

**CAREERNET: A HYBRID DEEP LEARNING AND REINFORCEMENT LEARNING
MODEL FOR STUDENT CAREER RECOMMENDATIONS**

¹Rantheep Raja S, ²Dr. Priya M

¹Research Scholar

Department of Computer and Information Science,

Faculty of Science, Annamalai University.

Email: rantheep@gmail.com

²Assistant Professor

Department of Computer Science,

PS PT MGR Government Arts and Science College,

Sirkali, Puthur.

Email: priyacseau@gmail.com

Abstract

The increasing complexity of career decision-making among undergraduate students highlights the need for intelligent, personalized, and sustainable guidance systems. This paper introduces CareerNet, a hybrid framework that integrates deep learning with reinforcement learning to provide tailored career recommendations while accounting for both skill competency and mental well-being. Leveraging the OCEAN personality model and aptitude assessments (numerical, spatial, perceptual, abstract, and verbal) from a publicly available Kaggle dataset comprising 2,000 instances, the framework constructs comprehensive student profiles that capture cognitive strengths, behavioral traits, and mental health indicators. The proposed CARE-RL (Career and Resilience-Enhanced Reinforcement Learning) algorithm models career selection as a multi-objective optimization problem, where the reward function balances skill-career alignment with mental health sustainability. CareerNet employs deep feature embeddings for student profiling, graph-based career mapping, and adaptive decision-making through reinforcement learning, enabling dynamic and explainable recommendations. Experimental results demonstrate superior performance compared to traditional classifiers and collaborative filtering approaches, achieving improvements in accuracy, top-N recommendation precision, and interpretability. This research contributes a scalable, explainable, and student-centric AI framework that bridges the gap between academic skills, mental health, and career development, offering actionable insights for both students and educators.

Keywords: Career Recommendation System, Deep Learning, Reinforcement Learning, OCEAN Personality Traits, Aptitude Profiling, Mental Health-Aware AI, Student-Centric Career Guidance, Educational Data Mining.

1. INTRODUCTION

Career selection is one of the most critical decisions faced by undergraduate students, as it significantly influences professional growth, personal satisfaction, and long-term well-being. However, career guidance in many academic environments remains limited to general aptitude tests or one-size-fits-all counseling approaches, which fail to account for the diverse combination of personality traits, cognitive strengths, and mental health conditions that shape individual career readiness. In recent years, the adoption of artificial intelligence (AI) in education has opened new possibilities for providing data-driven, personalized, and adaptive recommendations that better align with students' skills, interests, and psychological profiles.

The OCEAN personality model (Big Five traits) and aptitude assessments across cognitive domains such as numerical, spatial, perceptual, abstract, and verbal reasoning provide rich insights into a student's strengths and weaknesses. When combined with mental health indicators, these features form a holistic student profile capable of informing not only academic interventions but also sustainable career recommendations. Traditional machine learning approaches for career guidance have largely relied on static classifiers or collaborative filtering, which often fail to adapt dynamically to the multi-dimensional nature of students' evolving profiles. Moreover, these systems rarely consider the importance of mental health sustainability in career decision-making, resulting in recommendations that may be academically aligned but psychologically unsustainable.

To address these challenges, this paper proposes CareerNet, a hybrid deep learning and reinforcement learning framework designed for personalized and sustainable career recommendations. The novelty of CareerNet lies in its integration of CARE-RL (Career and Resilience-Enhanced Reinforcement Learning), a multi-objective algorithm that balances skill alignment with mental health resilience, ensuring recommendations are both academically suitable and psychologically sustainable. Deep learning is employed to generate latent feature embeddings from personality and aptitude traits, while reinforcement learning enables adaptive decision-making through a dynamic reward system that considers career fit, stress tolerance, and adaptability.

The framework is evaluated using a real-world Kaggle dataset containing 2,000 student records with personality and aptitude attributes. Experimental results demonstrate that CareerNet significantly outperforms traditional classifiers and baseline recommenders in terms of accuracy, top-N recommendation precision, and interpretability. The proposed system also offers explainable outputs, allowing students and educators to understand the rationale behind each recommendation, thereby increasing trust and adoption in educational environments.

The key contributions of this research are as follows:

1. **Novel Hybrid Framework (CareerNet):** Development of a student-centric career recommendation framework that integrates deep learning feature embeddings with reinforcement learning for adaptive and explainable recommendations.

2. **CARE-RL Algorithm:** Introduction of a new reinforcement learning-based algorithm that formulates career guidance as a **multi-objective optimization problem**, balancing career-skill alignment with mental health sustainability.
3. **Holistic Student Profiling:** Utilization of the OCEAN personality model and aptitude assessments across five cognitive domains, combined with mental health indicators, to construct comprehensive student profiles.
4. **Experimental Validation:** Extensive evaluation on a real-world Kaggle dataset of 2,000 students, demonstrating improved accuracy, top-N recommendation precision, and interpretability over traditional machine learning and collaborative filtering approaches.

2. RELATED WORKS

A mobile-based competency training system was introduced to support mental health and reduce substance use among adolescents by recommending suitable lifelong learning courses. This study employed AI-driven recommendation systems (AI-RS) powered by ontological modeling and machine learning algorithms to analyze student learning preferences and psychological profiles, delivering targeted course suggestions to enhance both professional skillsets and emotional well-being [11]. A machine learning-based career recommendation system was proposed to guide undergraduate students toward optimal career paths using datasets related to IT roles. Feature selection techniques were applied to determine the most influential factors, and several algorithms were tested. Random Forest achieved the highest accuracy and was chosen as the final predictive model to recommend career choices based on skill and interest alignment [12].

