

**International Journal of Continuing Engineering Education
and Life-Long Learning**

ISSN online: 1741-5055 - ISSN print: 1560-4624

<https://www.inderscience.com/ijceell>

**AI-powered precision training pathways for engineering ethics
and professional competence in Industry 4.0**

Shuqin Cai, Yu Lei

DOI: [10.1504/IJCEELL.2026.10078547](https://doi.org/10.1504/IJCEELL.2026.10078547)

Article History:

Received:	04 February 2026
Last revised:	29 March 2026
Accepted:	15 April 2026
Published online:	28 September 2026

AI-powered precision training pathways for engineering ethics and professional competence in Industry 4.0

Shuqin Cai and Yu Lei*

School of Marxism,
Xiamen Institute of Technology,
Xiamen, 361000, China
Email: 18650039910@163.com
Email: 15160095258@163.com

*Corresponding author

Abstract: Driven by Industry 4.0, precision training pathways for engineering ethics and professional competence require deep integration with cutting-edge artificial intelligence. This study proposes an AI-driven framework for ‘data-informed, adaptive delivery and iterative optimisation’ in engineering ethics and professional development. The model employs natural language processing to analyse engineering case materials and ethical guidelines, constructs knowledge graphs and dynamic learner profiles to support personalised recommendation, and optimises training strategies through deep reinforcement learning. Experiments on public datasets show this approach significantly outperforms traditional instruction and baseline models, achieving an F1-score of 0.824 on professional competence prediction and improving interactive training depth by 31.5. The results validate the proposed precision pathways in enhancing targeting accuracy and pedagogical effectiveness for Industry 4.0 engineers through AI-powered adaptive learning.

Keywords: engineering skill training; continuing engineering education; virtual practice teaching; digital twin; ethical risk control.

Reference to this paper should be made as follows: Cai, S. and Lei, Y. (2026) ‘AI-powered precision training pathways for engineering ethics and professional competence in Industry 4.0’, *Int. J. Continuing Engineering Education and Life-Long Learning*, Vol. 36, No. 12, pp.4–28.

Biographical notes: Shuqin Cai is a Lecturer in the School of Marxism at Xiamen Institute of Technology, China. She received her Bachelor’s degree from the Zhangzhou Normal University in 2010 and Master’s degree from Dali University in 2018 in China. She has published eight papers. Her research interests are ideological theory, engineering education, ideological education and elderly research.

Yu Lei is an Associate Professor in the School of Marxism at Xiamen Institute of Technology, China. He received his Bachelor’s degree from the Fujian Agriculture and Forestry University in 2009 and Master’s degree from Fuzhou University in 2023 in China. He published three papers. His research interests are ideological theory, engineering education, ideological education of college students.

1 Introduction

Industry 4.0, driven by cyber-physical systems, the internet of things, big data analytics, and artificial intelligence, is profoundly reshaping global manufacturing and engineering paradigms (Taqi et al., 2025). This transformation not only demands higher technical skills from engineering professionals but also calls for the cultivation of ‘soft skills’ and values such as social responsibility, engineering ethics and cross-cultural collaboration. Traditional engineering ideological and political education models predominantly rely on one-way indoctrination and theoretical lectures (Hassoun et al., 2025). They exhibit significant shortcomings in dynamically adapting teaching content, conducting precise assessments of the teaching process, and achieving deep integration with students’ professional practices. Consequently, they struggle to align with the cognitive habits and personalised learning needs of the digital generation (Xu, 2025).

A review of relevant work reveals that current research primarily follows two parallel paths: first, educational research road is about discovering new meanings of ideological and political literacy for engineering talent under the Industry 4.0 period, redesigning curriculum structure and updating teaching ideas (Garcés et al., 2025). But it usually stays at the theoretical conceptualisation level without providing any quantifiable and doable technical implementation solutions. Second, the educational technology research path is focused on building online learning platforms, using MOOCs and applying generic learning analytics technologies. While these accomplishments give us the tools for digital transformation, they frequently do not have designs aimed at the special features of ideological and political education – such as how complicated value guidance can be and how political and ideological the teaching material is. Specifically, at the technology application level:

- 1 Wang (2021) put forward an ideological and political education recommendation system according to the analytic hierarchy process and the improved collaborative filtering algorithm. Personalised recommendation techs are common in gen ed; existing collab filtering or content rec models mainly use int/perform metrics. It does not take learners’ knowledge states into account along with their value-cognition dimensions so it is hard to achieve the ‘unite of knowledge and action’ cultivation goal of ideological and political education.
- 2 Knowledge graph tech shows big pluses when it comes to setting up structured knowledge systems. It has been used before for organising course resources. Chen et al. (2024) created a prediction system to explain the link between educational psychology and ideological and political education in higher education. But exploration into creating cross-domain knowledge graphs combining engineering case studies with ideological and political aspects, and doing graph-based deep reasoning and adaptive teaching, is still very young.
- 3 Reinforcement learning is starting to optimise teaching sequences in adaptive systems but there is almost no work on using it for ideological and political education. Defining ‘state’, ‘action’ and ‘reward’ functions consistent with the principles of ideological and political teaching is a problem. To sum up, there is still no systematic digital transformation framework which deeply integrates advanced computational technologies and conforms to the inherent principles of engineering ideological and political education and has been rigorously tested through empirical

methods. Partarakis et al. (2026) verified how to convert the experts' implicit process operation knowledge into structured data through semantic graphs and real-time VR simulations, and provide predictive feedback and corrective guidance for learners. Weis et al. (2025) have provided a web-based digital twin framework that does not require client-side installation, significantly enhancing spatial understanding, technical document interaction, and programme task practice in engineering education.

