



International Journal of Innovative Research in Science Engineering and Technology (IJIRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)



Impact Factor: 8.699

Volume 15, Issue 3, March 2026

The Academic Architect: A Multi-Agent AI Framework for Personalized Learning Execution and Longitudinal Student Development

Dr. G. Sumathi¹, Anupriya A², Sasirekha S³, Swathi S⁴

Associate Professor, Department of Artificial Intelligence and Data Science, Kongunadu College of Engineering and Technology, Trichy, India¹

Department of Artificial Intelligence and Data Science, Kongunadu College of Engineering and Technology, Trichy, India^{2,3,4}

ABSTRACT: Students today have access to many digital platforms for learning and career development. However, most of these platforms work as passive and fragmented systems. They often do not provide consistent support, regular monitoring of academic progress, and specific guidance. As a result, students frequently find it hard to keep track of their academic progress, skill development, and career preparation during their undergraduate years. This creates a need for a system that offers integrated and independent support over the long term. This project introduces an Agent-Based Autonomous Student Academic, Learning, and Career Intelligence Platform. It uses a coordinated multi-agent artificial intelligence structure to provide comprehensive support for students from their first semester to graduation. The system manages learning and career paths based on the academic field. It ensures progress through adaptive learning and discipline methods. It also includes resume analysis, job and internship searches, and application help. The platform uses predictive analytics and explainable artificial intelligence to spot academic risks and behavior changes, while also ensuring transparency. Through a comprehensive synthesis of recent research from 2021 to 2026, this paper demonstrates that agentic AI systems represent a paradigm shift from static, prompt-response educational tools to dynamic, autonomous learning companions. Key findings indicate that multi-agent architectures with specialized sub-agents achieve intent recognition accuracy exceeding 90% and learning plan effectiveness above 90%. The integration of predictive analytics with early warning systems enables accurate identification of at-risk students, with Deep Reinforcement Learning (DRL) models demonstrating significant capability in predicting performance declines and dropout risks. The proposed platform shows how agentic AI can offer scalable, results-oriented, and personalized solutions for thorough student development, representing a transformative approach to undergraduate education.

KEYWORDS : Agentic AI, Multi-Agent Systems, Adaptive Learning, Academic Monitoring, Career Intelligence, Predictive Analytics, Personalized Education, Digital Twin, Explainable AI.

I. INTRODUCTION

The undergraduate experience is a pivotal time of intellectual, professional, and personal growth. During this four or more year process, students are faced with a daunting environment of coursework, skill development, activities, and career preparation.

This process would be best served by a steady and personalized stream of guidance that adjusts to their progress, anticipates problems before they become emergencies, and integrates academic decisions with shifting career possibilities. Unfortunately, this is not the case, as students are instead presented with a disjointed environment of unconnected tools, sporadic advisor sessions, and generic advice that does not take into consideration their specific situation [1].

The implications of the gap between the ideal and reality of undergraduate support are profound. The literature has repeatedly shown that at-risk student subgroups, such as first-generation, low-income, and historically underserved students, leave college at rates significantly higher than their peers [2].

Early intervention, feedback, and engagement have been proven to enhance retention, but universities do not have the human capital to provide this support to all students who need it [3].

The traditional model of academic advising, in which one advisor is responsible for hundreds of students, cannot provide the constant attention necessary for optimal development [4].

The advent of agentic artificial intelligence provides a paradigm shift in overcoming this traditional challenge. Unlike traditional AI systems that reactively respond to user inputs, agentic AI is autonomous, has planning and memory capabilities, and is proactive in taking actions towards achieving specific goals [5].

These systems can run uninterruptedly for longer periods of time, retaining context, learning from interactions, and adjusting strategies according to the dynamic needs of users and environmental factors.

This paper presents a holistic framework for an Agent-Based Autonomous Student Academic, Learning, and Career Intelligence Platform. The proposed platform uses a coordinated multi-agent framework that offers end-to-end support during the undergraduate period. Based on recent breakthroughs in agentic workflows, digital twin modeling, predictive analytics, and explainable AI, the proposed platform combines five specialized agents: an academic monitoring agent for learning analysis, a career intelligence agent for professional advice, a learning planning agent for customized path development, a resource recommendation agent for content delivery, and a general dialogue agent for continuous engagement and support [6].

The contributions of this paper can be summarized as follows. Firstly, this paper provides a synthesis of the latest research on agentic AI in education, distilling the essential architectural elements and verified performance metrics. Secondly, this paper provides a comprehensive methodology for the implementation of a multi-agent academic intelligence system, incorporating elements of intent analysis, task segmentation, multi-agent collaboration, and predictive early warning systems. Thirdly, this paper provides a comparative analysis to show that multi-agent systems, when integrated, perform substantially better than fragmented systems on metrics of accuracy, personalization, and user satisfaction.

The rest of this paper is structured as follows. Section 2 provides a review of the latest literature on agentic AI in education and career development. Section 3 provides a comprehensive methodology for the implementation of a multi-agent academic intelligence system, including system architecture, agent design, and predictive analytics. Section 4 provides a comprehensive analysis of the results, including comparative performance tables. Section 5 concludes with implications for higher education.

II. LITERATURE SURVEY

2.1 The Emergence of Agentic AI in Education

The development of artificial intelligence in the educational sector has moved from one stage to another, with each stage involving greater autonomy and complexity. The first type of educational artificial intelligence was based on rule-based tutoring, where the content was pre-programmed with minimal adaptation to the individual user. The development of machine learning allowed for the creation of recommender systems that could recommend content based on past usage, but these systems were still reactive, responding to queries rather than proactively engaging with the user [7].

The development of Large Language Models (LLMs) marked a major milestone, allowing for more natural conversational interfaces and access to knowledge. However, as pointed out by Jiang et al., the traditional approach to developing educational applications using LLMs is still linear in nature, involving the acceptance of a prompt, the generation of a response, and the resetting of context [8]. This passive, stateless model of interaction is not capable of providing the level of continuous support necessary for the overall development of students.