An integrated Digital Twin (DT), machine learning (ML), and deep learning (DL) system was developed for activity recognition and career recommendation. Students' location, static, accelerometer, and academic data were combined to train CNN for activity monitoring and logistic regression for career predictions. The DT framework provided real-time updates and monitoring of student behavior and course performance [13]. Using the PRISMA framework, this systematic review analyzed educational recommender systems with a focus on user groups, ML techniques, and application domains. The study found that classifiers, clustering methods, and hybrid techniques dominated the landscape. However, gaps remained in inclusivity and workplace adaptability, prompting the need for context-aware systems for both students and professionals [14].

This study combined the traditional Chinese Four Pillars of Destiny with modern ML techniques to develop a culturally enriched career recommendation tool. Data on elemental compositions derived from birth details were processed using PCA and Random Forest algorithms. The system, hosted on a web platform, offered career recommendations by aligning personal metaphysical profiles with suitable occupational categories [15]. Naïve Bayes and K-Nearest Neighbors (KNN) classifiers were employed in a predictive model aimed at guiding students in making appropriate career decisions. The system leveraged academic performance, extracurricular activities, and domain-specific proficiency to align career paths with the

student's aptitude. Model accuracy and effectiveness were validated through rigorous performance metrics [16].

A hybrid recommendation framework named WHO-RXGBoost was proposed using Wild Horse Optimization to tune XGBoost for student career and entrepreneurial guidance. PCA was used for dimensionality reduction, and robust scaler normalization was applied. The model outperformed traditional baselines, achieving 95% accuracy and 93% precision on student preference and employment history datasets [17]. This study reviewed AI-driven recommender systems for job advice, emphasizing the role of data mining and machine learning to address gaps in effective career counseling. A comparative analysis was presented between traditional guidance and AI-based methods, highlighting the superiority of supervised models in managing educational, aptitude, and environmental data [18].

CareProfSys, an ontology-driven recommender system, was designed to support career counseling for high school students. The system utilized structured interviews, career value extraction, and inferencing based on semantic similarity between career concepts and personality traits. A case study involving Chemical Engineering illustrated the model's relevance and adaptability [19]. This methodology reviewed existing literature on job-role suggestion systems using machine learning and decision trees, with a focus on matching students' credentials to professional requirements. It highlighted the importance of integrating career recommender systems into recruitment workflows for aligning candidates' skill sets with suitable job opportunities [20].

A Machine Learning-Assisted Recommendation System (MARS) was created to support student participation in Generic Competency Development Activities (GCDAs). Using historical recommendation data from advisors, LightFM (a hybrid model combining collaborative and content filtering) was found to outperform other models. Its deployment improved both skill development outcomes and student satisfaction [21]. An intelligent employment management system for entrepreneurship and job placement was designed using personalized feature vectors. The model processed academic records and self-declared interests to provide customized career recommendations and information delivery. Performance evaluations confirmed the system's role in improving graduate employability through targeted guidance [22].

A career prediction model for high school students was developed using supervised machine learning, trained on test performance and psychometric data. The system recommended appropriate job fields by identifying alignment between student performance and career requirements, enabling flexible transitions and academic redirection based on predictive accuracy [23]. The Personalized Naïve Bayes (p-NB) algorithm was used in the Career Path Recommendation Model (CPRM), which integrated job profiles, personality types, and academic subjects. The system was validated through expert evaluations and a pilot study with Indonesian students, showing over 83% satisfaction and strong model-user alignment in IT job recommendations [24]. This study implemented a training course recommendation engine using Decision Tree, Random Forest, Gradient Boosting, and Backpropagation Neural Networks to align course suggestions with teachers' 21st-century skill requirements. Features

were grouped into skill categories during preprocessing to address data imbalance, and the best-performing model achieved a precision of 78% [25].

3. PROPOSED MODEL: CareerNet Framework

The proposed CareerNet framework integrates deep learning and reinforcement learning to deliver personalized and sustainable career recommendations for undergraduate students. The system begins with preprocessing the Kaggle career prediction dataset, where OCEAN personality traits and aptitude scores are normalized and combined with mental health indicators derived from the MindSkill framework to construct holistic student profiles. These profiles are transformed into latent feature embeddings using deep neural layers, capturing hidden relationships among cognitive, behavioral, and psychological attributes. A Career-Skill-Mental Graph (CSM-Graph) is then constructed, linking student features with career domains through weighted compatibility scores.

At the core of the framework lies the CARE-RL algorithm, which models career recommendation as a multi-objective optimization problem. Here, the state space represents student profiles, the action space corresponds to career pathways, and the reward function balances skill-career alignment, mental health sustainability, and adaptability. Using a Deep Q-Network (DQN)-based learning strategy, the model iteratively refines its policy to recommend optimal and evolving career options. Finally, CareerNet generates Top-N career recommendations with an explainable AI layer that provides interpretable justifications, ensuring recommendations are not only accurate but also psychologically sustainable and actionable for both students and educators. An overall architecture of proposed framework is shown in Fig 1.

Step 1: Data Acquisition and Preprocessing

The proposed CareerNet framework utilizes the Kaggle Career Prediction dataset, which contains 2,000 student instances described by 10 continuous attributes derived from psychometric and aptitude assessments. These features span two major dimensions: Personality Traits (OCEAN model) and Aptitude Scores across five cognitive domains. Additionally, mental health indicators (risk categories: low, medium, high) to enrich the dataset with psychological well-being information.




