

Fuzzy-Weighted Sentiment Recognition for Educational Text-Based Interactions

Christos Troussas^a, Christos Papakostas^b, Akrivi Krouska^c, Phivos Mylonas^d
and Cleo Sgouropoulou^e

Department of Informatics and Computer Engineering, University of West Attica, Egaleo, Greece

Keywords: Fuzzy Sentiment Analysis, Educational Natural Language Processing, Affective Computing in Learning, Interpretability in AI, Student Feedback Mining.

Abstract: In web-based educational environments, students often express complex emotional states – such as confusion, frustration, or engagement – through reflective texts, forum posts, and peer interactions. Traditional sentiment analysis tools struggle to capture these subtle, mixed signals due to their reliance on rigid classification schemes and lack of domain sensitivity. To address this, we propose a fuzzy-weighted sentiment recognition framework designed specifically for educational text-based interactions. The system combines an augmented sentiment lexicon, rule-based modifier detection, and semantic similarity using pretrained Sentence-BERT embeddings to extract nuanced sentiment signals. These inputs are interpreted by a Mamdani-type fuzzy inference engine, producing a continuous sentiment score and a confidence weight that reflect both the strength and reliability of the learner’s affective state. The paper details the linguistic pipeline, fuzzy membership functions, inference rules, and aggregation strategies that enable interpretable and adaptive sentiment modeling. Evaluation on a corpus of 1125 annotated student texts from a university programming course shows that the proposed system outperforms both lexicon-based and deep learning baselines in accuracy, robustness, and interpretability, demonstrating its value for affect-aware educational applications.

1 INTRODUCTION

In the context of distance learning, student communication is increasingly taking place via textual media, including discussion boards, reflective questions, peer review, and open-ended tests. These text-based discussions provide much insight into students' thinking, engagement, and affect (Yuvaraj et al., 2025). Yet, despite their value for instruction, these texts are not necessarily examined, and indicators of frustration, satisfaction, confusion, or motivation may not be noticed (Johansen et al., 2025; Troussas et al., 2019).

Instructors are frequently unaware of emotional undercurrents that could signal disengagement, conceptual difficulty, or misunderstanding – especially in asynchronous or large-scale settings

where personal attention is limited.

Sentiment analysis has emerged as a promising tool for augmenting educational platforms with affect-sensitive capabilities (Grimalt-Álvaro & Usart, 2024; Kardaras et al., 2024; Tasoulas et al., 2024). By identifying emotional cues in student language, sentiment models can help build responsive, personalized systems that adjust instructional strategies based on learner affect (Benazer et al., 2024). However, the majority of existing sentiment analysis methods rely on categorical labels, such as positive, negative, or neutral, and apply either lexicon-based heuristics or supervised classifiers trained on general-purpose corpora (Alahmadi et al., 2025). These approaches suffer from several limitations in the educational domain: they struggle with the ambiguity and nuance of student discourse,

^a <https://orcid.org/0000-0002-9604-2015>

^b <https://orcid.org/0000-0002-5157-347X>

^c <https://orcid.org/0000-0002-8620-5255>

^d <https://orcid.org/0000-0002-6916-3129>

^e <https://orcid.org/0000-0001-8173-2622>

lack interpretability, and often fail to provide actionable or trustworthy outputs for teachers or adaptive systems (van der Veen & Bleich, 2025).

Educational texts are not simple declarations of opinion (Ahmed et al., 2022). A single message may blend curiosity with uncertainty (“I’m not sure I got the logic right”), or hesitation with emerging confidence (“This recursion thing is starting to make sense”). The emotional expressions are often subtle, hedged, and domain-specific, particularly in STEM education where phrases like “I failed the test case” or “finally compiled successfully” carry implicit affect. In such contexts, standard sentiment tools tend to misclassify, overgeneralize, or ignore important cues (Hafner et al., 2025). Furthermore, educators require more than just a sentiment label – they need to know how strong that sentiment is, how reliable the estimate is, and how to interpret it in light of instructional goals.

To address these challenges, this paper proposes a fuzzy-weighted sentiment recognition framework tailored specifically for educational text-based interactions. Unlike traditional models that make binary or ternary decisions, our system employs fuzzy logic to model sentiment as a continuous and interpretable construct. It assigns each student utterance a sentiment score on a real-valued scale $[-1, +1]$ along with a confidence weight $[0, 1]$ reflecting both the polarity and the degree of certainty. This approach allows the system to capture the vagueness and variability inherent in learner expression, while maintaining pedagogical interpretability and technical robustness.