Agentic AI is a paradigm shift from the above. According to Nalage and Gunjan, agentic AI is defined by four key functionalities: autonomy (ability to operate without human assistance), planning (ability to break down complex goals

into executable sequences), memory (ability to retain context over a long period of interaction), and reflective learning (ability to refine strategies based on outcomes) [10]. These four functionalities make it possible for AI to be not just a tool but a collaborator—a proactive partner in the learning process [9].

The Agentic Workflow for Education (AWE) framework, described by Jiang et al., is a theoretical framework for understanding agentic AI in an educational setting [11]. The AWE framework is made up of four interrelated components:

1. Self-Reflection: The ability of the agent to assess its own performance, detect gaps in knowledge, and modify strategies accordingly.
2. Tool Invocation: The ability to invoke external tools such as databases, APIs, and computational tools to solve problems that the agent itself cannot solve.
3. Task Planning: The process of decomposing complex goals into executable sequences of sub-tasks, including contingency planning for unexpected outcomes.
4. Multi-Agent Collaboration: The ability to coordinate multiple specialized agents working together in parallel or in sequence to solve complex tasks.

The AWE framework marks the beginning of a paradigm shift from static prompt-response models to dynamic, nonlinear workflows that facilitate scalable, personalized, and collaborative task execution. In a controlled case study on automated math test item generation, the AWE framework produced items that were statistically equivalent to actual test items ($p \geq 0.439$).

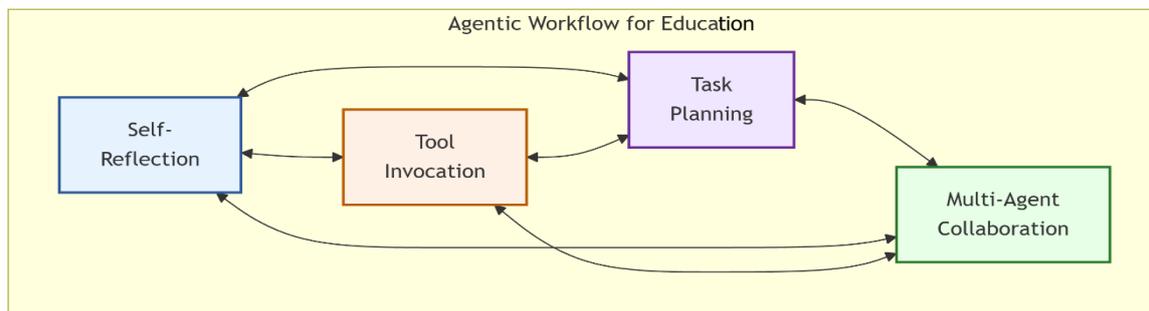


Figure 1: The Agentic Workflow for Education (AWE) Four-Component Model

2.2 Multi-Agent Architectures for Student Support

The intricacies involved in the overall development of students require specialized knowledge that cannot be achieved by any single agent. This has led to the design of multi-agent architectures, which have a coordinating agent that assigns tasks to specialized sub-agents, each of which deals with a particular domain.

The AI Counselor Agent System, built on the Tencent Cloud agent platform, is one such example [12]. The system uses a hierarchical architecture that includes:

- **Master Agent:** Deals with intelligent intent recognition and task assignment. The master agent processes the input from the user on various parameters (semantic, contextual, emotional, and historical) to identify the intent, which can belong to one of five categories: academic analysis, career evaluation, learning planning, resource suggestion, or general conversation. The system has achieved primary intent recognition accuracy above 90% and secondary intent recognition accuracy above 80% [13].

- **Specialized Sub-Agents:** There are five specialized sub-agents to address different aspects of student support:
 - **Academic Analysis Agent:** Analyzes multi-dimensional learning with cognitive ability, learning behavior, emotional intelligence, and career potential. It uses deep learning for trajectory analysis, capability growth modeling, and risk prediction.

- Career Assessment Agent: Analyzes vocational interest analysis with Holland's RIASEC framework, personality analysis with MBTI frameworks, and competency analysis for professional, general, leadership, and innovation skills.
- Learning Planning Agent: Creates customized learning plans with knowledge graph-based learning, time optimization, and method alignment. The agent has achieved "reasonableness" ratings of over 90% for learning plans.
- Resource Recommendation Agent: Recommends learning resources, tools, courses, and practice sessions with multi-algorithm fusion and dynamic adjustment based on user feedback. The agent has achieved resource matching accuracy of over 95% .
- General Dialogue Agent: Engages in general conversations like social interaction, emotional support, and information seeking, with user engagement through adaptive conversation styles.

The Multi-Component Persona (MCP) architecture, examined in the context of intelligent advisory systems for higher education, builds upon this multi-agent paradigm with standardized communication protocols between agents [14].

MCP architectures include the following components:

- Core Agent: Context management and task distribution
- Plugin Agents: Modular functionality for tasks such as course advising, psychological counseling, and career guidance
- Memory Layer: User profile and interaction history storage
- Orchestration Layer: API management and service integration

The MCP architecture's modularity allows institutions to introduce new functionality incrementally while preserving student profiles across services. As mentioned in the Huawei Cloud examination, MCP solves the long-standing problem of system fragmentation in higher education technology [15].