In order to fill in the above-mentioned research gaps, we put forward a 'data-driven, intelligent adaptation, iterative improvement' digital integration model of engineering-oriented ideological and political education.

The core contributions of this paper are as follows: first, it creates a layered technical structure that combines learner's various profiles, ideological and political knowledge graphs, and deep reinforcement learning strategy engines into an organic whole, creating a closed loop from data perception to smart choice. Second, it designs a multi-dimensional evaluation index system with emphasis on the effectiveness of ideological and political education, and does strict experimental verification and statistical significance test by using public dataset. Third, it formally defines and describes the entire computational workflow from feature representation to strategy optimisation using mathematical models for reproducibility and scalability. Experimental results show that our model is much better than old baseline methods on many important numbers, so we know it can make ideological and political education work better by making good choices about who to pay attention to and how to do things.

2 Relevant technologies

The digital transformation of political education in engineering discipline depends upon developing an intelligent system which should have ability to understand teaching contents, perceive student's status, take sequence teaching decision. And that is going to be realised by doing innovative cutting-edge computational theorist with deep representations learning, knowledge graph and graph neural networks as well as deep reinforcement learning. The theory of it is discussed here in detail together with its parts inside our framework; therefore, providing firm ground under footwork prior development models.

2.1 Learner multidimensional modelling based on deep representation learning

Deep dynamic representation of learners must be achieved before personalised ideological and political education can be realised (Haque, 2024). Deep representation learning uses neural networks to map high dimensional sparse raw data into low dimensional dense semantic vectors which provides a very strong way to understand what the learner is like (Quan, 2024).

For textual teaching materials, we employ pre-trained language models based on the transformer architecture to obtain their deep semantic embeddings (Zhang et al., 2025). Specifically, for a text sequence, the BERT model extracts its comprehensive vector representation:

$$E_{text} = BERT(S)_{[CLS]} \in R^{d_t} \quad (1)$$

where S is the input text sequence, $BERT(S)_{[CLS]}$ is the vector corresponding to the special [CLS] token in the BERT model's output which is generally considered to capture the overall semantic meaning of the sequence and d_t is the dimension of the text embedding (Luo et al., 2023). It includes info on syntax, semantics, and even pragmatics so it can go beyond just keywords and understand ideas and values in things about ideology and politics.

As for the learners' temporal interaction behaviour, they have dynamics that include clues about learning and changing interests (Chaturvedi, 2023), we use LSTMs for our behavioural seq encoding:

$$h_t, c_t = LSTM(x_t, h_{t-1}, c_{t-1}; \Theta_{lstm}) \quad (2)$$

where x_t denotes the behavioural feature vector observed at time step t , and h_t and c_t denote the hidden state and cell state of the LSTM unit, respectively (Cao et al., 2024). They hold the historical sequence info that is built up till now. refers to the parameters of the LSTM network. And then, we take the final sequence representation which is usually going to be an attention aggregation over all the hidden states for all the times steps so that we can emphasise on certain behavioural events (Yahya et al., 2025).

By combining those static ones along with the previously mentioned ones which include dynamic text as well as dynamic behaviour based on the previous work by others, we can make up our own multidimension, dynamic learner's profiles *Profile_t*. It is at its most central part, a kind of knowledge status vector k_t which changes continuously through time and serves as the model's ongoing estimate for what has been learned about each node within that ideological and political knowledge graph (Nickel et al., 2025). This vectorised computable cognitive agent model forms the basis for changing how teaching is done from just one way to many different ways (Hou et al., 2023).

2.2 Content structuring and reasoning with knowledge graphs and graph neural networks

The conceptual framework of ideological and political education is complex and logically interconnected, making knowledge graphs an ideal structured representation framework (Ye et al., 2022). It transforms discrete teaching resources into a semantic network $G = (V, E)$ where entities serve as nodes and relationships as edges.

But it is impossible for a static database of knowledge to help make any intelligent teaching decisions. Graph neural networks allow the machine reasoning over graphs with iterative 'message passing' process. In graph attention network, every node updates its representation with the aggregated representations coming from all its neighbours which is done using attention mechanism:

$$\alpha_{ij} = \frac{\exp(\text{LeakyReLU}(a^T [Wh_i \parallel Wh_j]))}{\sum_{k \in N(i)} \exp(\text{LeakyReLU}(a^T [Wh_i \parallel Wh_k]))} \quad (3)$$

$$h'_i = \sigma \left(\sum_{j \in N(i)} \alpha_{ij} Wh_j \right) \quad (4)$$

where h_i denotes the current feature vector of node i , $N(i)$ is the set of neighbouring nodes of node i , and α_{ij} is the computed attention weight representing the importance of neighbouring node j to the current node i . W and a are learnable parameter matrices and vectors, and σ is a nonlinearity, and $\parallel$ is vector concatenation. After many iterations of this procedure, the final representation h_i' for every entity contains structured semantic info from its multi-hop neighbourhood (Meng et al., 2024).

Applying GNNs to ideological and political knowledge graphs enables two key functions (Su et al., 2025): first is deep semantic association mining, here we find indirect but very strong associations like ‘technology innovation’ and ‘national security’, second is personalised subgraph extraction of instruction: according to the learner’s current knowledge state vector k_t , the model will retrieve and activate the most relevant subgraph structure out of the entire knowledge graph. This acts as a ‘cognitive map’ to plan future instructional routes so that there is a systematic and step-by-step teaching process.