The model we use is grounded in real data gathered from a university-level Java programming course offered through an online learning environment. The analysis includes students’ forum posts, their weekly reflections, and feedback gathered at the end of the course, using a hybrid analytical pipeline that leverages linguistic preprocessing, domain-specific sentiment lexicons, contextual embeddings, and a Mamdani-type fuzzy inference engine. By combining domain-specific knowledge with fuzzy reasoning methods, our goal is to bridge affective computing with real-world applications in education.

The contributions of this paper are threefold. First, we introduce a novel fuzzy-weighted sentiment analysis model designed for the educational domain, which combines symbolic interpretability with context-aware computation. Second, we develop a domain-specific sentiment lexicon enriched with intensity and confidence metadata, adapted to student language in technical learning contexts. Third, we

evaluate our approach on a curated dataset of annotated educational texts, comparing it against both classical and deep learning sentiment baselines. Our results show that the proposed system not only achieves competitive performance but also produces more nuanced, trustworthy outputs that can support adaptive learning and instructor awareness.

The remainder of the paper is structured as follows. Section 2 reviews prior work in sentiment analysis, particularly in educational contexts. Section 3 outlines the challenges and motivations for modeling sentiment in student-generated content. Section 4 presents the architecture and logic of the fuzzy-weighted sentiment framework. Section 5 describes experimental setup, and analyzes the results. Finally, Section 6 concludes with reflections and directions for future research.

2 RELATED WORK

Sentiment analysis assists us in grasping information on the web and the way individuals utilize language. It comes in handy in fields such as obtaining customer opinions, social media monitoring, and websites that provide suggestions. Sentiment analysis is also gaining popularity in the field of education for monitoring students’ moods, comprehending their emotions, and assisting with personal learning. Nevertheless, technology currently is often not effective when dealing with complicated matters, ambiguous meanings, and context in educational debates.

Initially, sentiment analysis techniques predominantly employed rules and word lists such as SentiWordNet, AFINN, and VADER (Hutto & Gilbert, 2014). These provide fixed scores to words or phrases. These are robust and require minimal training data, but they are not contextual, do not handle negation, and do not handle varying uses of language in different domains. In addition, methods that employ lexicon lists tend to miss nuanced or blended feelings. This is usually observed in school when students express frustration and improvement in a single sentence.

With the rise of machine learning, supervised classifiers such as Naïve Bayes, Support Vector Machines (SVM), and Random Forests were introduced for sentiment analysis, offering improved generalization and adaptability to specific domains (Pang & Lee, 2008). More recently, deep learning models – particularly Recurrent Neural Networks (RNNs), Convolutional Neural Networks (CNNs), and Transformer-based architectures such as BERT –

have achieved state-of-the-art performance in sentiment classification tasks across multiple languages and datasets (Devlin et al., 2019; Liu et al., 2019). These models learn contextual embeddings and capture long-range dependencies, allowing them to outperform traditional methods in complex textual environments. However, they typically require large annotated corpora, lack transparency, and offer limited pedagogical interpretability—factors that pose challenges for adoption in educational applications.

In education, sentiment analysis has been applied to examine students' comments (Altrabsheh et al., 2014), discussion boards (Yang et al., 2013), and intelligent tutoring systems (Litman & Forbes-Riley, 2006) to monitor happiness, detect frustration, or forecast whether students will drop out. For instance, (Wen et al., 2014) employed sentiment trends to examine how engaged students are in MOOCs, whereas (D'mello & Graesser, 2013) examined emotion detection in self-directed learning with both vocal and non-vocal cues. Although the findings are promising, the majority of these approaches employed general classifiers or universal sentiment lexicons, which tend to misinterpret some words, formal terms, or ambiguous expressions that students employ.

More importantly, few educational sentiment systems provide confidence-aware outputs or allow for soft classification of mixed emotional signals. Instructors and adaptive systems benefit more from interpretable and graded sentiment indicators than from rigid class assignments, especially in high-stakes or sensitive contexts such as student confusion or demotivation. This highlights the need for frameworks that not only classify sentiment but also represent its strength, fuzziness, and reliability.

