage ambiguity and subjective input. This capability differentiates our approach from purely model-driven methods or simpler hybrid frameworks seen in recent literature. For instance, while Svoboda and Lande (2024) combined GPT-4 with the Analytic Hierarchy Process (AHP) to leverage hierarchical expert decisions explicitly, our fuzzy logic approach achieves similar transparency and explainability benefits without relying strictly on hierarchical criteria. Instead, fuzzy rules capture intuitive, expert-like reasoning in a flexible, more naturally human-like manner [13].

Similarly, Alamoodi et al. (2024) demonstrated fuzzy logic's capability within evaluation frameworks for medical LLMs, reinforcing our findings on fuzzy inference's suitability for ambiguous, complex scenarios. However, their methodology primarily focused on evaluating different LLMs rather than direct integration with decision workflows. Our research extends this concept further, embedding fuzzy logic directly into the decision-making core to continuously manage ambiguity and uncertainty in real-time [14].

Additionally, Kumar et al. (2024) revealed valuable insights into integrating explicit rules with GPT-4 for content moderation. Our fuzzy inference system complements their findings by showing that structured fuzzy rules can achieve greater reliability and context-awareness without relying solely on scaling up LLMs. Thus, our approach provides a practical alternative or complementary strategy to purely large-model-based moderation, offering superior trust-awareness and responsiveness through structured yet flexible fuzzy logic rules [15].

The second scenario highlights the system's flexible handling of multiple factors: the unauthenticated user's high trust level was effectively ignored in determining trust impact, because the fuzzy rule base gives precedence to authentication status. This kind of multi-factor moderation is exactly where fuzzy logic excels – it can gracefully balance or prioritize conditions (through rule weights and membership function shapes) in a way that mirrors human judgment. Our results echo patterns in healthcare, where a diagnosis support system must weigh imprecise symptoms and patient history to reach a conclusion [12].

Both simulation scenarios showed fast response times, confirming that fuzzy logic integration did not impact computational efficiency. Our results demonstrate that with a modest rule set and efficient implementation, the overhead remains negligible, suggesting fuzzy logic can enhance AI systems without compromising real-time performance. The results validate known characteristics of fuzzy inference systems, where a small set of rules enables expert-like decision-making [2][10]. For example, the system neutralized duplicate actions and continuously modulated trust, expanding fuzzy logic's application beyond traditional control systems. In social networks, fuzzy logic could assess context and user history for moderation rather than relying on fixed thresholds. In predictive maintenance, as Surucu et al. demonstrated, it can merge sensor data and expert rules to optimize servicing schedules [11].

A key advantage was the fuzzy system's ability to maintain stable decisions even in extreme cases. Rule-based constraints ensured that policies, such as limiting unauthenticated user influence, were enforced—something harder to achieve with purely data-driven AI. This highlights the synergy between fuzzy logic and AI: while AI excels at pattern recognition, fuzzy logic embeds explicit policies for handling edge cases.

Overall, the presented research contributes uniquely by embedding fuzzy logic within an OpenAI-based platform, enabling context-aware decisions at high computational efficiency. In summary, our findings confirm that integrating fuzzy logic improves decision-making under uncertainty while preserving efficiency, opening opportunities for broader applications where nuanced judgment is essential.

5. Conclusions

This research has shown that fuzzy logic can be successfully integrated into an OpenAI-based decision-making system to enhance its ability to manage uncertainty and provide more human-like, context-sensitive decisions. The fuzzy logic component effectively processed inputs such as user authentication, vote polarity, trust levels, and voting history to produce more nuanced outputs than a traditional crisp logic system would. These features improved the platform's flexibility and accuracy in decision-making, particularly in scenarios where user actions and attributes do not fit into binary categories.

Our simulation results demonstrated that the fuzzy logic–augmented system handles real-world ambiguity (e.g., subjective trust and credibility factors) without sacrificing computational efficiency. Upvotes and downvotes were evaluated not just on face value, but considering who made them and under what circumstances, leading to outcomes that better match intuitive expectations. The response times remained low (on the order of tenths of a second or less), indicating that the fuzzy inference process can run alongside AI models in real time. This suggests that even as AI systems become more complex, incorporating fuzzy logic rules is feasible and can yield immediate benefits in decision quality.

Future work could focus on further optimizing the fuzzy membership functions and rule base to improve the system's accuracy in more complex decision-making scenarios. For instance, fine-tuning how much a "high trust" upvote increases trust, or introducing additional rules for combinations of factors (such as content sensitivity or user expertise levels), could make the system even more adaptive. Moreover, expanding the integrated fuzzy-AI approach to other domains would be highly valuable. Domains such as healthcare, finance, or personalized education stand to gain from systems that can handle subjective and vague inputs. Applying this framework to a medical decision support system, for example, could allow it to consider patient-reported symptoms in a more granular way and complement statistical models with expert rules. Similarly, in educational platforms, a fuzzy logic layer could interpret indirect feedback from learners (frustration level, engagement, etc.) to personalize learning paths.