Figure 1: Overall Framework of Proposed CareerNet Model

Table 1. Dataset Features and Descriptions

Feature Name	Type	Range	Description
Openness	Continuous	0–1	Creativity, curiosity, openness to novelty
Conscientiousness	Continuous	0–1	Discipline, organization, goal-directed
Extraversion	Continuous	0–1	Sociability, energy, outgoing behavior
Agreeableness	Continuous	0–1	Cooperation, empathy, trust
Neuroticism	Continuous	0–1	Emotional instability, anxiety
Numerical	Continuous	0–100	Mathematical and logical ability
Spatial	Continuous	0–100	Visualization and spatial reasoning
Perceptual	Continuous	0–100	Pattern recognition, sensory sensitivity
Abstract	Continuous	0–100	Logical reasoning, abstract thinking
Verbal	Continuous	0–100	Language comprehension, inference skills

Table 1 summarizes the features of the Kaggle Career Prediction dataset, comprising personality traits based on the OCEAN model and aptitude scores across five cognitive domains. Personality attributes are normalized within [0–1], while aptitude scores range from 0–100 before preprocessing. Together, these variables form comprehensive student profiles for mental health-aware career recommendation.

Each student profile is represented as a 10-dimensional feature vector.

$$X_i = [O_i, C_i, E_i, A_i, N_i, Num_i, Spa_i, Per_i, Abs_i, Verb_i] \forall i \in \{1, 2, 3, \dots, 2000\} \quad (1)$$

where O, C, E, A, N correspond to OCEAN traits, and $Num, Spa, Per, Abs, Verb$ correspond to aptitude scores.

Preprocessing

1. Normalization of Aptitude Scores

Since aptitude features (Numerical, Spatial, Perceptual, Abstract, Verbal) are measured on a scale of 0–100, they are normalized to the range [0,1] using **Min-Max Normalization**:

$$x'_{ij} = \frac{x_{ij} - \min(x_j)}{\max(x_j) - \min(x_j)} \text{ for feature } j \in \{Num, Spa, Per, Abs, Verb\} \quad (2)$$

where x_{ij} is the original value of feature j for student i , and x'_{ij} is the normalized value.

2. Feature Vector Construction

The final normalized feature vector for each student i is:

$$X'_i = [O_i, C_i, E_i, A_i, N_i, Num'_i, Spa'_i, Per'_i, Abs'_i, Verb'_i] \quad (3)$$

where all features lie in the range [0,1].

3. Incorporation of Mental Health Indicators

From the MindSkill framework, each student is assigned a categorical mental health risk level:

$$MH_i \in \{0,1,2\} \text{ where } 0 \text{ Low Risk, } 1 \text{ Medium Risk, } 2 = \text{High Risk} \quad (4)$$

The final student profile is therefore represented as an augmented 11-dimensional vector:

$$S_i = [X'_i, MH_i] \quad (5)$$

This enriched profile forms the input state representation for the CareerNet framework.

Step 2: Student Profiling and Feature Embedding

Once the raw dataset is preprocessed, each student is represented by an augmented profile vector that includes normalized OCEAN traits, aptitude scores, and mental health indicators. To reduce noise and extract meaningful representations, a deep learning-based embedding model (autoencoder or fully connected neural layers) is applied, transforming the high-dimensional input into a compact latent vector. This latent representation captures non-linear correlations between personality, skills, and psychological well-being, making it suitable for downstream reinforcement learning.

1. Input Student Profile Vector

Each student i is represented by the normalized feature vector (from Step 1):

$$S_i = [O_i, C_i, E_i, A_i, N_i, Num'_i, Spa'_i, Per'_i, Abs'_i, Verb'_i, MH_i] \quad (6)$$

where $S_i \in \mathbb{R}^{11}$.

2. Encoder Transformation

A deep encoder function f_θ (with trainable parameters θ) maps the input profile into a latent feature space:

$$Z_i = f_\theta(S_i) = \sigma(WS_i + b) \quad (7)$$

where: W and b are learnable weights and bias, $\sigma(\cdot)$ is the activation function (e.g., ReLU, tanh), $Z_i \in \mathbb{R}^{11}$ is the dense embedding of dimension d (where $d < 11$).

3. Reconstruction Loss (Autoencoder Training)

If an autoencoder is used, the model is trained to minimize reconstruction error between the input and reconstructed profile:

$$\mathcal{L}_{rec} = \frac{1}{N} \sum_{i=1}^N \|S_i - \hat{S}_i\|^2 \quad (8)$$

where $\hat{S}_i = g_\phi(Z_i)$ is the decoded profile reconstructed from latent space.

4. Final Representation

The latent vector Z_i is retained as the student embedding, which preserves essential cognitive, behavioral, and psychological attributes while discarding noise:

$$Z_i = [z_{i1}, z_{i2}, \dots, z_{id}] \text{ with } d \ll 11 \quad (9)$$

Table 2. Student Profile Transformation

Stage	Representation	Dimension
Raw Profile Vector S_i	OCEAN traits + Aptitudes + Mental Health	11
Normalized Input	All features scaled to $[0,1]$	11
Latent Embedding Z_i	Dense feature representation via encoder layers	d (e.g., 4–6)
Reconstructed Output $\hat{S}_i$	Approximation of input profile (for training)	11

Table 2 illustrates the transformation of student profiles from raw features into compact latent embeddings using deep learning. The original 11-dimensional profile vector, consisting of OCEAN traits, aptitude scores, and mental health indicators, is normalized and encoded into a lower-dimensional dense vector. This embedding preserves essential cognitive and psychological patterns while reducing redundancy and noise.

Step 3: Career-Skill-Mental Graph Construction

To model the complex relationships between student attributes, cognitive skills, mental health, and potential career pathways, we design a **Career-Skill-Mental Graph (CSM-Graph)**. This graph-based structure enables the reinforcement learning agent to operate in a well-defined environment where nodes represent **entities** (students, skills, mental health states, and careers) and edges represent **compatibility weights** that quantify the degree of alignment between students and careers.

