# ISALE: A Neuro-Symbolic Framework for Concealed Psychological Risk Assessment in K-12 Students

Xinyin Hu*

College of Education, Zhejiang University of Technology, Hangzhou 310023, China

*302023572080@zjut.edu.cn

**Abstract.** Traditional automated sentiment analysis often fails to detect concealed emotional distress in K-12 students due to the subtle nature of their metaphorical expressions. This paper presents ISALE, a neuro-symbolic framework that integrates a Neo4j knowledge graph with the DeepSeek-V3 large language model. By employing a dual-pathway weighted fusion mechanism, ISALE maps unstructured student narratives to structured psychological metaphors to calculate a comprehensive risk index. Preliminary evaluation on 25 samples demonstrates an 88% precision rate, with high expert alignment (Pearson $r=0.8626$, $p < 0.01$; Quadratic Weighted Kappa 0.8758). These results indicate that bridging neural perception with symbolic knowledge significantly enhances the interpretability and accuracy of adolescent psychological screening.

**Keywords:** Affective Computing; Knowledge Graph; Neuro-Symbolic AI; Psychological Risk Assessment; Metaphorical Recognition; K-12 Education

## 1 Introduction

Education informatization is entering an intelligent phase, with growing focus on AI-driven mental health prevention. Current authoritative data indicates that psychological issues among students are becoming increasingly younger and more concealed, posing a significant threat to K-12 well-being. To align with the "15th Five-Year Plan" and the Healthy China initiative, there is an urgent need for proactive and scientific screening mechanisms. However, current prevention relies heavily on the subjective judgment of non-professional teachers (e.g., class instructors). This approach is not only time-consuming but also prone to missing hidden distress due to a lack of clinical psychological expertise among frontline educators.

While automated detection using isolated NLP models like BERT is an emerging trend[5], these tools often only capture surface-level sentiment.They struggle to decode concrete metaphorical expressions,such as spatial metaphors (e.g., "living in an abyss") or somatic/physical metaphors (e.g., "my chest is filled with lead"),frequently used by children to articulate internal distress. Consequently, these models often misclassify such texts as neutral, lacking the structured guidance required for effective educational intervention. To bridge these gaps, this research introduces the ISALE (Integrated Symbolic-neural Affective Learning Engine) system featuring a layered


C. V. Garcés et al. (eds.), *Proceedings of the 6th International Conference on Internet, Education and Information Technology (IEIT 2026)*, Atlantis Highlights in Social Sciences, Education and Humanities 53,
https://doi.org/10.2991/978-94-6239-733-0_13

symbolic-neural architecture: a Knowledge-Symbolic Layer for mapping structured psychological symptoms via a specialized Knowledge Graph (KG), a Cognitive Reasoning Layer for dynamic emotion recognition and hybrid diagnostic inference, and an Interactive Visualization Layer that transforms complex backend data into intuitive evidence for educational intervention. Experimental evaluations demonstrate that the system achieves high diagnostic accuracy and significantly optimizes the preliminary screening workflow.

# 2 System design

## 2.1 Overall System Architecture of ISALE

The ISALE system is engineered as a cloud-native, hybrid symbolic-neural intelligence platform. To overcome the limitations of traditional isolated NLP models, the architecture integrates deterministic knowledge graph reasoning with generative cognitive analysis. The system follows a modular design consisting of three core functional layers: the Knowledge-Symbolic Layer, the Hybrid Diagnostic Engine, and the Interactive Presentation Layer(as illustrated in Fig. 1)[1].

**Knowledge-Symbolic Layer:** This layer serves as the structural foundation, utilizing Neo4j Aura DB to host a specialized mental health knowledge graph. It stores pre-defined psychological risk dimensions and metaphorical mapping rules. By employing Cypher query optimization, this layer provides stable and interpretable "symbolic evidence" to support subsequent AI reasoning.

**Hybrid Diagnostic Engine:** As the central processing unit, this engine orchestrates the interaction between DeepSeek-V3 (via RESTful APIs) and local sentiment analysis modules. It implements a Hybrid Scoring Mechanism that weights structural graph hits against nuanced semantic interpretations, allowing the system to decode complex, concealed metaphors in student narrative texts that surface-level sentiment analysis often overlooks.