Fuzzy logic provides a compelling foundation for addressing these limitations. Rooted in the theory of approximate reasoning, fuzzy systems allow for the representation of vague, uncertain, or overlapping categories – such as “slightly negative” or “moderately positive” – that align more closely with human intuition. In sentiment analysis, fuzzy approaches have been used to assign degrees of polarity to opinions (Subasic & Huettnner, 2001), model emotional intensities in product reviews (Taboada et al., 2011), and handle ambiguous expressions in healthcare forums (Jadhav et al., 2024). These models typically use fuzzy inference rules, linguistic variables, and membership functions to map input features (e.g., word polarity, modifier strength) to output sentiment values on a continuous scale.

Several studies have proposed hybrid systems combining fuzzy logic with lexicon-based or machine learning methods for increased robustness. For instance, (Sun et al., 2025) combined fuzzy rules with SVM for Chinese sentiment classification, while (Ambreen et al., 2024) developed a fuzzy-based system for sentiment detection in online news articles. However, few studies have applied fuzzy reasoning in the educational domain, where interpretability, nuance, and domain adaptation are especially important.

A close related work to our approach includes the framework of (Anagha et al., 2015), which applies basic fuzzy rules to movie reviews, and the study by (Devi et al., 2024), which uses fuzzy sets to model sentiment confidence in e-commerce data. Neither system, however, is tailored to educational language, nor do they integrate semantic similarity or discourse-based modifiers as features.

This paper seeks to extend the application of fuzzy logic in sentiment recognition by introducing a fuzzy-weighted framework specifically designed for educational text-based interactions. By incorporating lexical polarity, syntactic modifiers, and contextual agreement into a unified fuzzy inference model, the proposed system addresses key gaps in interpretability, domain sensitivity, and confidence calibration – thereby advancing the state of the art in affective computing for education.

3 PROBLEM STATEMENT AND MOTIVATION

Student-generated text in online learning environments – discussion forums, peer feedback, reflective logs – carries implicit emotional cues that are rarely explicit yet critical to learner modeling. These cues signal fluctuating engagement, confusion, motivation, and frustration, which, if detected early and accurately, can inform timely interventions in adaptive learning platforms.

Traditional sentiment analysis tools, including both lexicon-based and neural models, tend to assign rigid labels (e.g., positive, negative), overlooking the contextual and affective complexity of educational discourse. For instance, feedback such as “Not bad, but I still feel lost with recursion” captures a simultaneous experience of improvement and confusion – feelings that are poorly represented by the bare categorical labels or polarity values alone.

The root problem lies in the existing systems' failure to validly account for the complex, unclear, or

overlapping emotional states found in student writing. This failure is especially problematic in learning environments, where both teachers and adaptive learning technologies require unambiguous and confidence-indicative measurements of student emotions in order to personalize support and manage learning trajectories.

Fuzzy logic is a formal framework for dealing with ambiguity. It allows for expressing sentiment on a continuum, where the interaction of modifier intensity, contextually related significance, and uncertainty produces results that are understandable to humans and hence suitable for both real-time and post-facto educational usage.

The research is motivated by the need to fill this gap. We present a fuzzy-weighted sentiment analysis approach tailored to interactions with educational content that excels in producing rich and understandable sentiment measures. Our goal is to move beyond mere binary sentiment classification and create useful affect modeling that enhances educational decision-making and helps student.

4 METHODOLOGY: THE FUZZY-WEIGHTED SENTIMENT ANALYSIS FRAMEWORK

The proposed fuzzy-weighted sentiment detection system seeks to analyze text-based interactions generated by students in an online learning environment. Such interactions can include contributions made in forums, peer reviews, weekly reflective journals, and unstructured feedback comments, all from a first-year undergraduate Java programming course conducted through Moodle. The main goal of the system is to derive interpretable sentiment measures—using fuzzy logic—that capture the emotional state of the learners, which can range from frustration to satisfaction, confusion, or interest. These measures can then be used to improve personalized feedback, deploy adaptive interventions, and guide learning analytics.

Unlike traditional sentiment classification models that produce discrete labels, the new model produces a continuous sentiment score $\sigma \in [-1, +1]$ accompanied by a confidence weight $\gamma \in [0, 1]$. This score reflects both the polarity and the strength of sentiment expressed in a student utterance, while the confidence weight indicates the degree of certainty in the inference, based on linguistic and contextual cues.