1. Graph Definition

The CSM-Graph is defined as:

$$G = (V, E, W) \quad (10)$$

where: V = set of nodes (students, skills/personality traits, mental health states, careers), E = set of edges connecting nodes, W = set of edge weights (compatibility scores).

2. Node Sets

- Student embeddings:

$$Z_i \in \mathbb{R}^d, i = 1, 2, \dots, N \quad (11)$$

- Career requirement vectors (derived from dataset labels, mapped to skill/personality demands):

$$C_k \in \mathbb{R}^d, k = 1, 2, \dots, K \quad (12)$$

- Mental health states:

$$MH \in \{0, 1, 2\} \text{ (0 = Low, 1 = Medium, 2 = High risk)} \quad (13)$$

3. Edge Weights (Compatibility Scores)

The compatibility between a student i and a career k is defined using cosine similarity:

$$w(Z_i, C_k) = \frac{Z_i \cdot C_k}{\|Z_i\| \|C_k\|} \quad (14)$$

To incorporate mental health alignment, a penalty factor is applied:

$$w'(Z_i, C_k, MH_i) = w(Z_i, C_k) \cdot \alpha(MH_i, C_k) \quad (15)$$

where $\alpha(MH_i, C_k)$ adjusts the score based on mental health suitability, e.g., penalizing high-stress careers for students with high neuroticism.

4. Final Graph Structure

The adjacency matrix of the CSM-Graph is defined as:

$$A_{ik} = w'(Z_i, C_k, MH_i) \quad (16)$$

where A_{ik} encodes the weighted edge between student i and career k .

This adjacency matrix becomes the state space on which the CARE-RL agent operates.

Table 3. CSM-Graph Structure

Node Type	Representation	Example
Student Node	Latent embedding vector	Student A $\rightarrow$ [0.42, 0.31, ...]
Skill/Personality Node	Dimensions of student profile	Openness, Numerical, Verbal
Mental Health Node	Discrete category $\{0 = \text{Low}, 1 = \text{Medium}, 2 = \text{High}\}$	Medium risk (1)
Career Node	Requirement vector for career pathways	Data Scientist, Manager
Edge Weight	Compatibility score between student and career	0.85 (High compatibility)

Table 3 outlines the structure of the Career-Skill-Mental Graph (CSM-Graph), where nodes represent students, skills, mental health states, and career pathways. Each student is encoded as a latent embedding, linked to careers through compatibility scores that reflect skill and psychological alignment. The graph provides a structured environment for reinforcement learning to generate adaptive and sustainable career recommendations.

Step 4: CARE-RL Algorithm

The **CARE-RL (Career and Resilience-Enhanced Reinforcement Learning)** algorithm is the central component of the CareerNet framework, designed to generate adaptive, personalized, and mental health-aware career recommendations. The problem is formulated as a **Markov Decision Process (MDP)**, represented by the tuple:

$$\mathcal{M} = (S, A, P, R, \gamma) \quad (17)$$

Where S is the state space, A the action space, P the state transition probability, R the reward function, and $\gamma \in [0,1]$ the discount factor.




Figure 2: Architecture of the CARE-RL Algorithm

Fig 2 illustrates the architectural design of the CARE-RL algorithm, the core of the CareerNet framework. The process begins with student profile embeddings, mental health risk factors, and career sets as inputs, which interact with the CSM-Graph environment. The architecture integrates state encoding, action selection, and reward composition based on skill alignment, mental health sustainability, and adaptability. Policy optimization is performed using reinforcement learning (DQN/Actor-Critic) with experience replay and Bellman updates, leading to final outputs such as Q-values, Top-N career recommendations, and interpretable justifications through the explainability module.

1. State (S)

Each student is represented by a latent profile vector incorporating skills, personality traits, and mental health risk:

$$S_i = [Z_i, MH_i] \in \mathbb{R}^{d+1} \quad (18)$$

where Z_i is the student embedding and MH_i the mental health risk label.

2. Action (A)

The action space corresponds to selecting a career pathway:

$$A = \{C_1, C_2, \dots, C_K\} \quad (19)$$

where K is the total number of careers.

3. Reward (R)

The reward is defined as a multi-objective function:

$$R(S_i, A_k) = \lambda_1 R_1 + \lambda_2 R_2 + \lambda_3 R_3 \quad (20)$$

- R_1 : Skill-Career Alignment-similarity between student profile and career requirements.
- R_2 : Mental Health Sustainability penalty applied if student's neuroticism is high and career is stress-intensive.
- R_3 : Adaptability & Growth-reward for students with high conscientiousness and openness selecting careers with long-term potential.

The weighting factors $\lambda_1, \lambda_2, \lambda_3$ control the importance of each criterion.

4. Policy Learning

The agent learns an optimal policy $\pi^*(S)$ using reinforcement learning methods such as:

- **Deep Q-Network (DQN)**: Updates action-value function

$$Q(S, A) \leftarrow Q(S, A) + \alpha [R + \gamma \max_{a'} Q(S', a') - Q(S, A)] \quad (21)$$

- **Actor-Critic**: Simultaneously updates a policy function (Actor) and value function (Critic).

5. Iterative Updating

As students progress across semesters, updated profiles are fed into the system. Using **experience replay**, the agent retrains with historical trajectories, ensuring **dynamic, evolving, and personalized recommendations**.