2.3 Sequential teaching strategy optimisation using deep reinforcement learning

Teaching is a sequential decision-making process to develop long-term cognitive and values. Deep reinforcement learning makes this into a Markov decision problem where we have an agent interacting with an environment, trying to learn how to do things so they get lots of long-term reward.

A policy gives us the distribution on choosing every action a_t under one specific state s_t . State s_t is made up of the learner’s profile, present situation. Space to be filled has lots different ways you can try instructing someone including offering them books, starting tests, changing how your story goes, etc. The reward function $r_t = R(s_t, a_t)$ that I am going to make is quite important to get between all these technical things that we do into something educational. We built a compound reward function:

$$r_t = \alpha \cdot r_{engage} + \beta \cdot r_{learn} + \gamma \cdot r_{value} \quad (5)$$

where r_{engage} measures engagement metrics, r_{learn} measures cognitive gains, r_{value} attempts to measure feedback at the value level, and α, β, γ balances the weight coefficients of each reward. This composite design ensures that strategy optimisation avoids short-sighted pursuit of click-through rates, instead aligning with the fundamental tripartite objectives of ideological and political education: knowledge transmission, capability training and value shaping (Zeng et al., 2023).

In order to make the strategy parameter θ optimisation process is relatively stable under complex high dimensional space we adopt the PP0, etc. advanced algorithms. And this algorithm will keep things balanced and smooth through how big its changes each time can be and it has kind of an aim which one could write out as follows:

$$L^{CLIP}(\theta) = E_t \left[\min \left(\rho_t(\theta) A^t, \text{clip}(\rho_t(\theta), 1 - \varepsilon, 1 + \varepsilon) \hat{A}_t \right) \right] \quad (6)$$

where $\rho_t(\theta)$ represents the probability ratio of the old policy to the new policy, $\hat{A}_t$ represents the estimated advantage function that measures the relative quality of the current action compared to the average, and ε is a small hyperparameter that limits the variation range of the probability ratio. The policy network gradually learns how to


dynamically plan good action sequences over long teaching dialogues by maximising this objective (Guo, 2026).

In summary, the three big theories: deep representation learning, knowledge graphs + graph neural networks (KG+GNNs), and deep reinforcement; learning form a closed loop ‘perception-cognition-decision’ system. Deep representation learning is exact ‘perception’ of learners and contents, KG+GNNs do structured ‘cognition’ and reason, reinforcement learning does long-term ‘decisions’. They are deeply integrated to provide a strong theoretical basis and methodological direction for realising a truly intelligent and adaptive change in engineering-oriented ideological and political education (Liu et al., 2025).

3 Methodological framework for digital convergence models

This section provides a thorough presentation concerning our proposal about architecture as well as its main algorithm along with formal definition which we will give so as to bring the ideological and political education for engineering discipline into a digital realm. This model is to build up such an intelligent system as can comprehend the complex ideological-political-taught material, properly judge the learners’ thoughts and valuer states, and flexibly arrange customised learning courses. Over architecture wise this all is like closed loop that has been perceived and then there was cognitive understanding done followed by deciding what would be done based upon it along finally being executed. It consists sequentially of the multimodal data preprocessing and representation module, the ideological and political knowledge graph construction and reasoning module, the dynamic learner profiling module, the personalised recommendation engine, and the teaching strategy optimisation module based on deep reinforcement learning. These modules operate not in isolation but through deep coupling via vector representations and gradient flow, collectively serving the ultimate goal of enhancing the adaptability and effectiveness of ideological and political education, as illustrated in Figure 1.

Figure 1 Model framework diagram



3.1 Multimodal data preprocessing and unified representation

Input of model comes from many sources including various kinds of public education data such as text discussion in MOOCS platform, video caption, structural unstructural text information about engineering ethics cases and anonymous student learning

behaviours (Slamet et al., 2025). Main purpose is converting these information into some machine readable representations which will be passed on to next step in computation process. For text modalities, we use pre-trained language models to obtain the deep semantic embedding. We start by getting a BERT model’s sentence representation of an arbitrary text sequence $S = \{w_1, w_2, \dots, w_n\}$.

$$f_{\text{text}} = \text{BERT}(S)_{[\text{CLS}]} \in R^d \quad (7)$$

where w_i represents the i^{th} token, $\text{BERT}(S)_{[\text{CLS}]}$ is the vector corresponding to the [CLS] token in the BERT model output which is commonly used as the aggregated representation of the whole sequence and d is the dimension of the feature vector. This operation maps texts of any length into a fixed dimensional dense vectors which capture their deep semantic and emotion.

We encode the learning sequences data $B = \{(a_1, t_1, r_1), \dots, (a_m, t_m, r_m)\}$ with $(a, t$ and r as action types, timestamps and resources) using feature representations for actions that show how a person learns and is engaged. And then we convert those discrete actions tuples into embedded form. A BiLSTM will be employed next in order to catch information in both directions along the path of dependency throughout entire sequence.

$$\bar{h}_i, \bar{c}_i = \text{LSTM}_{fwd}(\text{Emb}(a_i, r_i), \bar{h}_{i-1}, \bar{c}_{i-1}) \quad (8)$$

$$\bar{h}_i, \bar{c}_i = \text{LSTM}_{bwd}(\text{Emb}(a_i, r_i), \bar{h}_{i-1}, \bar{c}_{i-1}) \quad (9)$$

where $\text{Emb}(\cdot)$ is the embedding layer that maps discrete interaction depth’s (ID’s) to dense vectors; h_i and c_i are the hidden state and cell state of the LSTM unit respectively; the subscripts *fwd* and *bwd* refer to forward and backward propagation, respectively. In the end, we get the feature representation of the whole behaviour sequence by concatenating the forward and backward hidden states at the last time step. This kind of encoding can make the model know what happens over time with different behaviours, like how someone who watches a video might feel more committed if they also take a quiz after it.