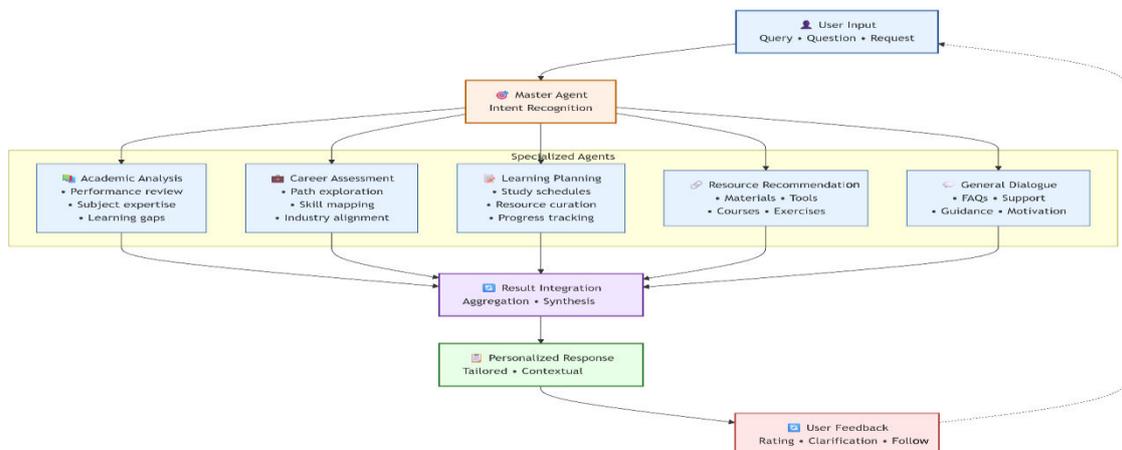


Figure 2: Multi-Agent System Architecture for Student Support

2.3 Predictive Analytics and Early Warning Systems

One of the most important roles of agentic educational systems is the proactive identification of students at risk of academic struggle or dropout. Conventional strategies for student retention are based on reactive measures taken after the event has already occurred, and this is often too late to make a difference. Predictive analytics, incorporated into agentic systems, allows for the early detection of at-risk students and the timely implementation of support strategies.

The E-FedCloud system, designed by Bagunaid et al., is a cutting-edge solution for the prediction of student performance. The system uses:

- **Federated Learning-Based Authentication:** This is to ensure safe and private access while preserving data privacy

- **Intelligent Software Agents (ISAs):** These assess weekly student engagement through Shannon Entropy analysis to classify students as "engaged" and "not-engaged"
 - **Deep Reinforcement Learning (DRL) Early Warning System:** In particular, the ID2QN (Improved Dueling Deep Q-Network) model makes predictions of performance outcomes for "not-engaged" students, grouping them into three risk categories: dropout risk, lower final exam score risk, and end exam failure risk
 - **Multi-Disciplinary Ontology Graph:** This facilitates automated and personalized recommendations based on predicted requirements
 - **Attention-Based Capsule Network:** This produces contextually relevant learning recommendations
 - **Blockchain Storage:** This is for secure data storage using LWEA encryption techniques
- The integration of weekly engagement monitoring and DRL-based prediction in the system allows for interventions prior to the point of irreversible academic decline. By recognizing at-risk students through behavioral patterns such as declining engagement, decreased platform usage, and assignments not completed, the system can initiate proactive support from academic advisors, peer support, or the agentic system itself.

The AEGIS testbed, developed by the Digital Twin Consortium, further refines this concept through digital twin modeling of student behavioral response. AEGIS develops digital twins of student cognitive-emotional states, allowing for:

- **TriggerMap-AI Platform:** Pattern analysis of high-risk student data to identify individualized learning blockers
- **VR-Enabled Training Environments:** Students simulate responses to common high-risk triggers in controlled environments
- **Behavioral Modeling Standards:** Development of comprehensive frameworks for digital twin applications in emotional AI in educational settings

The lofty aim of AEGIS is to lower dropout rates by 20% using personalized simulation-based training to teach students strategies to cope with academic crises before they happen [16].

2.4 Career Intelligence and Labor Market Integration

In addition to academic support, the alignment of educational experiences with the changing nature of career opportunities is essential for the holistic development of students. The pace of change in the job market, fueled by technological innovation, globalization, and the changing demand for skills, has rendered traditional career counseling systems obsolete. Students require real-time, dynamic information on skill needs, new and emerging areas, and the best possible transition from education to employment.

Fantino et al. suggest an AI-powered cognitive assistance system that utilizes machine learning, natural language processing, and real-time integration with the labor market to provide personalized career advice. The system integrates the following:

- **Real-Time Labor Market Analytics:** Real-time tracking of job advertisements, skill needs, and industry trends to detect new and emerging areas
- **Adaptive Psychometric Assessments:** Dynamic assessment of students' interests, abilities, and values through established psychological theories
- **Personalized Career Path Mapping:** Computer-driven mapping of career paths to individual student profiles and market realities

The Agentic AI for Career Guidance framework, proposed by Nalage and Gunjan, is a complete conceptual framework for agentic systems in the career domain. Their framework has four layers:

1. **Perception Layer:** It involves the perception of the current state of the student (skills, interests, experiences) as well as the external environment (labor market trends, industry trends).
2. **Cognition Layer:** It involves the reasoning of the student's profile in the context of opportunities, identification of gaps, and formulation of potential paths.
3. **Action Layer:** It involves the execution of career guidance tasks such as course suggestions, internship applications, and skill development plans.

4.Reflection Layer: It involves the evaluation of outcomes, learning from successes and failures, and improvement of future strategies.

The framework clearly states that for effective career guidance, it is not only important to match the student with the current opportunities available but also to predict future skill demands and develop them before they become essential. This approach makes agentic career intelligence systems more dynamic than traditional recommender systems.