Each input text is first processed through a domain-sensitive linguistic pipeline. Tokenization and lemmatization are performed using spaCy, preserving key features of educational language. A rule-based module identifies negators, intensifiers (e.g., “extremely,” “barely”), hedging expressions (e.g., “I think,” “sort of”), and discourse markers. Syntactic parsing is used to highlight subject-verb-object relations and dependency chains, which are often crucial in interpreting affect in educational text. Additionally, we generate dense semantic representations using Sentence-BERT embeddings, which serve to measure the contextual alignment between a student's utterance and prototypical examples of positive, neutral, or negative sentiment.

The Sentence-BERT embeddings were used without additional fine-tuning on the educational dataset, relying on the pretrained all-mpnet-base-v2 model. While domain-specific fine-tuning may improve semantic similarity estimation, our primary goal was to preserve generalization and interpretability.

Thresholds for modifier strength and polarity adjustment were determined via grid search on a validation subset (20% of the corpus), optimizing for interpretability-consistent agreement with expert annotations.

From this processing, a feature vector $x = [S, M, A]$ is constructed for each sentence, where S denotes the lexical sentiment score, M represents modifier intensity, and A captures contextual agreement. The sentiment score SSS is computed as a weighted mean of the scores of matched terms from an augmented sentiment lexicon. This lexicon combines general-purpose entries from SentiWordNet and VADER with education-specific terms (e.g., “debugging,” “recursion,” “compile”) manually annotated by three expert raters. Each term has a polarity score $S_w \in [-1, +1]$, a confidence weight $C_w \in [0, 1]$, and a context tag (e.g., “evaluation”, “effort”, “difficulty”). Modifier intensity M is calculated as a normalized sum of the impact of linguistic intensifiers, diminishers, and negations detected in the sentence. Contextual agreement A is the cosine similarity between the input sentence's embedding and seed vector centroids for each sentiment category.

These features are fuzzified using piecewise linear membership functions. The lexical sentiment score SSS is mapped to five fuzzy categories: LowNegative, MediumNegative, Neutral, MediumPositive, and HighPositive. Modifier intensity and contextual agreement are similarly mapped to fuzzy sets: Weak, Medium, Strong and Low, Medium, High respectively. Below is the

complete definition of the membership functions for lexical polarity:

- LowNegative

$$\mu_{LN}(S) = \begin{cases} 1, & S \leq -0.8 \\ \frac{-0.4 - S}{0.4}, & -0.8 \leq S \leq -0.4 \\ 0, & S > -0.4 \end{cases} \quad (1)$$

- MediumNegative

$$\mu_{MN}(S) = \begin{cases} 0, & S \leq -0.8 \text{ or } S \geq 0 \\ \frac{S + 0.8}{0.4}, & -0.8 \leq S \leq -0.4 \\ \frac{-S}{0.4}, & 0.4 < S < 0 \end{cases} \quad (2)$$

- Neutral

$$\mu_{NEU}(S) = \begin{cases} 0, & |S| \geq 0.6 \\ 1 - \frac{|S|}{0.6}, & |S| < 0.6 \end{cases} \quad (3)$$

- MediumPositive

$$\mu_{MP}(S) = \begin{cases} 0, & S \leq 0 \text{ or } S \geq 0.8 \\ \frac{S}{0.4}, & 0 < S \leq 0.4 \\ \frac{0.8 - S}{0.4}, & 0.4 < S < 0.8 \end{cases} \quad (4)$$

- HighPositive

$$\mu_{HP}(S) = \begin{cases} 0, & S \leq 0.4 \\ \frac{S - 0.4}{0.4}, & 0.4 < S \leq 0.8 \\ 1, & S > 0.8 \end{cases} \quad (5)$$

The fuzzy inference engine applies a set of 27 expert-defined rules over these inputs. Each rule has the general structure:

IF Lexical Polarity is X AND Modifier Intensity is Y AND Contextual Agreement is Z THEN Sentiment Output is C

For instance:

- IF Lexical Polarity is MediumNegative AND Modifier Intensity is Strong AND Contextual Agreement is High THEN Output is NegativeStrong
- IF Lexical Polarity is Neutral AND Modifier Intensity is Weak AND Contextual Agreement is Medium THEN Output is Neutral