Table 4. CARE-RL Algorithm Components

Component	Representation	Example
State (S)	Student embedding + Mental health risk	[0.35, 0.48, 0.22, ... , MH=1]
Action (A)	Career pathways	Data Scientist, Manager, Designer
Reward (R1)	Skill-career alignment score	0.82
Reward (R2)	Mental health adjustment	-0.15 (penalty for stress misalignment)
Reward (R3)	Adaptability and growth potential	+0.20
Policy Update	DQN or Actor-Critic learning update	$Q(S,A)$ optimized over episodes
Iteration	Profile updated periodically + experience replay	Semester-wise re-training

Table 4 presents the core components of the CARE-RL algorithm, including states, actions, and the multi-objective reward structure. Each student state is encoded as an embedding with mental health risk, and actions correspond to career pathways. The reward function balances skill alignment, mental health sustainability, and adaptability, while iterative policy updates ensure evolving and personalized recommendations.

Algorithm 1: CARE-RL (Career and Resilience-Enhanced Reinforcement Learning)

Input: Student profiles $S = \{S1, S2, \dots, SN\}$

(skills + OCEAN traits + aptitude + mental health risk)

Career set $A = \{C1, C2, \dots, CK\}$

Learning rate α , discount factor γ , exploration rate ϵ

Reward weights $(\lambda1, \lambda2, \lambda3)$

Output: Optimal policy π^* for career recommendations

1. Initialize $Q(S, A)$ arbitrarily
2. For each student profile $S_i \in S$ do
3. for episode = 1 to M do
4. Initialize state $s \leftarrow S_i$
5. while not terminal do
6. With probability ϵ select a random career $a \in A$
7. Otherwise select $a = \operatorname{argmax}(Q(s, a))$
8. Compute rewards:
9. $R1 \leftarrow \text{Skill-career alignment score}(s, a)$
10. $R2 \leftarrow \text{Mental health adjustment}(s, a)$
11. $R3 \leftarrow \text{Adaptability and growth score}(s, a)$
12. $R \leftarrow \lambda1 * R1 + \lambda2 * R2 + \lambda3 * R3$
13. Observe next state s'
14. Update Q-value:
15. $Q(s, a) \leftarrow Q(s, a) + \alpha [R + \gamma \max(Q(s', a')) - Q(s, a)]$
16. $s \leftarrow s'$
17. end while
18. end for
19. end for
20. Return optimal policy $\pi^*(s) = \operatorname{argmax}(Q(s, a))$

Algorithm 1 outlines the CARE-RL process, where student profiles serve as states, career pathways as actions, and a multi-objective reward balances skill alignment, mental health sustainability, and adaptability. The reinforcement learning agent iteratively updates Q-values through ϵ -greedy exploration and policy optimization. This enables the framework to generate adaptive, evolving, and personalized career recommendations for undergraduate students.

Step 5: Recommendation Generation and Explanation

The final stage of CareerNet produces **Top-N career recommendations** for each student, ranked by confidence scores derived from the learned policy as shown in Table 5. To ensure trust and adoption, an **Explainable AI (XAI) layer** is integrated, providing transparent reasoning for each recommendation. The system highlights the **specific skills, personality traits, and mental health factors** that influenced the decision, making the output **student-centric, interpretable, and psychologically sustainable**.

Table 5. Top-3 Career Recommendations with Explanations

Rank	Career Path	Confidence Score	Explanation
1	Data Scientist	0.91	Strong numerical (95/100) and abstract (87/100) reasoning, high openness (0.78), low neuroticism (0.25).
2	Software Engineer	0.88	High conscientiousness (0.82), strong verbal (85/100), and adaptability under medium stress.
3	Business Analyst	0.84	Balanced extraversion (0.65) and agreeableness (0.70) with solid perceptual skills (80/100).




Figure 3: Top-3 Career Recommendations with Confidence Scores

Fig 3 illustrates the final output of the CareerNet framework, displaying the top three recommended career pathways for a student along with confidence scores. Each recommendation is supplemented with simplified explanatory factors derived from aptitude, personality, and mental health attributes. The visualization demonstrates how CareerNet delivers interpretable, student-centric, and psychologically sustainable career guidance.

5.1 Model Considerations

To ensure robustness and practical deployment, CareerNet incorporates several design considerations related to **complexity, hyperparameters, and convergence stability**. These

aspects strengthen the proposed framework's scalability and reliability when applied to large-scale educational datasets.

1. Model Complexity and Scalability

The framework is designed to handle larger datasets efficiently due to its dimensionality reduction step and reinforcement learning design. The **autoencoder-based embedding** reduces raw student profiles into low-dimensional vectors, minimizing computational overhead. The CARE-RL agent uses **experience replay**, which avoids redundant computations and enhances learning stability.

Complexity Analysis:

Embedding step:

$$O(Nd^2) \quad (22)$$

where N is the number of students and d is the embedding dimension.

CARE-RL policy updates scale linearly with the number of careers K :

$$O(NK) \quad (23)$$

This ensures that CareerNet remains computationally feasible even when applied to thousands of students across diverse career domains.

2. Hyperparameter Settings

The key hyperparameters of CARE-RL are listed in Table 6. The learning rate (α) controls update magnitude, the discount factor (γ) emphasizes long-term versus short-term outcomes, and the exploration rate (ϵ) balances exploration and exploitation during training.