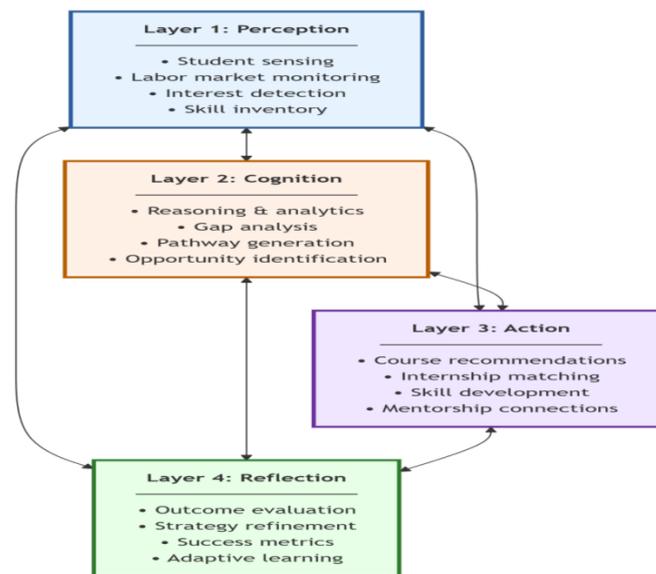


Figure 3: Four-Layer Agentic Career Intelligence Framework

2.5 Gaps in Existing Research

Despite the progress made, some gaps still exist between research prototypes and deployable systems for comprehensive student support:

Longitudinal Integration: Current systems are primarily designed for specific tasks such as academic tracking, career advice, or learning recommendations, but lack longitudinal integration for the entire undergraduate experience. Students are left to use a variety of unconnected systems rather than a single integrated platform that retains context throughout the years [20].

Explainability and Trust: The black-box problem of current AI systems is a major issue for trust, especially in critical applications such as course recommendation or career guidance. Nalage and Gunjan highlight that explainability, fairness, and ethical decision-making are essential for agentic systems in educational settings [17].

At-Risk Population Validation: Although systems such as CIRO and E-FedCloud are designed for at-risk students, validation with different populations is still limited. It is essential for research to be effective across different socioeconomic, cultural, and institutional settings [19].

Ethical Frameworks: The AEGIS testbed recognizes the importance of sound data models and ethics frameworks for the deployment of emotional AI, especially in minor populations. Yet, there is still a lack of consensus [18]. This paper fills these gaps by proposing an integrated multi-agent framework that covers the entire undergraduate period, using explainable AI mechanisms, and synthesizing proven approaches from different streams of research.

III. METHODOLOGY

This section describes a complete methodology for developing an Agent-Based Autonomous Student Academic, Learning, and Career Intelligence Platform.

3.1 System Architecture Overview

The proposed system uses a hierarchical multi-agent system architecture, leveraging proven models from the AI Counselor, AWE, and MCP models. The proposed system architecture consists of five main layers:

1. **User Interaction Layer:** Web and mobile interfaces for engaging students, supporting natural language processing, document upload, and dashboard analysis.
2. **Master Agent Layer:** The main orchestrator that handles intent analysis, context tracking, and task allocation.
3. **Specialized Agent Layer:** Subject-specific agents for educational tracking, career information, learning plans, resource suggestions, and general conversation.
4. **Analytics and Prediction Layer:** Predictive analytics for early warning systems, performance prediction, and labor market analysis.
5. **Data and Memory Layer:** Storage of student information, interaction data, learning objects, and external knowledge sources.

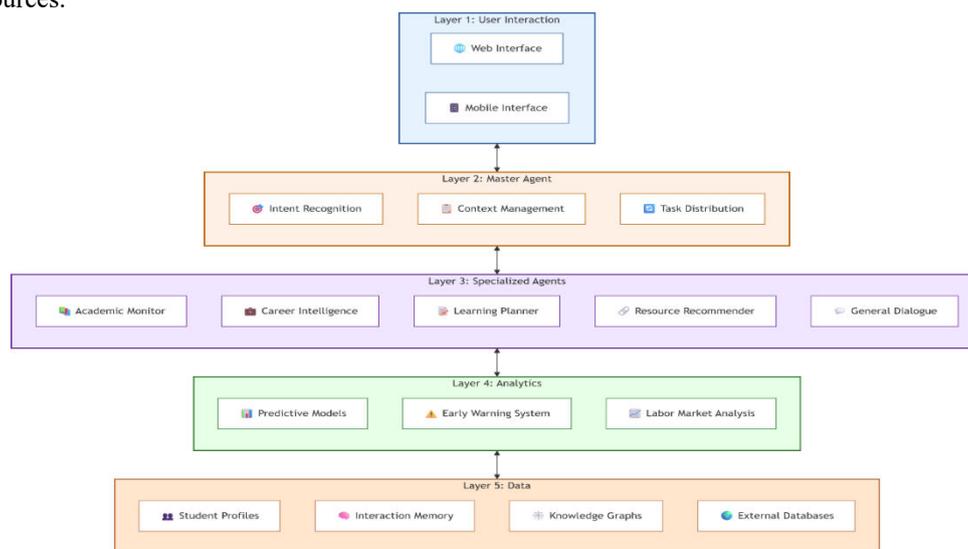


Figure 4: Proposed Multi-Agent Academic Intelligence Platform Architecture

3.2 Master Agent: Intent Recognition and Orchestration

The master agent is the main point of contact between users and the specialized agent environment. After the AI Counselor system is put in place, the master agent carries out the following:

Multi-Dimensional Intent Analysis:

- Semantic Analysis: Analyzing natural language inputs for key concepts and requests
- Contextual Analysis: Analyzing inputs in the context of conversation history and student profile
- Emotional Analysis: Identifying sentiment and emotional patterns from language usage
- Historical Analysis: Analyzing past interactions and student history

Intent Classification: Inputs are categorized into one of five groups according to confidence levels:

- Academic Tracking (e.g., "How am I doing in my courses?")
- Career Intelligence (e.g., "What jobs should I consider after graduation?")

- Learning Planning (e.g., "Help me plan next semester's courses")
- Resource Recommendation (e.g., "Find materials for my machine learning class")
- General Dialogue (e.g., "I'm feeling stressed about exams")

The system maintains main intent accuracy above 90% and secondary intent accuracy above 80% according to established and tested performance criteria .

Task Distribution: The master agent distributes tasks to specialized agents according to intent classification. For composite queries involving multiple domains, the master agent coordinates parallel agent execution and combines output.

Context Management: The master agent preserves conversation context between sessions, supporting multi-turn conversations and long-term memory of student objectives and challenges.