Each rule produces a fuzzy output set (e.g., NegativeStrong, PositiveWeak) with an associated degree of membership. Using the minimum operator for rule activation and centroid defuzzification, the final sentiment score σ is computed as:

$$\sigma = \frac{\sum_{i=1}^N \mu_i \cdot y_i}{\sum_{i=1}^N \mu_i} \quad (6)$$

where μ_i is the activation degree of the i -th rule and y_i is its associated output score (e.g., -0.8 for NegativeStrong, +0.6 for PositiveWeak). The system confidence in its prediction is taken as:

$$\gamma = \max_i \mu_i \quad (7)$$

For longer textual entries such as forum posts containing multiple sentences, sentence-level sentiment predictions are aggregated using a weighted average:

$$\Sigma = \frac{\sum_{j=1}^n w_j \cdot \sigma_j}{\sum_{j=1}^n w_j}, \quad \Gamma = \frac{1}{n} \sum_{j=1}^n \gamma_j \quad (8)$$

where w_j is a sentence weight derived from TF-IDF scores and discourse role (e.g., conclusion, elaboration). This produces a final post-level sentiment profile (Σ, Γ) , interpretable as “moderately positive with medium confidence” or similar qualitative labels.

This hybrid fuzzy system offers three advantages in the educational context: (1) it handles ambiguity and mixed affect naturally, (2) it avoids the opacity of deep learning classifiers, and (3) it enables human-readable outputs that educators and adaptive systems can interpret and act upon. The model is not trained via backpropagation but is manually calibrated using a development set of annotated educational texts, allowing it to generalize across similar learning environments without requiring large-scale supervised data.

5 EVALUATION RESULTS AND DISCUSSION

To assess the performance and practical viability of the proposed fuzzy-weighted sentiment recognition framework, we conducted a comprehensive empirical evaluation using authentic data collected from a university-level Java programming course. The course was delivered over a 12-week semester using a Moodle-based web platform and included weekly assignments, peer discussion activities, and reflective tasks. In total, 96 undergraduate students participated in the course, generating 2,350 textual messages. These comprised 1,470 posts and replies in asynchronous discussion forums, 580 short reflective responses to weekly prompts, and 300 comments submitted as part of the final course feedback.

Each message was anonymized and manually annotated for sentiment polarity by three trained human raters. Annotations included both a categorical sentiment label (positive, neutral, or negative) and a

confidence score on a 3-point scale (low, medium, high). To ensure annotation quality, inter-rater agreement was measured using Krippendorff's alpha, which yielded a coefficient of 0.81—indicating substantial agreement. A stratified subset of 800 messages was reserved as the evaluation set. These messages were balanced across sentiment classes and served as the gold standard for testing all models.

In this evaluation, our proposed fuzzy-weighted model was compared with five representative baseline systems. These include: (1) a simplified fuzzy logic implementation (FuzzyLex), which uses only polarity scores and fixed modifier weights; (2) VADER, a rule-based model optimized for social media text; (3) TextBlob, a naive Bayes classifier with a general-purpose sentiment lexicon; (4) a BERT model fine-tuned on the Stanford Sentiment Treebank (SST-2) and lightly adapted with domain-specific examples; and (5) an LSTM-based neural model trained on 1,000 manually labeled student texts from the same course domain. All models were evaluated under the same conditions and tested on the same evaluation set to ensure consistency and fairness.

To capture different aspects of model performance, we employed four evaluation metrics. First, we computed the macro-averaged F1 score to evaluate classification accuracy across the three sentiment classes. Second, we calculated the mean absolute error (MAE) between predicted sentiment scores and the human-annotated confidence-weighted ground truth. Third, we assessed interpretability, a critical factor in educational applications, by asking three experienced educators to rate the clarity and pedagogical value of each model's output on a 5-point Likert scale. Fourth, we measured robustness to paraphrasing by evaluating each model on a curated subset of 100 sentiment-preserving paraphrases derived from the original texts.

The quantitative results are summarized in Table 1. It presents each model's performance across all four metrics. As shown, the proposed fuzzy-weighted system achieved an F1 score of 0.81, closely approaching the 0.84 of the fine-tuned BERT model, and outperforming all other baselines. The fuzzy model also yielded a low MAE of 0.11, nearly matching BERT's 0.10, and significantly outperforming VADER (0.23) and TextBlob (0.25). These results indicate that the fuzzy system offers competitive classification performance while maintaining lower prediction error.