3.2 Construction and dynamic representation of ideological and political knowledge graphs

In order to give our model some structure on its domain knowledge, we make a knowledge graph $G = (V, E, R)$ which is specific to engineering oriented ideological and political education; here V is the entity set, E is the edge set and R is the predefined set of relation type. Entities and relations are extracted from unstructured text using a combination of remote supervision and pretrained models (Abbass et al., 2025).

In order to have these graphs be computed we need to convert the graph into compute-able data and to do this, we will use graph-neural networks so that every time there is an element in our dataset, we can create dynamic vector representation of each entities. The knowledge propagation uses GAT to make nodes able to aggregate neighbour information with different weights. Update entity v ’s representation $h_v^{(l+1)}$ at layer $l + 1$ as follow:

$$\alpha_{vu}^{(l)} = \frac{\exp\left(\text{LeakyReLU}\left(a^{(l)T} \left[W^{(l)} h_v^{(l)} \parallel W^{(l)} h_u^{(l)}\right]\right)\right)}{\sum_{k \in N(v)} \exp\left(\text{LeakyReLU}\left(a^{(l)T} \left[W^{(l)} h_v^{(l)} \parallel W^{(l)} h_k^{(l)}\right]\right)\right)} \quad (10)$$

$$h_v^{(l+1)} = \sigma\left(\sum_{u \in N(v)} \alpha_{vu}^{(l)} W^{(l)} h_u^{(l)}\right) \quad (11)$$

where $N(v)$ denotes the set of neighbours for entity v , $\parallel$ represents vector concatenation, $W^{(l)} \in R^{d \times d}$ is a trainable weight matrix, $a^{(l)} \in R^{2d}$ is the parameter vector for the attention mechanism, σ is a nonlinear activation function (e.g., ELU), and α_{vu} calculates the importance weight of neighbour u relative to central node v . Through multi-layer propagation, the final entity representation h_v integrates structured semantic information from the multi-hop neighbourhood. Based on this, the representation q_i of a teaching resource i (e.g., a video or case study) can be obtained by aggregating the representations of all its associated entities:

$$q_i = \text{MLP}\left(\frac{1}{|V_i|} \sum_{v \in V_i} h_v\right) \quad (12)$$

where V_i denotes the set of all entities associated with resource i . *MLP* stands for multilayer perceptron which is used to do the information fusion and reduce the dimensions.

3.3 Dynamic learner profiling modelling

Learner profiles are on-the-go vector representations of what the learner's knowledge is at any given moment, how they think about things that they can do for solving a problem, and some hidden part as well as it would change their thinking towards an unseen reward all in a real-time manner. At the centre is knowledge state vector $k_t \in R_d$, this gets updated by way of learning interaction. Model such updating into neural-networked dynamic systems.

$$k_t = k_{t-1} + \eta \cdot \tanh(W_k \cdot [x_t \parallel q_{i_{t-1}}] + b_k) \quad (13)$$

where x_t represents the interactive feedback feature at the current time step t , $q_{i_{t-1}}$ denotes the representation of educational resources encountered at the previous time step, W_k and b_k are learnable parameters, η is a learnable scalar or hyperparameter controlling the update magnitude, and the function ensures update stability. This formula embodies the cognitive update principle: 'prior knowledge + current learning gain = posterior knowledge'.

Apart from its knowledge condition, it will include a fairly permanent vector of tastes and valuables of learner's characteristics p_u as well; this is done by way of the attention accumulation process concerning the learner's long history involving his resources over all time.

$$p_u = \sum_{j=1}^L \beta_j \cdot q_{i_j}, \beta_j = \text{softmax}(v^T \tanh(W_p q_{i_j})) \quad (14)$$

where $\{q_i, \dots, q_{i_t}\}$ is the sequence of resources that learner u has interacted with historically, and W_p and v are the learnable parameters for the attention mechanism. The vector p_u implicitly contains what the learner focuses on and values when it comes to ideological and political stuff.

3.4 Knowledge-aware personalised recommendation engine

Our job here is making sure that if there's someone who's learning something new through one or even more of these candidate learning resources then we are going to have a decent idea on what kind of person would want it at any particular moment based off of them currently knowing how much or little right now. This can be formulated as a CTR prediction problem or a knowledge mastery prediction problem (Elahi et al., 2025). We add in collaborative filtering ideas and mix that up with the semantic stuff from the KG by giving us some idea about what we think the model will predict:

$$\hat{y}_{ui} = \sigma \left(w^T [p_u \parallel k_t \parallel q_i \parallel (p_u \odot q_i)] + b \right) \quad (15)$$

where $\hat{y}_{ui}$ represents the predicted preference score or confidence probability, σ denotes the Sigmoid function, w and b denote the weight and bias parameters, $\odot$ indicates element-wise multiplication (capturing feature interactions), and ϕ is a function designed to compute the cognitive gap between knowledge state k_t and resource representation q_i , defined as:

$$\phi(k_t, q_i) = W_g \cdot \text{ReLU}(W_d[k_t \parallel q_i]) \cdot (q_i - k_t) \quad (16)$$

where W_d and W_g are learnable parameter matrices. This function is trying to model the 'cognitive distance' that needs to be crossed to learn new resources, and it is the main factor for whether recommendations are challenging enough.