**Interactive Presentation Layer:** Developed using the Streamlit framework, this layer handles asynchronous data coordination and visualization. It leverages streamlit-agraph for real-time graph rendering, providing educators with a "bridging visualization" that transforms complex backend computations into intuitive diagnostic evidence.

The system is deployed using a serverless cloud paradigm, ensuring high availability and elastic scaling without the overhead of local hardware maintenance. Communication between modules is implemented via secured REST APIs and environment-level encryption. By utilizing asynchronous streaming for report generation, the architecture ensures a responsive user experience even during complex multi-dimensional analysis, maintaining high service reliability across diverse campus network environments.

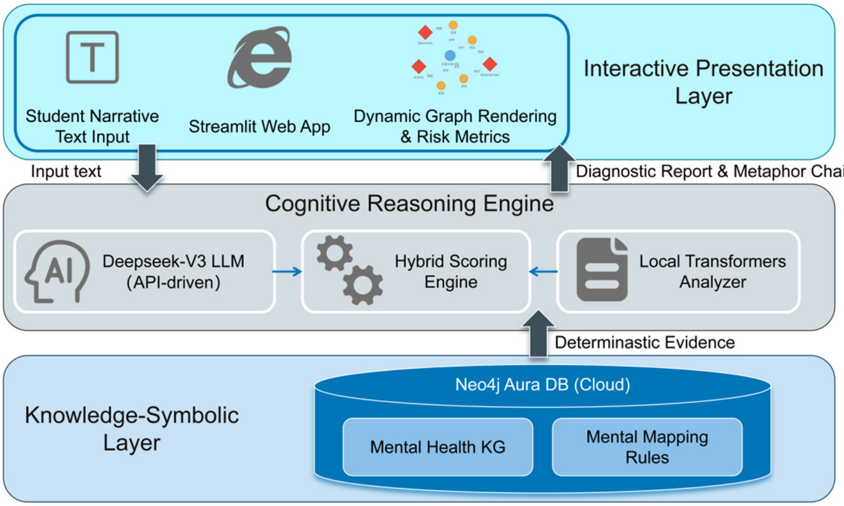


**Fig. 1.** The ISALE Layered Symbolic-Neural Architecture

## 2.2 Mental Metaphor Knowledge Graph Construction

The ISALE system employs the graph database Neo4j to construct and store the specialized mental metaphor knowledge graph. The construction of this knowledge base is strictly grounded in authoritative clinical frameworks, such as the DSM-5 (Diagnostic and Statistical Manual of Mental Disorders) and standard psychological assessment scales[2]. To bridge the gap between clinical diagnostic criteria and unstructured student writings, the graph translates complex psychological states into a structured topological network.

Currently, the prototype graph contains 63 precise knowledge nodes and 102 relationship edges. The nodes are logically categorized into two primary labels: 6 "Risk" nodes (representing core psychological distress dimensions derived from the DSM-5, such as depressive mood or anxiety) and 57 "Metaphor" nodes (capturing the implicit textual imagery frequently found in K-12 narrative texts, such as "abyss" or "withered"). The relationships between these entities are explicitly defined using "INDICATES" edges (102 in total), establishing a direct, quantitative mapping from a specific linguistic metaphor to its corresponding clinical risk dimension. This evidence-based extraction ensures high diagnostic validity. Furthermore, the graph database supports dynamic expansion, allowing for the continuous integration of new metaphor-to-risk mappings as more educational text data is processed[3]. Figure 2 illustrates the topological overview of the currently deployed knowledge graph.