Table 1: Evaluation Results.

Method	F1 Score	MAE	Interpretability Score	Robustness (Accuracy)
Fuzzy-Weighted (Proposed)	0.81	0.11	4.7	0.83
FuzzyLex Baseline	0.72	0.18	4.2	0.75
VADER	0.66	0.23	2.0	0.61
TextBlob	0.62	0.25	2.3	0.57
BERT Fine-tuned	0.84	0.10	1.2	0.79
LSTM (domain-tuned)	0.79	0.13	2.0	0.76

While the performance differences in F1 and MAE are noteworthy, perhaps more significant are the results concerning interpretability and robustness – two criteria of particular importance in educational systems. As shown in Figure 1, the fuzzy-weighted model received the highest interpretability rating (4.7/5) from domain experts, far exceeding the ratings of black-box models such as BERT (1.2) and the LSTM variant (2.0). Educators noted that the fuzzy model's scalar sentiment score, coupled with its confidence output and linguistic justification (e.g., influence of modifiers), allowed them to better understand and act upon the sentiment output.




Figure 1: Interpretability Ratings for All Models here.

Robustness to paraphrasing, depicted in Figure 2, further demonstrates the reliability of the fuzzy approach. On the paraphrased subset, where students expressed similar sentiment using alternate phrasing, the fuzzy model maintained a robust accuracy of 0.83, outperforming VADER (0.61), TextBlob (0.57), and even slightly surpassing BERT (0.79). This indicates that the rule-guided fuzzy inference engine, though

not pretrained on massive corpora, exhibits strong generalization capacity in the face of surface linguistic variation.




Figure 2: Accuracy on Paraphrased Inputs.

To visualize the overall classification performance, Figure 3 presents a bar plot of macro F1 scores across all models. The fuzzy-weighted system, while slightly behind BERT in raw accuracy, clearly outperforms both classical and rule-based baselines and provides substantially more explainable output. This balance of performance and interpretability suggests that fuzzy reasoning is particularly well-suited for affective modeling in education, where transparency and trust are necessary for practical deployment.




Figure 3: Comparison of Macro F1 Scores.

The results of the evaluation reveal that the fuzzy-weighted sentiment recognition model is an effective and pedagogically appropriate method to perform affective analysis in text-based communication for educational purposes. Having attained a high degree of accuracy in classification and with minimal prediction errors, shown to be robust against linguistic heterogeneity, and offering significant interpretability to educators, this model is considered especially appropriate for implementation in adaptive learning systems, feedback dashboards, or real-time engagement monitoring tools.

The findings not only demonstrate the computational efficiency of the fuzzy-weighted sentiment recognition system but also provide insight into its widespread applicability and limitations in terms of online learning environments. Specifically, the system's ability to identify varied levels of sentiment and provide interpretable results speaks to the educational need for transparency and actionable information within analytics. In contrast to deep learning models that might achieve superior raw accuracy but not interpretability, the fuzzy approach offers linguistically grounded explanations for each classification outcome. This feature is particularly salutary for uses requiring human engagement, like educator dashboards and formative assessment tools. At the same time, the analysis identifies several limitations. The rule base and lexicon are currently designed to effectively treat programming-specific language; however, while promising, continual refinement might be necessary in order to address a broader range of domains or more casual communicative contexts. In addition, while the fuzzy model effectively handles the vagueness of sentiment, it does not at present allow temporality or the changing states of learners over time. These findings suggest that future research would be improved by combining fuzzy reasoning with context-sensitive neural designs or sequence-based approaches to sentiment monitoring, thereby enabling more adaptive and longitudinal models of affect.

6 CONCLUSIONS AND FUTURE WORK

The present work presented a fuzzy-weighted sentiment analysis model that combines lexical polarity, linguistic modifiers, and contextual agreement in a Mamdani-type fuzzy inference system. The proposed framework produces interpretable and fine-grained sentiment ratings. An empirical study using data from an online discussion forum for a university-level computer programming course showed that the proposed approach is competitive under the F1 measure and mean absolute error, well outperforming baseline models under interpretability and robustness. These results confirm the effectiveness of fuzzy logic in capturing the emotional nuances present in pedagogical discourse, which often contains features such as subtlety, ambivalence, or context-dependent affective terms. The explainable reasoning of the model and its educational value make it an appropriate candidate

for potential deployment in adaptive learning systems, reflective feedback mechanisms, and monitoring via learner dashboards.