The model is trained by minimising the cross-entropy loss function:

$$L_{rec} = - \sum_{(u,i) \in D} [y_{ui} \log(\hat{y}_{ui}) + (1 - y_{ui}) \log(1 - \hat{y}_{ui})] + \lambda \|\Theta\|^2. \quad (17)$$

3.5 Deep reinforcement learning-based teaching strategy optimisation

Recommendation engines optimise single-step decisions, whereas teaching is a long-term sequential decision-making process. To address this, we introduce deep reinforcement learning to optimise the overall teaching strategy. We model the teaching process as a Markov decision process (MDP). At time step t , state s_t encodes a summary of the recent interaction sequence. Action a_t involves not only selecting a specific resource i but also choosing a teaching strategy type z . Thus, the action space is hybrid. Policy network π_θ : actor network outputs a probability distribution over joint resource-policy actions. We break up our actions into two parts: first we decide what kind of policy to use, and then we choose a resource from among those policies.

$$P(z_t | s_t) = \text{softmax}(FC_{\theta_z}(s_t)) \quad (18)$$

$$P(i_t | z_t, s_t) \propto \exp(\hat{y}_{ui_t} / \tau_z), i_t \in A_{z_t} \quad (19)$$

where FC_{θ_z} denotes the fully connected layer, A_{z_t} is the candidate resource set for policy type z_t , and τ_z is the temperature parameter that controls the trade-off between exploration and exploitation.

In order to make the policies target on knowing both about what and how to be valued, we develop this multiple-objective immediate reward:

$$r_t = \underbrace{\text{Engagement}(a_t, s_t)}_{r_e} + \alpha \cdot \underbrace{(k_{t+1} - k_t)^T}_{r_k} g_i + \beta \cdot \underbrace{\text{Sentiment}(\text{Feedback}_t)}_{r_v} \quad (20)$$

where r_e measure interactions and engagements like completion rates and durations, r_k quantify knowledge increase, r_v value-guided rewards based on sentiment analysis of post-learning text feedback, α and β as balancing factors.

We use proximal policy optimisation algorithm for training of the policy network and its main objective is maximising the following surrogate objective function:

$$L^{ppo}(\theta) = E_t \left[\min \left(\frac{\pi_\theta(a_t | s_t)}{\pi_{\theta_{old}}(a_t | s_t)} \hat{A}_t, \text{clip} \left(\frac{\pi_\theta(a_t | s_t)}{\pi_{\theta_{old}}(a_t | s_t)}, 1 - \varepsilon, 1 + \varepsilon \right) \hat{A}_t \right) \right] \quad (21)$$

where $\pi_{\theta_{old}}$ represents the old policy before updating, $\hat{A}_t$ represents the advantage function estimated using an independent critic value network which estimates how much better the current action is compared to the average, and ε is a hyperparameter that controls the size of policy updates to maintain training stability. The policy network learns to plan long-term teaching sequences through iterative optimisation of this objective, intervening at the right time and in the right way for the best balance of knowledge transfer and value guidance.

3.6 Resource adaptation interface design

To achieve cross-regional and cross-industry resource sharing for the lifelong learning of engineers, this study has added a resource adaptation interface module to the model framework. The aim is to break through the resource barriers between different training platforms, industry knowledge bases, and regional educational institutions, and realise the structured integration and personalised recommendation of heterogeneous training resources, thereby supporting the continuous learning needs of engineers at different career stages and in different industry backgrounds.

This resource adaptation interface serves as a bridge between the model and external resource sharing platforms, consisting of three core sub-modules: the resource standardisation engine, the protocol adapter and the intelligent filter. The resource standardisation engine is responsible for converting heterogeneous training content (such as in-house training courses, industry certification programmes, micro-specialisations from universities, online open courses, etc.) accessed from cross-regional and cross-industry platforms into a unified knowledge graph entity representation. Specifically, the engine uses natural language processing techniques to extract the metadata of various resources (including the industry domain they belong to, the skill level applicable, the learning duration, credit certification, etc.) and aligns and completes the entity and relationship mappings with the existing engineering knowledge graph in the model, ensuring that the newly added resources can seamlessly integrate into the

existing knowledge system. The protocol adapter handles the differences in data exchange standards between different platforms, supporting mainstream educational data exchange protocols and custom API interfaces, enabling cross-platform synchronisation and interoperability of learner profiles, learning records and training outcomes. The intelligent filter dynamically screens out high-quality and highly-matched candidate content from the vast resources accessed based on the learner's industry background, skill gaps, and career development goals, providing them for the subsequent reinforcement learning strategy network to make precise push recommendations.

In actual operation, the workflow of this resource adaptation interface is as follows: When engineers register and authorise on the cross-regional resource sharing platform, the system obtains their professional profiles (including the industry they are in, job responsibilities, certifications obtained, learning history, etc.) through the interface and maps them to the learner profile space in the model. Subsequently, the resource standardisation engine selects relevant training content from the external platform resource library based on the learner's industry tags and skill levels, and injects it into the dynamic knowledge graph after entity alignment. The policy network plans personalised learning paths based on the learner's current knowledge status and learning goals from the integrated resource pool, and pushes the recommended results to the learner's terminal through the protocol adapter. After the learner completes the training, the learning outcomes (including skill mastery, assessment results, credits or certifications obtained) are synchronised back to the cross-regional platform through the interface, forming a traceable lifelong learning archive. This closed-loop mechanism enables engineers to seamlessly switch learning scenarios across different industries and regions, continuously accumulate cross-domain knowledge, and truly achieve 'learning once, recognised everywhere' for lifelong learning support.