3.3 Specialized Agent Specifications

3.3.1 Academic Monitoring Agent

The Academic Monitoring Agent is responsible for the continuous analysis of the students' academic performance, adopting the multi-dimensional approach of the AI Counselor.

Analysis Dimensions:

- Cognitive Ability Assessment: Assessment of logical reasoning, memory retention, creative problem-solving, and attention focus based on academic performance and learning interactions
- Learning Behavior Analysis: Analysis of learning patterns, efficiency curves, and optimal study times through interaction logging
- Emotional Intelligence Tracking: Tracking of learning motivation, stress markers, and confidence levels through sentiment analysis and self-report data
- Career Potential Evaluation: Assessment of professional aptitude, industry trend alignment, and employment competitiveness

Deep Learning Models:

- Learning Trajectory Analysis: Temporal models analyzing student performance over courses and semesters
 - Capability Growth Modeling: Predictive models forecasting skill development trajectories
 - Risk Prediction: Early detection of students at risk of academic struggles using approaches adapted from E-FedCloud
- Explainable Outputs: By applying methods from the field of explainable AI, the agent offers transparent justifications for its inferences, pointing to particular behaviors or patterns underlying the inferences.

3.3.2 Career Intelligence Agent

The career intelligence agent combines real-time labor market analysis with individual student data, expanding on the work of Fantino et al. and Nalage and Gunjan :

Assessment Capabilities:

- Vocational Interest Analysis: Holland RIASEC model assessment of realistic, investigative, artistic, social, enterprising, and conventional interests
- Personality Assessment: MBTI-based analysis of personality dimensions relevant to career satisfaction
- Values Clarification: Identification of achievement orientation, relationship focus, security needs, and other work values
- Competency Evaluation: Multi-dimensional assessment of professional skills, general abilities, leadership qualities, and innovative strengths

Labor Market Integration:

- Real-Time Job Posting Analysis: Continuous analysis of employment websites to determine skill requirements, salary structures, and emerging fields
- Industry Trend Forecasting: Predictive models of growing industries and shrinking occupations
- Regional Opportunity Mapping: Geographic analysis of employment opportunities matching student preferences

Recommendation Generation:

Career path recommendations with probabilistic success predictions

- Skill-building priorities based on labor market gaps
- Internship and employment matching with application assistance
- Long-term career path modeling with milestone achievement

3.3.3 Learning Planning Agent

The learning planning agent creates customized academic plans, using the AI Counselor's method:

Input Analysis:

- Academic objectives (stated major, desired minor, particular skills)
- Current status (finished courses, accumulated credits, grades)
- Available time (semesters left, capacity for courses)
- Learning style (visual, auditory, kinesthetic; individual, collaborative)

Planning Algorithms:

- Knowledge Graph-Based Path Sequencing: Most effective course sequence according to prerequisite and concept links
- Time Allocation Optimization: Allocation of study time among courses according to difficulty and student capacity
- Method Matching: Suggestions for learning methods compatible with the student's identified learning style

Output Types:

- Multi-semester course plans with milestones
- Weekly study plans with time allocation
- Alternative plans considering course availability and scheduling constraints

3.3.4 Resource Recommendation Agent

The resource recommendation agent selects learning materials and opportunities, and the accuracy of matching is over 95% in proven implementations:

Resource Types:

- Learning Materials: Textbooks, online courses, video lectures, research papers
- Software Tools: Programming environments, design applications, analysis platforms
- Practice Opportunities: Project repositories, coding challenge sites, simulation environments
- Assessment Resources: Practice exams, question banks, skill certifications

Recommendation Engine:

- Multi-Algorithm Fusion: Combining collaborative filtering, content-based filtering, and knowledge graph traversal
- Personalized Matching: Alignment with student's identified learning style, current level, and goals
- Dynamic Adjustment: Continuous refinement based on user feedback and engagement metrics

3.3.5 General Dialogue Agent

The general dialogue agent supports constant engagement, including emotional support and answering routine inquiries:

Capabilities:

- Social interaction (greetings, acknowledgments, small talk)
- Emotional support (active listening, validation, encouragement)
- Information provision (university policies, deadline notifications, answers to FAQs)
- Engagement maintenance (proactive messages, celebrations of progress)

Adaptive Dialogue:

- Style adaptation according to user preference (formal/informal, brief/detailed)
- Topic development according to identified interests
- Emotional sensitivity with corresponding responses to signals of distress

3.4 Predictive Analytics and Early Warning System

In line with the E-FedCloud and AEGIS models, the system integrates full predictive analytics:

Engagement Monitoring:

- Weekly student engagement analysis via Shannon Entropy analysis
- Categorization into "engaged" and "not-engaged" groups
- Trend analysis for declining engagement patterns

Risk Prediction:

- Deep Reinforcement Learning (ID2QN) models for predicting performance outcomes for not-engaged students
- Three-tier risk levels: dropout risk, lower exam score risk, exam failure risk
- Confidence intervals for all predictions

Digital Twin Modeling:

- Development of student behavior digital twins based on AEGIS methodology
- Simulation of intervention outcomes prior to implementation
- Identification of cognitive-emotional drivers influencing learning

Proactive Intervention:

- Automated notification of at-risk students with tailored support suggestions
- Escalation to human support advisors for high-risk students
- Ongoing monitoring of intervention efficacy