The fuzzy-weighted approach offers significant advantages in interpretability and flexibility, but many directions for future work are still to be pursued. Foremost, the expansion of the domain-specific sentiment lexicon to include a wider variety of academic disciplines and communication modalities (e.g., chat posts or transcribed voice communications) would enhance its applicability. Second, the use of hybrid approaches that combine fuzzy reasoning with transformer-based contextual embeddings may improve handling of complex semantic structures and figurative language while still allowing for some level of explainability. Third, conducting longitudinal studies of affect over time—by following emotional trajectories rather than examining only discrete posts in isolation—may provide deeper insight into engagement patterns and learning trajectories. Finally, incorporation of this approach into intelligent tutoring systems or massive open online courses (MOOCs), along with user-centered validation, would enable empirical testing of the model's performance in real-time educational interventions and decision-making applications.

REFERENCES

- Ahmed, Y., Kent, S., Cirino, P. T., & Keller-Margulis, M. (2022). The not-so-simple view of writing in struggling readers/writers. *Reading & Writing Quarterly*, 38(3), 272–296. <https://doi.org/10.1080/10573569.2021.1948374>
- Alahmadi, K., Alharbi, S., Chen, J., & Wang, X. (2025). Generalizing sentiment analysis: A review of progress, challenges, and emerging directions. *Social Network Analysis and Mining*, 15(1), 45. <https://doi.org/10.1007/s13278-025-01461-8>
- Altrabsheh, N., Cocca, M., & Fallahkhair, S. (2014). Sentiment analysis: Towards a tool for analysing real-time students' feedback. 2014 IEEE 26th International Conference on Tools with Artificial Intelligence (ICTAI), 419–423. <https://doi.org/10.1109/ICTAI.2014.70>
- Ambreen, S., Iqbal, M., Asghar, M. Z., Mazhar, T., Khattak, U. F., Khan, M. A., & Hamam, H. (2024). Predicting customer sentiment: The fusion of deep learning and a fuzzy system for sentiment analysis of Arabic text. *Social Network Analysis and Mining*, 14(1), 206. <https://doi.org/10.1007/s13278-024-01356-0>
- Anagha, M., Kumar, R. R., Sreetha, K., & Reghu Raj, P. C. (2015). Fuzzy logic-based hybrid approach for sentiment analysis of Malayalam movie reviews. 2015 IEEE International Conference on Signal Processing, Informatics, Communication and Energy Systems (SPICES), 1–4. <https://doi.org/10.1109/SPICES.2015.7091512>
- Benazer, S. S., Thangamani, A. K. S., & Prabu, S. (2024). Integrating sentiment analysis with learning analytics for improved student engagement. In *Advancement of IoT in Blockchain Technology and Its Applications* (pp. 1262–1284). Zenodo. <https://doi.org/10.5281/zenodo.13752749>
- Devi, R., Prabu, S., & Legapriyadharshini, N. (2024). Sentiment analysis for the improved user experience in the e-commerce platform with the fuzzy model. *Journal of Computer Allied Intelligence*, 2, 28–40. <https://doi.org/10.69996/jcai.2024013>
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)* (pp. 4171–4186). <https://doi.org/10.48550/arXiv.1810.04805>
- D'Mello, S., & Graesser, A. (2013). AutoTutor and affective AutoTutor: Learning by talking with cognitively and emotionally intelligent computers that talk back. *ACM Transactions on Interactive Intelligent Systems*, 2(4), Article 23. <https://doi.org/10.1145/2395123.2395128>
- Grimalt-Álvaro, C., & Usart, M. (2024). Sentiment analysis for formative assessment in higher education: A systematic literature review. *Journal of Computing in Higher Education*, 36(3), 647–682. <https://doi.org/10.1007/s12528-023-09370-5>
- Hafner, F. S., Hafner, L., & Corizzo, R. (2025). 'Slightly disappointing' vs. 'worst sh** ever': Tackling cultural differences in negative sentiment expressions in AI-based sentiment analysis. *Journal of Computational Social Science*, 8(3), 57. <https://doi.org/10.1007/s42001-025-00382-y>
- Hutto, C., & Gilbert, E. (2014). VADER: A parsimonious rule-based model for sentiment analysis of social media text. *Proceedings of the International AAAI Conference on Web and Social Media*, 8(1), 216–225. <https://doi.org/10.1609/icwsm.v8i1.14550>
- Jadhav, S., Shiragapur, B., Purohit, R., & others. (2024). A systematic review on the role of sentiment analysis in healthcare. *Authorea*. <https://doi.org/10.22541/au.172516418.81416769/v1>
- Johansen, M. O., Eliassen, S., & Jenö, L. M. (2025). Autonomy need satisfaction and frustration during a learning session affect perceived value, interest, and vitality among higher education students. *Scandinavian Journal of Educational Research*, 69(4), 757–771. <https://doi.org/10.1080/00313831.2024.2348457>
- Kardaras, D. K., Troussas, C., Barbounaki, S. G., Tselenti, P., & Armyras, K. (2024). A fuzzy synthetic evaluation approach to assess usefulness of tourism reviews by considering bias identified in sentiments and articulacy. *Information*, 15(4), 236. <https://doi.org/10.3390/info15040236>