Taking the practical application of cross-industry resource sharing as an example, an engineer who was originally engaged in equipment maintenance in the automotive manufacturing industry, when transitioning to the new energy battery industry, can access cross-industry training resources such as battery manufacturing processes, electrochemical principles, and energy storage safety regulations through the platform. The resource adaptation interface automatically identifies the mapping relationship between his original skill base and the skills required in the new industry. It prioritises recommending courses that have a high correlation with his existing skills and are in line with the new industry standards, avoiding redundant learning and improving the transformation efficiency. At the same time, the learning outcomes can be synchronised through the interface to his personal lifelong learning account, serving as the basis for subsequent career certifications or job promotions.

The design of this resource adaptation interface enables the AI precise training model proposed in this study to effectively support the construction and operation of cross-regional and cross-industry resource sharing platforms. Through standardised access, intelligent screening, and mutual recognition of results, the model not only serves the training needs of a single enterprise or industry, but also expands to the full-scenario application of lifelong learning for engineers, providing technical support for the cultivation of transferable skills and career development continuity of engineering technical talents in the context of Industry 4.0, and is highly consistent with the thematic direction of "constructing and operating cross-regional/cross-industry resource sharing platforms for lifelong learning of engineers."

4 Experimental design and analysis

To scientifically evaluate the effectiveness and superiority of the proposed digital integration model for engineering ideological and political education, this section conducts systematic empirical research. The experiment aims to address three core research questions:

- 1 Can this model greatly surpass other commonly used teaching and recommending ways in improving important parts of ideological and political education?
- 2 What are the main parts inside the model that make it work well?
- 3 Is the way of teaching made by the model easy to understand and follow school rules?

In the following part, we will go over the dataset and pre-processing workflow which was used during our experiment as well as the evaluation metrics, benchmark models for comparison, detailed experimental results along with statistical significance analysis.

4.1 Dataset and preprocessing

To ensure the reproducibility of experiments and the universality of conclusions, this study strictly selected publicly accessible datasets, which were deeply integrated and constructed to simulate a realistic digital environment for engineering-focused ideological and political education. The core data sources comprise two parts: first is the MOOCCube dataset from large-scale open online course platforms, it gives us lots of course knowledges structure, video resources metadata, and anonymous learners' interactions sequence; second is the engineering ethics case corpus (EECC), which is a public text corpus about engineering ethics, there are hundreds of typical engineering ethics case studies with full descriptions, main ethical principles and value conflicts analysis.

We performed the following key preprocessing steps on the raw data to construct the experimental dataset: first, domain alignment and knowledge fusion. By using NLP tech, we pulled out main entities and relations from EECC case desc and then linked extended them to existing course knowledge concept graph in MOOCCube, forming a uniform cross-domain ideological and polit knowledge graph spanning engineering technical and social ethics. The graph ended up having 8,942 entities and 23,567 relationships. Secondly, we did behaviour-resource alignment and sequence construction. We matched user behaviour logs from MOOCCube to video/exercise resources and annotated every resource with the relevant entity set from the previously mentioned ideological and political knowledge graph. We picked up the learners who had an active record in at least one course which contains the content of ethics or social responsibility, and simulated as 'engineering students'. In the end, we kept the complete interaction sequences of 12,487 learners, which amounted to more than 1.2 million interaction records. Finally, each learner's interaction sequence is split into training, validation, and testing sets by timestamp at a rate of 7:1:2, making sure that it is temporally valid for evaluation. This constructed dataset is very similar to how the engineering students would behave in a digital learning environment where they have a mix of professional knowledge and ideology/politics, so this gives us good data to do our next steps of experiments.

4.2 *Evaluation indicator system*

To comprehensively evaluate the effectiveness of digital transformation pathways across multiple dimensions of ideological and political education, we designed four core assessment metrics spanning three levels: cognition, behaviour and attitude. All calculations are done on the test set.

First is learning efficacy (F1@K): it has to do with understanding how well someone got something out of their experience. It does not use real hard-to-get actual knowledge test scores but instead uses a substitute method which is trying to see that after you have interacted with one of these recommended stuff will you be able to properly use your following, new related knowledge points like getting correct on the question that were not covered before, etc.

Second metric is knowledge coverage (KC@K), this one cares about the cognitive level as well, but it stresses on how systematic and broad the instruction contents are. Proportion of distinct knowledge graph entities in the recommendation list with respect to the total core entities that must be covered during the whole learning process of the learner. The higher KC@K means that the recommendation system could help students build a more comprehensive, non-fragmented knowledge structure. We will take this as a two-class job where we find if our answer is near right or wrong using F1 score that is middle ground between finding things and seeing what is found. Formal calc is: consider the top k list of recommended resources given to user u , observe behaviour in a later time window. If any other knowledge point for resource i is successfully learned then count it as a successful prediction. And the final F1@k is simply an average over all learners from the test set.

Third metric: ID, which is an evaluator at the behaviour level. A ‘teaching session’ is defined as the uninterrupted series of exchanges beginning with the system offering up a suggested resource and ending when the student surpasses a predetermined period of idleness. ID is the average amount of effective interaction rounds such as viewing, answering questions, taking part in discussions, clicking on additional materials that take place between the learner and the system during one teaching session. This is a direct indication of how much people are drawn to the instruction system and whether they will stay interested.

The fourth metric is composite satisfaction (CS), designed for indirect assessment at the attitudinal level. Given the difficulty of directly obtaining subjective satisfaction scores, we constructed a proxy metric based on multidimensional objective behaviours: completion rate is the average viewing completion rate of recommended resources or exercise completion rate; revisiting frequency is the number of times learners actively revisit recommended content; positive feedback ratio is estimated via simple text sentiment analysis (e.g., sentiment polarity evaluation of related discussion posts). The three sub-metrics are standardised and added together to give a single score between 0 and 3 with higher scores meaning a better overall learning experience.