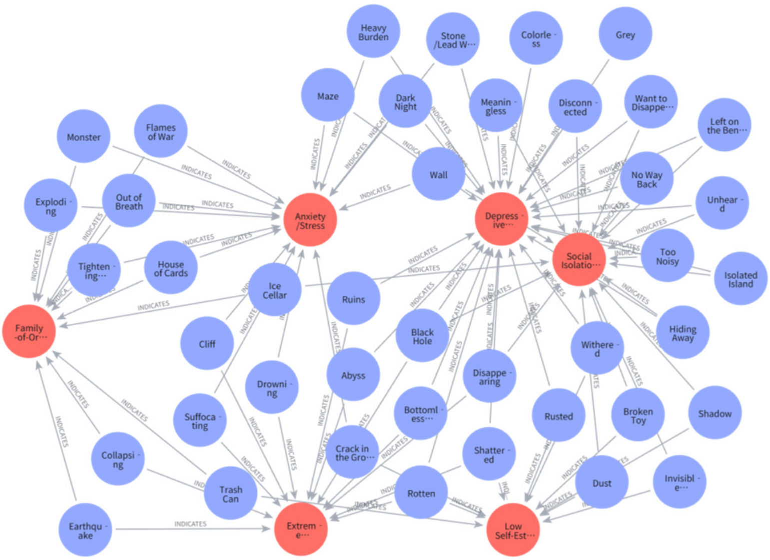


**Fig. 2.** Visualization of the Neuro-Symbolic Knowledge Mapping in ISALE.

### 2.3 Affective Computing and Linguistic Extraction Module

To implement the integration of affective computing and symbolic knowledge, ISALE employs a hybrid scoring engine[4]. The system first aggregates a base penalty for negative sentiment ($P_{sent} = 0.3$) and weights from Neo4j KG hits ($W_{hit} = 0.25\ per\ hit$) to capture the total risk energy ($E$). An exponential decay function is then applied to perform non-linear normalization of the objective graph-based score ($S_{graph}$), ensuring the value asymptotically approaches 1.0. This reflects the clinical characteristic of diminishing marginal utility in risk certainty:

$$S_{graph} = 1 - e^{-E}, \quad where\ E = P_{sent} + \sum W_{hit}$$

Ultimately, an Equal-Weight Fusion algorithm integrates the deterministic graph-based score ($S_{graph}$) with the LLM agent's deep semantic score ($S_{agent}$) to calculate the final risk index ($S_{final}$)[1]:

$$S_{final} = 0.5 \times S_{graph} + 0.5 \times S_{agent}$$

This hybrid architecture ensures that the system maintains high sensitivity toward complex rhetorical expressions (via the agent) while remaining anchored by clinical evidence and interpretability (via the graph), thereby significantly reducing the risk of false positives in automated psychological screening.

# 3 System Testing and Preliminary Evaluation

## 3.1 Test Environment and Methodology

To verify the feasibility of the ISALE framework, we established a high-performance computational environment for deployment and testing. The system's backend, including the DeepSeek-V3 engine and Neo4j database, was deployed on a server equipped with an NVIDIA RTX 4090 GPU (24GB VRAM), an Intel Core i9-13900K CPU, and 64GB of RAM to ensure low-latency processing of natural language inputs.

Given the high sensitivity and privacy constraints of minor mental health data, we constructed a pilot testing corpus comprising 25 high-fidelity synthetic essays. These samples were generated using a structured prompt strategy designed to simulate the linguistic patterns and metaphorical preferences of K-12 students. The corpus specifically incorporates subtle distress markers and complex metaphors, providing a rigorous baseline for assessing the system's diagnostic accuracy compared to human expert consensus[6].

## 3.2 Preliminary Case Study and Validity Analysis

**Qualitative Case Analysis.**

To demonstrate the superiority of the ISALE neuro-symbolic approach, we analyzed a representative case:"The sun is bright, but I am living in a room without windows."

Traditional Baseline (BERT-based): A standard sentiment analysis model (e.g., RoBERTa-base) primarily focuses on surface-level keywords like "bright" and "sun," mistakenly classifying the text as "Neutral" or even "Slightly Positive." It fails to decode the spatial metaphor of "a room without windows."BERT models often rely on surface-level valence (sentiment scores), failing to map the source domain (physical objects) to the target domain (psychological states).