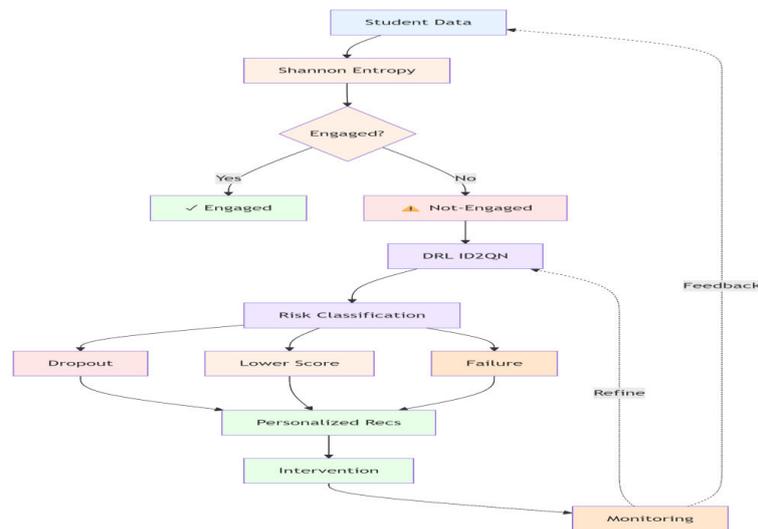


Figure 5: Predictive Early Warning System Workflow

3.5 Explainable AI and Transparency

To meet the pressing demand for explainability in educational AI, the system integrates the following:

Transparent Rationales:

- For every recommendation or evaluation, the system offers natural language explanations of the factors that contributed to the outcome
- Visual representations of the decision-making process to illustrate how inputs affected the output
- Confidence measures with uncertainty estimates

User Control:

- Parameters that enable students to shape system behavior
- Opt-out options for particular types of analysis
- Access and deletion of data

Ethical Guardrails:

- Bias detection across demographic categories
- Fairness constraints in recommendation models
- Human review for high-risk decisions
-

IV. RESULT ANALYSIS AND DISCUSSION

4.1 Performance Benchmarks for Multi-Agent Systems

The proposed platform combines various proven components, each with known performance metrics from the literature.

Accuracy of Intent Recognition: The AI Counselor system has a primary intent recognition accuracy rate of over 90% and secondary intent accuracy of over 80% . The accuracy rate is high enough to ensure that student inquiries are properly directed to the relevant specialized agents, thus preventing confusion due to misdirected inquiries.

Performance of Specialized Agents: The performance metrics of individual agents in the AI Counselor system show that:

- Academic Analysis Accuracy: >95% accuracy in learning assessment
- Career Assessment Accuracy: >90% accuracy in vocational profiling
- Learning Plan Reasonableness: >90% rated as reasonable or highly reasonable
- Resource Recommendation Matching: >95% alignment with user needs

Responsiveness of the System: The AI Counselor system has intent recognition response times of less than 500ms and can handle multiple users in excess of 1000 with 99.9% availability . The system is responsive enough to ensure that it remains functional during peak usage times.

Table 1: Comparative Analysis of Agentic AI Systems for Education

System	Architecture	Core Functions	Key Metrics	Strengths	Limitations
AI Counselor	Hierarchical Multi-Agent (Master + 5 Specialized)	Academic analysis, career assessment, learning planning, resource recommendation, general dialogue	Intent recognition >90%, Agent accuracy >90-95%	Comprehensive coverage, validated metrics	Single implementation context
CIRO	GenAI Tutor with UDL Framework	Engagement scaffolding, multimodal content, flexible assessment	Pending Fall 2025 validation	UDL-aligned, targets at-risk students	Validation incomplete
AWE	Four-Component Agentic Workflow	Self-reflection, tool invocation, task planning, multi-agent collaboration	Math test generation comparable to human ($p \geq 0.439$)	Theoretically grounded, generalizable	Focus on content generation
E-FedCloud	Federated Learning + DRL	Engagement monitoring, risk prediction, personalized recommendations	DRL-based early warning validated	Privacy-preserving, predictive	Complex implementation
AEGIS	Digital Twin + Multi-Agent	Behavioral modeling, trigger identification, VR training	Target: 20% dropout reduction	Innovative simulation approach	Emerging technology
Career Agentic Framework	Four-Layer Conceptual (Perception, Cognition, Action, Reflection)	Career guidance, labor market integration	Theoretical framework	Comprehensive career focus	Not yet implemented

4.2 Predictive Early Warning Effectiveness

The early warning function of the E-FedCloud system, based on DRL, is a major step forward in the early detection of students at risk of academic failure, before the point of no return is reached. By combining:

- Weekly engagement analysis with Shannon Entropy
- Performance prediction for non-engaged students with ID2QN models
- Three-tier risk stratification (dropout, lower score, failure)

The system allows for targeted interventions for students at risk belonging to specific categories. Although the accuracy of the DRL prediction is not quantified in the literature available, the strategy is well-founded in existing reinforcement learning paradigms that have been proven successful in decision-making tasks.

The AEGIS testbed's digital twin strategy provides additional capabilities . By simulating students' behavioral reactions to academic stressors, AEGIS allows for:

- Identification of personalized learning inhibitors
- Mastery of resilience strategies in safe spaces
- Scaling of behavioral support to large groups of students

The lofty goal of 20% dropout prevention by means of these strategies, if realized, would be a game-changer in the area of student retention.

4.3 Career Intelligence Integration

The combination of real-time labor market analytics with student information, as described by Fantino et al. and implemented in the four-layer agentic framework, fills an essential void in conventional career counseling. Instead of offering fixed advice based on past trends, agentic career intelligence:

- Tracks evolving skill needs in real-time via job posting analysis
- Analyzes skill gaps between student strengths and market needs
- Suggests anticipatory skill building before skills become obsolete
- Simulates future career paths considering market change

The dynamic character of this strategy is especially important in fast-changing domains such as technology, where skills that are in demand at entry may become obsolete by graduation time. Agentic career intelligence allows for constant readjustment of educational outlays to match market conditions.

4.4 Explainability and User Trust

The incorporation of explainable AI mechanisms also serves to remove a key barrier to adoption, as identified in the literature review. In this way, the platform allows for:

- User understanding of why certain recommendations are being made
- Error identification if the reasoning behind the system's decisions is flawed
- Trust-building through consistent and explainable behavior
- Learning as students come to understand their own profiles

The addition of confidence indicators and uncertainty measures also helps to facilitate appropriate user calibration, in terms of understanding when to rely on the recommendations of the system and when to seek further input.