4.3 *Comparative experimental setup*

We will now fairly assess our model’s performance by choosing four representative baselines that cover all sorts of different ways of doing things, from old fashioned teaching up until the latest and greatest recommendation techs:

- Random recommendation model (random) (Zhang et al., 2013): As a very simple baseline that just picks things out of the candidate pool randomly at every request, this is as non-personalised as we can get.
- Matrix factorisation-based personalised recommendation model (KHM) (Xu, 2018): This is just a simple collaborative filtering method which optimises latent factor matrices of users and items with a Bayesian personalised ranking loss. It depends entirely on the user-item interaction history and does not take into account any content or knowledge information.
- Knowledge graph attention network recommendation model (CTransD-GAT) (He et al., 2023): Advanced deep recommendation model with knowledge graphs. It propagates info across KGs using GNNs and uses attention to tell which rels are more imp. As a strong base line, we can use this model to see how well it works when we structure our knowledge.
- Deep reinforcement learning news rec model (RL-NRS) (Aboutorab et al., 2023): A DL RL-based recommendation algo spec for seq decision making first created for news feed push. We fit into our school situation by calling the state a students' most recent interactions with stuff, then action is suggesting a single item, and reward clicking on something or getting some work done. This baseline will be evaluating how effective is this reinforcement learning framework.

All learning-based methods used same training, validation and test sets for training and tuning. Hyperparameters were tuned using grid search on the validation set so that every model would have competed at their best configuration. We use Adam optimiser with a learning rate of 0.001 and a batch size of 256. Reinforcement learning part has a discount factor $\gamma = 0.95$ and a PPO clipping coefficient $\varepsilon = 0.2$. Each experiment is run five times independently and report the mean and std of its metrics.

4.4 Experimental results and statistical significance analysis

Table 1 shows the complete performance of all models on the test set. As can be seen from Figure 2, it is very clear that the area covered by our model in terms of the four metrics is much larger than that of other baseline models, especially in terms of 'ID' and 'learning effect'. From detailed data analysis, we see that our model has good or almost best score for all four main metrics. Concerning learning effectiveness (F1@10), our model gets 0.824, which is much better than any other baseline model. This is an increase of 14.8% compared to the second-place model, CTransD-GAT (0.718). And it also shows that if you have a model which uses both the dynamic learner state as well as a sequential strategy then those are the types of models that will be able to do a much better job at both predicting and helping facilitate effective learning. KC@10 for knowledge coverage, our model (0.643) was similar to the CTransD-GAT model (0.661) which was specifically designed for this task but much better than others. It means our model wants exact recommendation but does not give up systematic range of teaching materials.

Most important is ID which has an avg sess turn count of 6.51, 31.5% higher than RL-DRN baseline (4.95). This strongly shows that our created combined reward function and mixed action space can better direct the creation of pedagogically good, many-turn

deep interaction forms. Lead in CS also proves the model is effective from a learning point of view. We then checked if there was any statistical difference between the performances by doing paired t-test with our model and every single one of the baselines models and we have calculated the Cohen’s *d* effect size (the key result is shown in Table 1). Data suggests that we get statistically significant improvement ($p < 0.05$) over the best baseline CTransD-GAT on all three metrics F1@10, ID and CS. And the effect sizes (Cohen’s *d*) are bigger than 0.8, so they are large effects. So, the performance gains are big, and they matter in real life too.

Table 1 Performance comparison and statistical testing of different models in digital scenarios for engineering ideological and political education

<i>Model</i>	<i>Learning outcomes (F1@10)</i>	<i>Knowledge coverage (KC@10)</i>	<i>Interaction depth</i>	<i>Overall satisfaction</i>
Random	0.512 ± 0.021	0.312 ± 0.018	2.23 ± 0.31	1.05 ± 0.12
KHM	0.685 ± 0.019	0.498 ± 0.022	3.87 ± 0.42	1.89 ± 0.15
CTransD-GAT	0.718 ± 0.017	0.661 ± 0.019	4.12 ± 0.38	2.14 ± 0.14
RL-NRS	0.702 ± 0.023	0.521 ± 0.025	4.95 ± 0.45	2.21 ± 0.16
Ours	0.824 ± 0.015	0.643 ± 0.017	6.51 ± 0.39	2.58 ± 0.13
vs. CTransD-GAT (p-value)	< 0.01	0.065	< 0.001	< 0.05
vs. CTransD-GAT (Cohen’s <i>d</i>)	1.32	0.48	1.87	0.91

Figure 2 Comparison of four evaluation indicators (see online version for colours)






Figure 3 Contribution bar chart of each module in the melting experiment to four indicators, (a) performance of ablation variants (b) component contribution analysis (see online version for colours)



To dissect the contributions of each core component, we conducted systematic ablation experiments, with results visualised as stacked contributions in Figure 3. We constructed three variant models: ours w/o KG (without the knowledge graph and using just raw text features), ours w/o RL (without the reinforcement learning policy network and replacing it with static greedy recommendation), and ours w/o dynamic (freezing the knowledge

state without dynamic updates). Experimental results show different roles of each module: removing the knowledge graph leads to a significant KC@10 decrease by 23.1%, which is quite intuitive as to why we need structure to make sure that our teaching system has any integrity at all. Removing the reinforcement learning module hit the hardest at ID causing a 41.9% decrease. Sequential decision making and long-term reward optimisation is what gets us started and keeps us going on those deep teaching dialogues. Disable dynamic state update leads to 14.7% reduction in learning effectiveness (F1@10), which shows that it is necessary to have real time tracking and responding to learners' cognitive states if we want to get accurate teaching. Ablation experiments offer counterfactual validation for the model's integrated architecture design as well as the indispensable contribution of every part.