ISALE Diagnosis: Through the fetch_mkg_context interface, ISALE successfully identified the metaphorical entity "Room without windows" and matched it to the "Social Isolation" and "Depressive Tendency" nodes in the Neo4j Knowledge Graph. By integrating this symbolic evidence with the LLM's semantic reasoning, the system assigned a high-risk score of 0.82, accurately flagging the hidden emotional distress that traditional models overlooked.This aligns with recent findings that children's metaphor comprehension is a critical indicator of their internal cognitive state[7].

**Quantitative Performance and Validity.**

The system's performance on the 25-sample pilot dataset confirms its robustness in identifying concealed psychological issues. Quantitatively, the ISALE system achieved a Precision of 88.0% in identifying high-risk cases.

As detailed in our reliability analysis (Table 1), the high Quadratic Weighted Kappa of 0.8758 and the strong Pearson Correlation of 0.8626 further validate the system's alignment with expert judgments. The trend consistency between the automated ratings and the expert consensus mean is visualized in Fig. 3. These preliminary results demonstrate that by bridging the gap between neural sentiment perception and symbolic knowledge validation, ISALE significantly enhances the validity of automated psychological screening in educational settings.

Furthermore, a preliminary error analysis of false positives and false negatives reveals the practical robustness of the system. Traditional LLMs often generate false positives by misinterpreting common hyperbolic expressions (e.g., "I am drowning in homework") as severe distress. ISALE mitigates this: if the LLM flags a neutral metaphor, but the symbolic knowledge graph (KG) detects no corresponding clinical nodes, the dual-pathway fusion scales down the final risk score, effectively reducing false alarms. Conversely, occasional false negatives occurred primarily when students employed highly idiosyncratic or localized slang not yet mapped in the current prototype graph.

**Table 1.** Performance Metrics for Affective Computing Module (N=25)

| Metric Category | Evaluation Metric | Value | Interpretation/Level |
|---|---|---|---|
| Trend | Pearson Correlation | 0.8626 | High Correlation |
| Accuracy | Mean Absolute Error | 0.4000 | Low Error |
| Agreement | Simple Kappa | 0.5455 | Moderate Agreement |
| Weighting | Quadratic Weighted Kappa | 0.8758 | Almost Perfect |

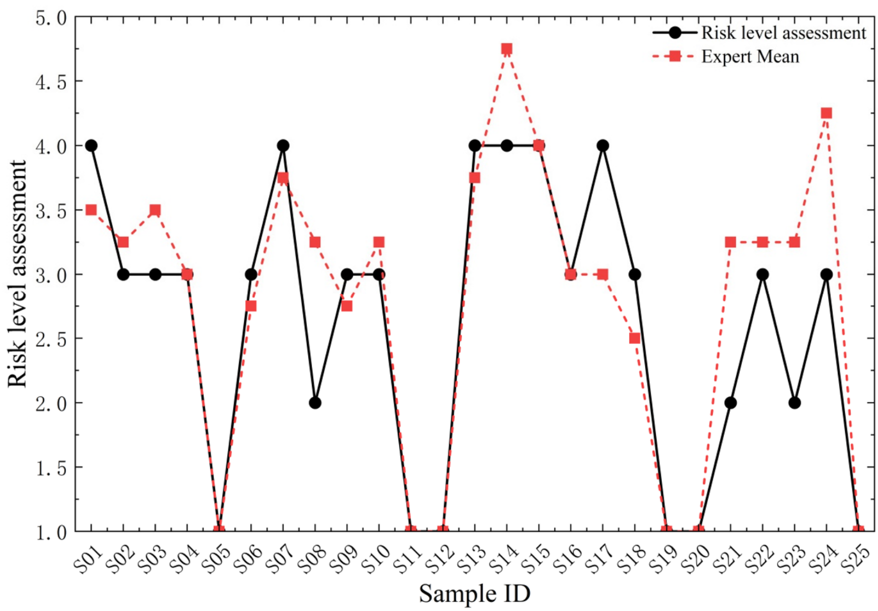


**Fig. 3.** Trend consistency analysis between ISALE scores and expert consensus mean (N=25).

The plot illustrates the scoring fluctuations across 25 test samples. The high degree of overlap between the ISALE automated ratings and the expert mean curve (Pearson $r=0.8626$) validates the system's evaluative stability across varying risk intensities.