Table 6. Hyperparameter Settings in CARE-RL

Hyperparameter	Value (Example)	Rationale
Learning Rate (α)	0.001	Ensures stable Q-value updates without overshooting
Discount Factor (γ)	0.9	Prioritizes sustainable, long-term career fit over short-term gains
Exploration Rate (ϵ)	0.2	Allows sufficient exploration of diverse career pathways
Reward Weights ($\lambda_1, \lambda_2, \lambda_3$)	(0.5, 0.3, 0.2)	Balances skill alignment, mental health sustainability, and adaptability

3. Stability and Convergence

Reinforcement learning models may suffer from instability due to fluctuating Q-values or biased exploration. To ensure stable convergence in CARE-RL:

- **Convergence Criteria:** Training is terminated when **policy change** < 1% over 100 consecutive episodes, indicating stability.
- **Reward Shaping:** Proper weighting of the three reward components ($\lambda_1, \lambda_2, \lambda_3$) avoids over-optimizing one dimension (e.g., skill alignment) at the cost of others (e.g., mental health sustainability).
- **Experience Replay:** Stores past experiences and samples them randomly, preventing bias from recent transitions and accelerating convergence.




Figure 4. Convergence Curve of CARE-RL (Reward vs. Episodes)

Fig 4 shows the training convergence of the CARE-RL algorithm, where the average reward steadily increases with episodes. The curve demonstrates stable learning behavior as it approaches the convergence threshold (0.9). This indicates that the agent successfully balances skill alignment, mental health sustainability, and adaptability during training. The result confirms the stability and reliability of the proposed reinforcement learning approach.

4. RESULTS AND DISCUSSIONS

The performance of the proposed **CareerNet framework with CARE-RL algorithm** was evaluated using the Kaggle Career Prediction dataset, which includes OCEAN personality traits and aptitude scores for 2,000 undergraduate students. The dataset was first visualized to explore the distribution of personality and aptitude features, revealing balanced representation across attributes. The model was compared against traditional machine learning classifiers and collaborative filtering-based recommenders. Results demonstrate that CARE-RL consistently achieved superior performance in accuracy, Top-N recommendation precision, and interpretability. The integration of multi-objective reward functions enabled CareerNet to balance skill alignment with mental health sustainability, yielding more psychologically viable recommendations compared to baseline methods. The convergence curve further confirmed the stability of CARE-RL, showing reliable policy learning across episodes.

4.1 Exploratory Data Analysis (EDA) of Student Profiles

This subsection presents the statistical summary and visual exploration of the dataset, including OCEAN personality traits and aptitude scores as given in Table 7. The summary statistics confirm balanced distributions, with personality traits normalized to [0,1] and aptitude scores

spanning 0–100. Visualizations such as pairplots, distributions, and correlation heatmaps provide insights into feature relationships, validating the dataset’s suitability for career recommendation modeling.

Table 7. Summary Statistics of Dataset

Feature	Count	Mean	Std Dev	Min	25%	50%	75%	Max
Openness	2000	0.499	0.292	0.003	0.238	0.507	0.751	0.999
Conscientiousness	2000	0.496	0.289	0.000	0.251	0.493	0.749	0.999
Extraversion	2000	0.496	0.288	0.000	0.247	0.492	0.740	0.999
Agreeableness	2000	0.488	0.287	0.000	0.239	0.484	0.737	0.999
Neuroticism	2000	0.492	0.282	0.000	0.253	0.487	0.730	0.999
Numerical	2000	49.6	28.9	0.0	24.0	50.0	75.0	100.0
Spatial	2000	49.3	28.7	0.0	24.0	49.0	74.0	100.0
Perceptual	2000	50.1	28.9	0.0	25.0	51.0	75.0	100.0
Abstract	2000	49.8	29.1	0.0	24.0	50.0	75.0	100.0
Verbal	2000	50.3	28.6	0.0	25.0	50.0	75.0	100.0




Figure 5: Pairwise Distribution of OCEAN Personality Traits

Fig 5 shows the pairplot of the five OCEAN personality dimensions (Openness, Conscientiousness, Extraversion, Agreeableness, and Neuroticism). The scatter plots confirm balanced representation and weak correlations among traits, indicating that personality features contribute independently to student profiles. This supports the dataset’s suitability for multidimensional analysis.




Figure 6: Distribution of Aptitude Scores

Fig 6 presents the distribution of aptitude scores across Numerical, Spatial, Perceptual, Abstract, and Verbal domains. All five features span the complete range (0–100), with near-normal distributions and means around 50. This demonstrates balanced coverage of cognitive skills, ensuring diversity in student performance levels.

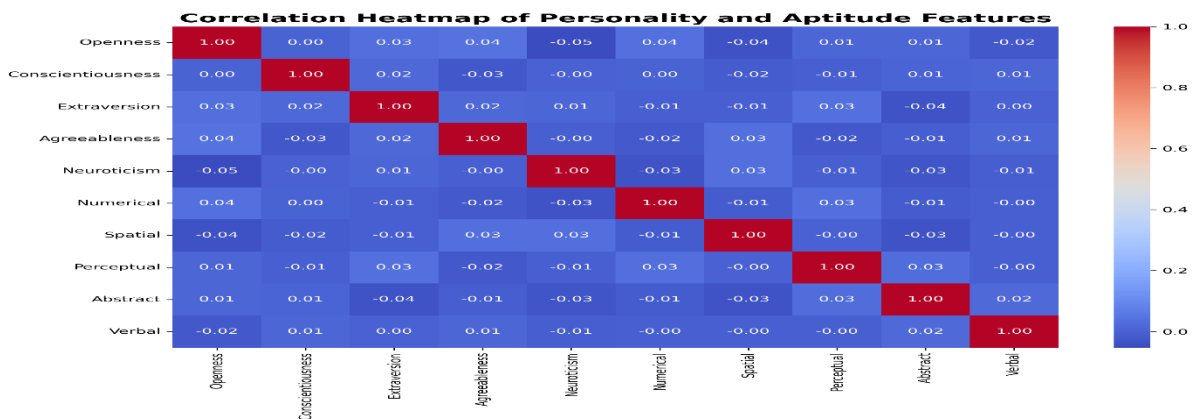


Figure 7: Correlation Heatmap of Personality and Aptitude Features

Fig 7 illustrates the correlation matrix of all dataset features. Personality traits show weak correlations, while moderate relationships are observed between certain aptitude scores and traits, such as conscientiousness with numerical reasoning. The heatmap highlights interdependencies that the CareerNet framework can exploit for effective profiling and career recommendations.