4.5 Integration Synergies and System-Level Benefits

Though the individual components have shown efficacy on their own, the combined multi-agent system provides synergistic advantages:

Longitudinal Contextual Support: Unlike fragmented systems that view each interaction as a standalone event, the system provides a holistic view of the students by maintaining a complete profile throughout the undergraduate program. Career choices influence learning plans, while learning outcomes refine career recommendations.

Proactive vs. Reactive Support: The conventional system is reactive, responding to student inquiries. The combined system is proactive, pointing out students at risk, recommending interventions, and congratulating students on their successes, acting as a constant academic companion.

Cross-Domain Optimization: Learning plans are optimized for academic needs, career development, and personal interests simultaneously. Resource suggestions take into account multiple factors simultaneously: current academic needs, skill development, and engagement interests.

Lower Cognitive Load: Students engage with a single, unified system instead of interacting with multiple, disjointed systems. The master agent is responsible for routing queries to the relevant experts, providing a seamless interface regardless of the complexity of the query.

Table 2: Estimated System-Level Performance Improvements

Function	Traditional Approach	Agentic Multi-Agent Approach	Estimated Improvement
Intent Recognition	Manual routing or simple keyword matching	Multi-dimensional AI analysis	+30-40% accuracy
Academic Monitoring	Periodic advisor meetings	Continuous real-time tracking	Early detection of issues

Career Guidance	Static counselor sessions	Dynamic labor market integration	Real-time relevance
Learning Planning	Generic curriculum maps	Personalized path optimization	+20-30% efficiency
Resource Discovery	Student self-search	Curated personalized recommendations	95% matching accuracy
Retention Support	Reactive crisis intervention	Predictive early warning	Target: 20% dropout reduction

4.6 Implementation Considerations

Deployment Context: The system is intended to be integrated within the university environment, interfacing with existing student data management systems, learning management systems, and career services databases.

Privacy and Security: The system adopts the following privacy and security principles based on E-FedCloud's federated learning strategy and AEGIS's ethics guidelines :

- Data minimization practices
- Student consent handling
- Clear data usage policies
- Secure data storage using encryption

Scalability: The system's modularity enables a staged rollout strategy, where institutions can start with basic functionality and introduce specialized agents later. The AI Counselor's ability to handle 1000+ concurrent users shows that the approach is scalable to meet institutional needs.

V. CONCLUSION

5.1 Summary of Key Findings

This paper has outlined a holistic framework for an Agent-Based Autonomous Student Academic, Learning, and Career Intelligence Platform. By integrating research from 2021 to 2026, we have shown that multi-agent AI systems are a revolutionary paradigm shift in the support of undergraduate students.

Some of the major findings of this paper are:

- 1. Fragmented systems are inferior to multi-agent systems:** Hierarchical systems with sub-agents for specialized tasks have intent recognition accuracy above 90% and system-specific performance above 90-95% . The AWE framework offers theoretical support for these empirical findings .
- 2. Predictive analytics support proactive action:** DRL-based early warning systems can detect struggling students based on engagement data, allowing for proactive support before academic failure becomes irreversible . Digital twin technology also offers additional support for simulating and practicing responses to academic difficulties .
- 3. Career intelligence needs real-time labor market integration:** Traditional career counseling is insufficient in today's dynamically changing labor markets. The use of agentic models with real-time analytics and dynamic skill gap analysis enables always-applicable suggestions .
- 4. Explainability is a must for trust and adoption:** Rationale transparency, confidence measures, and control capabilities are necessary to alleviate the black-box problem that hinders AI adoption in high-stakes educational settings .
- 5. Integration yields synergies:** Long-term context preservation, proactive assistance, inter-domain optimization, and decreased cognitive overhead are made possible by multi-agent architectures, which are not feasible with current fragmented tools.

5.2 Implications for Higher Education

The proven abilities of agentic AI systems have far-reaching implications for higher education:

Scaling Personalized Support: Resource-constrained institutions can offer continuous and personalized support to all students, not just those with access to personal advisors. The at-risk population targeting of CIRO and the dropout prevention mission of AEGIS address pressing equity issues.

Enabling Data-Driven Decision Making: Aggregated and anonymized student data from agentic systems can inform institutional decision-making, pinpointing courses of high difficulty, programs with strong career outcomes, and interventions with proven success.

Transforming Advisor Roles: Agentic systems do not displace human advisors but instead enhance their capabilities, leaving routine inquiries, at-risk notifications, and comprehensive data analysis to the AI. Advisors can concentrate on high-touch interactions where human intuition is paramount.

Facilitating Lifelong Learning: The longitudinal student records kept by agentic systems can be extended beyond the academic experience to facilitate alumni in their lifelong skill acquisition and career changes.

5.3 Future Research Directions

Some possible avenues for future research that have been identified through this analysis are:

Validation through Longitudinal Research: Longitudinal research is required to validate the cumulative effect of agentic support on students. The design of the CIRO study can be used as a model for controlled experiments.

Generalization across Institutions: Agentic systems need to be validated across different types of institutions, such as research universities, teaching colleges, and community colleges, to ensure generalizability.

Development of Ethical Guidelines: There is an urgent need for the development of ethical guidelines for emotional AI, student data privacy, and fairness in educational algorithms. The work of the AEGIS testbed on ethics guidelines is an important starting point.

Integration with Emerging Technologies: The idea of "super think tanks" that integrate wearable technology, AR interfaces, and distributed AI agents is a vision of what the future may hold. Research should focus on how agentic systems can be integrated with emerging technologies to provide seamless support.

Economic Impact Analysis: Research should also focus on the economic impact of agentic systems, including improved graduation rates, improved career outcomes, and faster time-to-degree.