### 3.3 Analysis of Current Limitations

Despite its promising performance, several limitations of the ISALE system warrant further investigation. First, the current validation relies on a small-scale pilot corpus ($N = 25$) of synthetic samples; while effective for initial testing, it may not fully capture the linguistic diversity and cultural nuances of real-world student writing across different regions. Second, the symbolic knowledge graph, while accurate, currently covers only a prioritized set of psychological metaphors, potentially limiting its recall for localized or evolving slang. Finally, the affective module utilizes a generalized DeepSeek-V3 engine without domain-specific fine-tuning on massive K-12 psychological datasets. To mitigate potential domain shift, the current system heavily relies on advanced prompt engineering, specifically, injecting expert-persona instructions and dynamically retrieving contexts from the Neo4j graph (acting as a limited Retrieval-Augmented Generation paradigm), to explicitly anchor the LLM's reasoning within clinical boundaries. Future research will focus on expanding the knowledge base and fine-tuning the transformer layers to enhance diagnostic robustness in authentic educational environments.

## 4 Conclusion

*System Availability:* The ISALE system is deployed and accessible at https://isalemax.streamlit.app/ Users can access the diagnostic dashboard for research purposes.

The ISALE (Metaphorical-Knowledge-Graph-based Affective Computing) system has demonstrated significant potential in the automated screening of concealed psychological distress among K-12 students. In pilot trials with 25 diverse narrative samples, the system achieved an 88% precision rate in high-risk case identification and maintained an "almost perfect" agreement with expert consensus (Quadratic Weighted Kappa of 0.8758). By bridging the gap between neural semantic perception and symbolic knowledge validation, the system successfully identified complex psychological metaphors that traditional NLP models frequently overlooked. Supported by a specialized knowledge graph and the DeepSeek-V3 cognitive engine, ISALE provides a highly interpretable and evidence-based diagnostic framework for educational practitioners. Future development will focus on expanding the metaphorical knowledge base to include regional cultural nuances, fine-tuning the transformer layers on authentic adolescent psychological datasets, and integrating multi-modal interaction analysis to further enhance diagnostic robustness.

## References

1. Debets, T., Banihashem, S. K., et al. (2025). Chatbots in education: A systematic review of objectives, underlying technology and theory, evaluation criteria, and impacts. Comput. Educ., 234, 105323. https://doi.org/10.1016/j.compedu.2025.105323

2. Wang, Q., Yang, F., Quan, L., Fu, M., Yang, Z., & Wang, J. (2025). Knowledge graph and its application in the study of neurological and mental disorders. Front. Psychiatry, 16, 1452557. https://doi.org/10.3389/fpsyt.2025.1452557
3. Park, C., Lee, H., & Jeong, O. (2024). Leveraging Medical Knowledge Graphs and Large Language Models for Enhanced Mental Disorder Information Extraction. Future Internet, 16(8), 260. https://doi.org/10.3390/fi16080260
4. Schlicher, M., Li, Y., Murthy, S. M. K., Sun, Q., & Schuller, B. W. (2025). Emotionally adaptive support: a narrative review of affective computing for mental health. Front. Digit. Health, 7. https://doi.org/10.3389/fdgth.2025.1657031
5. Hossain, M. M., Hossain, M. S., Mridha, M. F., et al. (2025). Multi task opinion enhanced hybrid BERT model for mental health analysis. Sci. Rep., 15, 3332. https://doi.org/10.1038/s41598-025-86124-6
6. Themistocleous, C. K., Andreou, M., & Peristeri, E. (2024). Autism Detection in Children: Integrating Machine Learning and Natural Language Processing in Narrative Analysis. Behav. Sci. (Basel), 14(6), 459. https://doi.org/10.3390/bs14060459
7. Starr, A., Petersen, T., & Sridhar, S. (2025). Metaphor comprehension in preschool children. Language Development Research, 5(2), 1–27. https://doi.org/10.34842/starr