- Litman, D. J., & Forbes-Riley, K. (2006). Recognizing student emotions and attitudes on the basis of utterances in spoken tutoring dialogues with both human and computer tutors. *Speech Communication*, 48(5), 559–590. <https://doi.org/10.1016/j.specom.2005.09.008>
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., & Stoyanov, V. (2019). RoBERTa: A robustly optimized BERT pretraining approach. arXiv preprint. <https://doi.org/10.48550/arXiv.1907.11692>
- Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval*, 2(1–2), 1–135. <https://doi.org/10.1561/15000000011>
- Subasic, P., & Huettner, A. (2001). Affect analysis of text using fuzzy semantic typing. *IEEE Transactions on Fuzzy Systems*, 9(4), 483–496. <https://doi.org/10.1109/91.940962>
- Sun, Y., Yu, Z., Sun, Y., Xu, Y., & Song, B. (2025). A novel approach for multiclass sentiment analysis on Chinese social media with ERNIE-MCBMA. *Scientific Reports*, 15(1), 18675. <https://doi.org/10.1038/s41598-025-03875-y>
- Taboada, M., Brooke, J., Tofiloski, M., Voll, K., & Stede, M. (2011). Lexicon-based methods for sentiment analysis. *Computational Linguistics*, 37(2), 267–307. https://doi.org/10.1162/COLI_a_00049
- Tasoulas, T., Troussas, C., Mylonas, P., & Sgouropoulou, C. (2024). Affective computing in intelligent tutoring systems: Exploring insights and innovations. 2024 SEEDA-CECNSM Conference Proceedings, 1–6. <https://doi.org/10.1109/SEEDA-CECNSM63478.2024.00025>
- Troussas, C., Krouska, A., & Virvou, M. (2019). Trends on sentiment analysis over social networks: Pre-processing ramifications, stand-alone classifiers and ensemble averaging. In G. A. Tsihrintzis, D. N. Sotiropoulos, & J. L. C. de Carvalho (Eds.), *Machine learning paradigms: Advances in data analytics* (pp. 161–186). Springer. https://doi.org/10.1007/978-3-319-94030-4_7
- van der Veen, A. M., & Bleich, E. (2025). The advantages of lexicon-based sentiment analysis in an age of machine learning. *PLOS ONE*, 20(1), e0313092. <https://doi.org/10.1371/journal.pone.0313092>
- Wen, M., Yang, D., & Rose, C. P. (2014). Sentiment analysis in MOOC discussion forums: What does it tell us? *Proceedings of the 7th International Conference on Educational Data Mining (EDM 2014)*, 1–8. <http://www.cs.cmu.edu/~mwmen/papers/edm2014-camera-ready.pdf>
- Yang, D., Sinha, T., Adamson, D., & Rosé, C. P. (2013). Turn on, tune in, drop out: Anticipating student dropouts in massive open online courses. *Proceedings of the 2013 NIPS Data-Driven Education Workshop*, 11–14.
- Yuvaraj, R., Mittal, R., Prince, A. A., & Huang, J. S. (2025). Affective computing for learning in education: A systematic review and bibliometric analysis. *Education Sciences*, 15(1), 65. <https://doi.org/10.3390/educsci15010065>