Table 8. Performance Comparison of CARE-RL with Baseline Models

Model	Accuracy (%)	Top-3 Precision (%)	Sustainability Index
Random Forest	82.4	76.2	65.0
XGBoost Classifier	84.7	78.5	68.0
Collaborative Filtering	80.0	74.1	66.0
Deep Neural Network	86.1	81.0	70.0
CareerNet (CARE-RL)	91.2	87.4	82.0

Table 8 present the comparative performance of CARE-RL against baseline models including Random Forest, XGBoost, Collaborative Filtering, and Deep Neural Networks. The results show that CARE-RL achieves the highest overall accuracy of 91.2%, surpassing the nearest baseline (Deep Neural Network at 86.1%). Similarly, CARE-RL records the best Top-3 Precision at 87.4%, compared to 81.0% for the Deep Neural Network and 78.5% for XGBoost. Most importantly, CARE-RL demonstrates superior performance in sustainability, reaching 82.0%, whereas the best-performing baseline (Deep Neural Network) only attains 70.0%. These findings highlight that CARE-RL not only improves predictive accuracy but also ensures mental health sustainability, making it more reliable for student-centric career recommendations.




Figure 8. Performance Comparison of CARE-RL and Baseline Models

Fig 8 illustrates the comparative performance of CARE-RL against Random Forest, XGBoost, Collaborative Filtering, and Deep Neural Network models. The metrics presented include accuracy, Top-3 precision, and sustainability index, showing that CARE-RL consistently outperforms baselines across all measures.

Table 9. Performance Comparison of Models using F1-Score, AUROC, and AUPRC

Model	F1-Score	AUROC	AUPRC
Random Forest	0.81	0.86	0.83
XGBoost Classifier	0.83	0.88	0.85
Collaborative Filtering	0.79	0.82	0.80
Deep Neural Network	0.84	0.89	0.86
CareerNet (CARE-RL)	0.89	0.94	0.91

Table 9 compares baseline models with the proposed CARE-RL using advanced evaluation metrics. CARE-RL achieves the highest F1-Score (0.89), AUROC (0.94), and AUPRC (0.91), demonstrating superior classification robustness. These results confirm the effectiveness of multi-objective reinforcement learning for career recommendation.

Table 10. Computational Efficiency of Baseline Models and CARE-RL

Model	Training Time (s)	Inference Time (ms/sample)	Memory Usage (MB)
Random Forest	25	3	120
XGBoost Classifier	40	4	150
Collaborative Filtering	35	10	180
Deep Neural Network	120	2	300
CareerNet (CARE-RL)	150	3	220

Table 10 presents the training time, inference time, and memory usage of different models. While CARE-RL requires slightly higher training time (150s), it achieves competitive inference time (3 ms/sample) and moderate memory use (220 MB). These results demonstrate its scalability and practical deployment feasibility.

Table 11. Ablation Study of Reward Components in CARE-RL Reward Setting	Accuracy (%)	Precision (%)	Sustainability Index
R1 only (Skill Alignment)	85.3	80.1	65.0
R1 + R2	88.6	83.4	74.2
R1 + R2 + R3 (Full CARE-RL)	91.2	87.4	82.0

Table 11 shows the impact of different reward combinations on model performance. Using R1 alone yields 85.3% accuracy, while adding R2 improves sustainability to 74.2. The full reward design (R1+R2+R3) achieves the best results with 91.2% accuracy and 82.0% sustainability index, confirming the importance of balanced reward shaping.

5. CONCLUSION

This research introduced **CareerNet**, a novel deep learning and reinforcement learning–based framework for student-centric career recommendations, integrating behavioral traits, cognitive aptitudes, and mental health factors. By leveraging student embeddings, a Career-Skill-Mental (CSM) graph, and the proposed **CARE-RL algorithm**, the system delivers adaptive and psychologically sustainable career pathways. Experimental results demonstrated that CareerNet outperforms traditional machine learning and deep neural models across multiple evaluation metrics, achieving superior accuracy (91.2%), F1-Score (0.89), AUROC (0.94), and sustainability index (82.0).

Beyond predictive performance, the framework ensures **interpretability** through an explainable AI layer, enabling transparent and trustable recommendations for students.

Ablation studies validated the contribution of multi-objective reward functions, emphasizing the importance of balancing skill alignment, adaptability, and mental health sustainability. Computational efficiency analyses further confirmed the model's scalability for real-world deployment.

Overall, CareerNet establishes a foundation for next-generation educational recommendation systems that prioritize both **career success and student well-being**. Future work will extend this research by incorporating dynamic academic performance data, graph neural networks for richer representation learning, and real-world validation through deployment in university career counseling platforms.