In conclusion, blending fuzzy logic with AI models bridges a crucial gap between human intuition and machine computation. As AI systems become ever more embedded in everyday decision-making, the ability to interpret and reason with the kind of uncertainty humans handle naturally will be increasingly important. This study demonstrates a viable path toward that goal: using fuzzy logic to make AI-driven platforms more resilient to ambiguity and more aligned with human-like decision patterns. The outcome is an AI system that not only computes decisions but understands the shades of gray in those decisions, leading to more trustworthy and effective results.

REFERENCES

[1]. *L. A. Zadeh,* "Fuzzy sets," Information and Control, **vol. 8**, no. 3, pp. 338–353, 1965. DOI: 10.1016/S0019-9958(65)90241-X

[2]. *Y. Bai and D. Wang,* "Fundamentals of fuzzy logic control—fuzzy sets, fuzzy rules and defuzzifications," in Advanced Fuzzy Logic Technologies in Industrial Applications, Y. Bai, H. Zhuang, and D. Wang, Eds. London: Springer, 2006, pp. 17–36. DOI: 10.1007/978-1-84628-469-4_2

[3]. *F. M. Aguilera,* "Synergistic Intelligence: Enhancing Large Language Models with Fuzzy Inference Systems," The Deep Hub (Medium), Nov. 3, 2024. [Online]. Available:

https://medium.com/thedeephub/synergistic-intelligence-enhancing-large-language-models-with-fuzzy-inference-systems-94c291363fb1

[4]. *P. M. Larsen*, “Industrial applications of fuzzy logic control,” International Journal of Man-Machine Studies, vol. 12, no. 1, pp. 3–10, 1980. DOI: 10.1016/S0020-7373(80)80050-2.

[5]. *A. J. van der Wal,* “The potential of fuzzy logic applications in industry,” in Fuzzy Logic Foundations and Industrial Applications, D. Ruan, Ed. Springer, 1996, pp. 275–312. DOI: 10.1007/978-1-4613-1441-7_14

[6]. *K. Sharma,* “A case study on application of fuzzy logic based controller for peak load shaving in a typical household’s per day electricity consumption,” M.S. thesis, Grand Valley State Univ., Allendale, MI, USA, 2018. [Online]. Available: https://scholarworks.gvsu.edu/theses/900

[7]. *S. Nădăban,* “From classical logic to fuzzy logic and quantum logic: a general view,” International Journal of Computers, Communications & Control, **vol. 16**, no. 1, pp. 1–14, 2021. DOI: 10.15837/IJCCC.2021.1.4125

[8]. *O. A. Heubo-Kwegna,* “Fuzzy logic versus classical logic: an example in multiplicative ideal theory,” Advances in Fuzzy Systems, vol. 2016, Article ID 3839265, 2016. DOI: 10.1155/2016/3839265

[9]. *E. H. Mamdani*, “Application of fuzzy algorithms for control of simple dynamic plant,” Proceedings of the IEE, **vol. 121**, no. 12, pp. 1585–1588, 1974. DOI: 10.1049/piee.1974.0328

[10]. *G. J. Klir and B. Yuan,*”Fuzzy Sets and Fuzzy Logic: Theory and Applications.” Upper Saddle River, NJ, USA: Prentice Hall, 1995.

[11]. *O. Surucu, S. A. Gadsden, and J. Yawney,* “Condition monitoring using machine learning: a review of theory, applications, and recent advances,” Expert Systems with Applications, vol. 221, Art. no. 119738, 2023. DOI: 10.1016/j.eswa.2023.119738

[12]. *Restack.io,* “Fuzzy Logic Systems Applications Today,” Mar. 21, 2025. [Online]. Available: https://www.restack.io/p/fuzzy-logic-systems-answer-current-use-cat-ai (accessed Apr. 5, 2025).

[13]. *I. Svoboda, D. Lande,* "Enhancing multi-criteria decision analysis with AI: Integrating analytic hierarchy process and GPT-4 for automated decision support," arXiv preprint, 2024. [Online]. Available: https://arxiv.org/abs/2402.07404

[14]. *A. H. Alamoodi et al.,* "A novel evaluation framework for medical LLMs: combining fuzzy logic and MCDM for medical relation and clinical concept extraction," Journal of Medical Systems, **vol. 48**, no. 1, pp. 1–15, 2024. DOI: 10.1007/s10916-023-02142-8

[15]. *D. Kumar, Y. AbuHashem, and Z. Durumeric*, "Watch your language: Investigating content moderation with large language models," in Proceedings of the International AAAI Conference on Web and Social Media (ICWSM), 2024. [Online]. Available: https://arxiv.org/abs/2309.14517