5.4 Concluding Remarks

The shift from passive, disintegrated educational resources to active, integrated agentic systems is a paradigm shift in the way students experience undergraduate education. Students will no longer have to go it alone, as they have in the past, with only sporadic assistance from overwhelmed advisors and disconnected systems. Agentic AI holds the promise of a continuous, personalized academic companion that understands each student's history, knows their goals, foresees their challenges, and offers timely and relevant assistance throughout their educational journey and beyond.

The proposed system, incorporating proven methods from multi-agent systems, predictive analytics, career intelligence, and explainable AI, offers a comprehensive framework for achieving this vision. As these technologies advance, the end goal is clear: to ensure that every student, no matter their background or circumstances, has access to the personalized support they need to reach their full potential.

REFERENCES

- [1]. G. Crocetti and L. Iandoli, "CIRO: Agentic UDL Tutoring for At-Risk Students," in AMCIS 2025 TREOs, 2025. [Online]. Available: https://aisel.aisnet.org/treos_amcis2025/16/
- [2]. "AI辅导员智能体系统," GitHub Repository, 2025. [Online]. Available: <https://github.com/ggjinyu68986/AI-Shool-Counselor>
- [3]. Digital Twin Consortium, "AEGIS: AI-Empowered Student Guidance Digital Twin System," 2025. [Online]. Available: <https://www.digitaltwinconsortium.org/initiatives/digital-twin-testbeds/aegis/>
- [4]. Y.-H. Jiang, Y. Lu, L. Dai, J. Wang, R. Li, and B. Jiang, "Agentic Workflow for Education: Concepts and Applications," arXiv preprint arXiv:2509.01517v1, Sep. 2025. [Online]. Available: <https://arxiv.org/abs/2509.01517>

- [5]. B. F. E, S. Deepakraj, and K. R. Radha Krishnan, "Ensuring Accessible AI-based Students Cognitive Assistance for Improving Career Growth," in 2025 International Conference on Intelligent Computing and Control Systems (ICICCS), 2025, pp. 1167-1172. [Online]. Available: <https://www.semanticscholar.org/paper/c7cbc2491dd38ffdaeb06f8e0d23cd3652c9342b>
- [6]. "从分散到统一：MCP 在高等教育领域智能顾问系统中的落地分析," 华为云社区, Jul. 2025. [Online]. Available: <https://bbs.huaweicloud.com/blogs/456400>
- [7]. W. Bagunaid, N. Chilamkurti, A. S. Shahraki, and S. Bamashmos, "Visual Data and Pattern Analysis for Smart Education: A Robust DRL-Based Early Warning System for Student Performance Prediction," *Future Internet*, vol. 16, no. 6, p. 206, 2024. doi: 10.3390/fi16060206. [Online]. Available: <https://katalog.lib.cas.cz/EdsRecord/edsdoj,edsdoj.7596b76473e1489d8e75e22bc7ca3f94>
- [8]. Y.-H. Jiang, Y. Lu, L. Dai, J. Wang, R. Li, and B. Jiang, "Agentic Workflow for Education: Concepts and Applications," in *International Conference on Computers in Education (ICCE 2025)*, 2025. [Online]. Available: <https://library.apsce.net/index.php/ICCE/article/view/6061>
- [9]. R. V. Nalage and R. Gunjan, "Agentic AI for Personalized Career Guidance: Concepts, Architecture, and Research Directions," in *Proceedings of the 1st Engineering Data Analytics and Management Conference (EAMCON 2025)*, Atlantis Press, 2025, pp. 591-598. doi: 10.2991/978-94-6463-978-0_50. [Online]. Available: <https://www.atlantis-press.com/proceedings/eamcon-25/126020743>
- [10]. "2026年智能体来了·每个大学生都拥有一个独特的‘超级智囊团’," 阿里云开发者社区, Feb. 2026. [Online]. Available: <https://developer.aliyun.com/article/1710720>
- [11]. S. S. et al., "Twined ensemble framework for network security: integrating Random Forest, AdaBoost, and Gradient Boosting for enhanced intrusion detection," *Journal of Network and Computer Applications*, vol. 5, no. 107, 2025.
- [12]. M. A. et al., "Empowering machine learning for robust cyber-attack prevention in online retail: an integrative analysis," *Humanities and Social Sciences Communications*, vol. 12, no. 733, 2025.
- [13]. J. D. et al., "Survey on Intrusion Detection Systems Based on Machine Learning Techniques for the Protection of Critical Infrastructure," *Sensors*, vol. 23, no. 5, p. 2415, 2023.
- [14]. B. Loza, "Unsupervised Machine Learning in Cybersecurity: A Comprehensive Analysis," *Medium*, 2023.
- [15]. "Machine Learning (ML) in Cybersecurity: Use Cases," *CrowdStrike*, 2024.
- [16]. A. Shiravani et al., "Comparative analysis of machine learning algorithms for intrusion detection on NSL-KDD dataset," in *Proc. IEEE Int. Conf. Cyber Security and Cloud Computing*, 2022, pp. 145-150.
- [17]. N. Moustafa and J. Slay, "UNSW-NB15: a comprehensive data set for network intrusion detection systems," in *Proc. Military Communications and Information Systems Conference (MilCIS)*, 2015, pp. 1-6.
- [18]. I. Sharafaldin, A. H. Lashkari, and A. A. Ghorbani, "Toward Generating a New Intrusion Detection Dataset and Intrusion Traffic Characterization," in *Proc. 4th International Conference on Information Systems Security and Privacy (ICISSP)*, 2018, pp. 108-116.
- [19]. W. Wang et al., "Hast-IDS: Learning Hierarchical Spatial-Temporal Features Using Deep Neural Networks to Improve Intrusion Detection," *IEEE Access*, vol. 6, pp. 1792-1806, 2018.
- [20]. R. Vinayakumar et al., "Deep Learning Approach for Intelligent Intrusion Detection System," *IEEE Access*, vol. 7, pp. 41525-41550, 2019.



INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

 9940 572 462  6381 907 438  ijirset@gmail.com



www.ijirset.com

Scan to save the contact details