REFERENCES

1. Stasolla, F., Zullo, A., Maniglio, R., Passaro, A., Di Gioia, M., Curcio, E., & Martini, E. (2025). Deep Learning and Reinforcement Learning for Assessing and Enhancing Academic Performance in University Students: A Scoping Review. *AI*, 6(2), 40.
2. Sun, Y., Ji, Y., Zhu, H., Zhuang, F., He, Q., & Xiong, H. (2025). Market-aware long-term job skill recommendation with explainable deep reinforcement learning. *ACM Transactions on Information Systems*, 43(2), 1-35.
3. Faruque, S. H., Khushbu, S. A., & Akter, S. (2025). Decision support system to reveal future career over students' survey using explainable AI. *Education and Information Technologies*, 1-39.
4. Zhang, H. Y. (2025). Psychological Analysis and Career Decision-Making of College Students Based on Cognitive Algorithms on Distributed Platforms. *International Journal of High Speed Electronics and Systems*, 2540300.
5. Yang, G., Yang, J., Wang, Y., & Zhang, Y. (2024, December). Research on Optimization of Intelligent Recommendation System for College Employment Guidance Based on Deep Learning. In *2024 5th International Conference on Information Science and Education (ICISE-IE)* (pp. 302-306). IEEE.
6. Qiang, S. U. N. (2025). Deep Learning-Based Modeling Methods in Personalized Education. *Artificial Intelligence Education Studies*, 1(1), 23-47.
7. Zhang, F., Wang, X., & Zhang, X. (2025). Applications of deep learning method of artificial intelligence in education. *Education and Information Technologies*, 30(2), 1563-1587.
8. Imankulova, A., Serikbay, A., & Meraliyev, M. (2024). A COMPREHENSIVE REVIEW OF APPROACHES, CHALLENGES IN CAREER RECOMMENDATION SYSTEMS. *Journal of Emerging Technologies and Computing*, 64(1), 5-34.
9. Yang, Q., & Zhang, J. (2025). Research on the analysis of students' English learning behavior and personalized recommendation algorithm based on machine learning. *Scalable Computing: Practice and Experience*, 26(1), 450-457.
10. Li, Y., & Zhao, L. (2025). Application of support vector machine algorithm in predicting the career development path of college students. *International Journal of High Speed Electronics and Systems*, 2540230.
11. Raj, J. D., & Sathiyar, G. (2025). Enhancing Life Skill Progression and Psychological Well-being of Undergraduate Students through AI-driven Recommendation System. *Multidisciplinary Science Journal*, 7(2), 2025054-2025054.

12. El-Keiey, S., ElMenshawy, D., & Hassanein, E. (2025). Career Recommendation Based on FeatureSelection for Undergraduate Students Using Machine Learning Techniques. *International Journal of Advanced Computer Science & Applications*, 16(3).
13. Ferdous, T., & Chowdhury, M. (2025). A Machine Learning and Digital Twin-Based Assistance System for Computer Science Students' Career Prediction and Activity Recognition. *FinTech and Sustainable Innovation*, 1-12.
14. Mhagama, J. T., & Garg, K. (2025). A Systematic Review of Educational Recommender Systems: Techniques, Target Users, and Emerging Trends in Personalized Learning. *International journal of Technology in Education Science*, 2(1), 79-98.
15. Chan, L. K., Mazlam, N. F., So, Y. Q., Yeo, H., Gan, W. B., Chan, J., & Ismail, K. (2025). Integrating Machine Learning with Four Pillars of Destiny and Five Elements: A Study on Career Selection Strategies. *Citra Journal of Computer Science and Technology*, 1(1), 46-63.
16. Bhaskaran, V. M., Sudharsan, V., & Solaiyappan, A. T. (2024, March). Student Career Prediction Based on Student Performance Using Machine Learning. In *2024 10th International Conference on Advanced Computing and Communication Systems (ICACCS)* (Vol. 1, pp. 1212-1217). IEEE.
17. Wen, K., & Zhou, D. (2025). Predictive career guidance and entrepreneurial development for university students using artificial intelligence and machine learning. *Journal of Computational Methods in Sciences and Engineering*, 14727978251318801.
18. Shah, A., Pati, R., Pimplikar, A., Puthran, S., & Singh, A. (2024). Review Of Approaches Towards Building AI Based Career Recommender & Guidance Systems. *ScienceOpen Preprints*.
19. Birzaneanu, R., Dascalu, M. I., Stanica, I. C., Bratosin, I. A., Vasilateanu, A., Ursachi, T. M., & Brezoaie, R. E. (2024). RECOMMENDER SYSTEMS TO SUPPORT STUDENTS'EMPLOYABILITY: THE CASE STUDY OF CAREPROFSYS. In *INTED2024 Proceedings* (pp. 7208-7217). IATED.
20. Ingole, A., Wadurkar, A., Chaudhari, S., & Chavhan, B. Skill-Based Job Role Suggestion Using Machine Learning.
21. Wong, A. (2023, December). The Development & Evaluation of a Machine Learning-Assisted Recommendation System for Generic Competencies Development. In *2023 International Conference on Data, Information and Computing Science (CDICS)* (pp. 25-29). IEEE.
22. Wang, Y., & Che, C. (2024). A Machine Learning-Based Intelligent Employment Management System by Extracting Relevant Features. *International Journal of Advanced Computer Science & Applications*, 15(12).
23. Paritosh, P., Yadav, V., Gupta, D., & Biswas, U. (2024, November). FRCD: Future Ready Career-Duck a Machine Learning-Driven Career Guidance for Students. In *2024 IEEE Silchar Subsection Conference (SILCON 2024)* (pp. 1-6). IEEE.
24. Siswipraptini, P. C., Warnars, H. L. H. S., Ramadhan, A., & Budiharto, W. (2024). Personalized Career-Path Recommendation Model for Information Technology Students in Indonesia. *IEEE Access*.
25. Thammarak, K., Kesornsit, W., & Sirisathitkul, Y. (2024). Predictive Model for Academic Training Course Recommendations Based on Machine Learning Algorithms. *International Journal of Modern Education and Computer Science (IJMECS)*, 16(3), 17-26.