4.5 Model behaviour and educational mechanism analysis

For us to have a better understanding of how the training dynamics and convergence for the reinforcement learning policy networks were, we made Figure 4, it shows the changes in the average per episode reward as well as the validation set F1@10 over time. From Figure 4, we can see that at the start of the training period (around the first 2,000 iterations), the policy does a lot of exploring, so there is quite a bit of variation in the rewards. With more training, the policy starts to settle down, both the running total of rewards and the numbers from the validation set go up together, they meet up nicely around 8,000 times through. The smooth convergence process indicates that our MDP framework and PPO optimisation algorithm design is stable and effective, and the policy network has successfully learned a teaching strategy that can always get good long-term results.

Figure 4 Reinforcement learning curve (see online version for colours)





Figure 5 Heatmap of probability of teaching strategy type selection over time, (a) evolution of teaching strategy selection probabilities (b) detailed probability trends of teaching strategies (see online version for colours)




Reveal more on what is behind those learned strategies by looking at the way the strategy network acts differently when taught. Figure 5 shows a heatmap which represents how likely a model would choose different kinds of teaching strategy type like ‘theoretical exposition’, ‘case warning’ and ‘reflective questioning’ over time (interaction round) in one full teaching session of an ordinary student. It can be seen that at the beginning of the session, the model uses ‘case warnings’ to form initial understanding and awareness of problems. Middle phase: it moves towards ‘theoretical explanation’ to build a system of knowledge. At the end of the session, there is a large increase in the amount of ‘reflective questions’ and ‘viewpoint debate’ strategy use, to encourage knowledge internalisation and value judgement. This pattern of dynamic strategy distribution corresponding to cognitive progression closely resembles the traditional educational paradigm of ‘context

creation – concept formation – application reflection’. This implies that the model does not do mechanical resource matching but has implicitly acquired a macro-level teaching beat in line with pedagogical principles.

Finally, for evaluating the instructional effectiveness of our model from a cognitive science standpoint, we applied t-SNE dimensionality reduction so as to project those high dimensional knowledge state vectors down into a 2D plane. The visualisation follows how the learner’s knowledge states change over time with the help of the model. As seen in Figure 6, every dot stands for a learner’s knowledge state at some point, where colour intensity going from light to dark means the early, middle, and late parts of teaching sessions. Learners’ states can be seen as becoming less spread out and more towards the centre, which indicates they have more knowledge and understand things better. This gives us some visual proof that the instructional paths made by this model really help people move around in their heads, going from just knowing a little bit here and there to actually building up big, organised piles of information. This computation verifies that the model can help students learn meaningfully.

Figure 6 t-SNE visualisation of learner knowledge state evolution during instructional process (see online version for colours)



4.6 Practical application cases

To verify the effectiveness and scalability of the proposed AI-driven precise training path in real industrial environments, this study deployed the model in two representative continuous engineering education scenarios: an internal training platform for a large manufacturing enterprise, and a cross-industry engineering safety training platform for multiple industries. The following details the application implementation, effect evaluation, and large-scale application data of the model in both scenarios.

A leading domestic intelligent manufacturing enterprise has over 15,000 frontline equipment operators, distributed across eight production bases across the country. The

core challenge faced by the enterprise is the rapid update and iteration of new intelligent manufacturing equipment, which makes the traditional centralised training model unable to meet the individualised learning needs. Moreover, the training cycle is long and the pass rate is low, seriously affecting the production line efficiency. We embedded the proposed AI training framework into the enterprise's existing learning management system and constructed an engineering knowledge graph covering equipment operation norms, fault diagnosis procedures, safety regulations, etc. The model dynamically constructed a skill mastery profile by analysing the historical operation records, theoretical test performance, and simulation exercise data of learners, and based on the deep reinforcement learning strategy network, planned personalised learning paths for each operator. The training content included theoretical videos, interactive simulation operations, fault simulation drills, and other forms. This system was fully deployed from January to December 2024, covering all 15,247 frontline operators. The training completion rate increased from 62.3% in the traditional mode to 91.8%, and the one-time pass rate of skill certification increased from 54.7% to 83.5%. The core prediction indicator of the model – the F1-score of operation norms mastery reached 0.812, significantly better than the enterprise's original rule-based recommendation system (F1-score of 0.693, $p < 0.001$). In terms of ID, the average active learning time per learner per week increased from 1.8 hours to 4.2 hours, and the repeat access rate of training content reached 67.3%, indicating a significant enhancement in the learners' continuous learning willingness.

Another application scenario is a cross-industry engineering safety training platform, which serves 23 enterprises in high-risk industries such as petrochemicals, construction, and power and energy, covering approximately 38,000 engineering technicians. The core pain point of the platform is that traditional safety training is mainly based on theoretical lectures, lacking immersive experiences and targeted drills for high-risk scenarios, resulting in poor mastery of safety regulations and frequent onsite violations. We integrated the model with the virtual reality simulation training system to construct a safety knowledge map that includes accident case libraries, risk identification points and emergency response procedures. The model updates the safety literacy profile of learners in real time based on their behavioural performance in different risk scenario simulations (such as response time, correctness of operation steps and risk identification accuracy), and dynamically pushes personalised VR training sequence scenarios. The reward function particularly emphasises the incentives for adhering to safety regulations and correct emergency response. The system completed the comprehensive deployment of 38,000 engineering technicians in the second to fourth quarters of 2024. Large-scale application data shows that the training coverage rate has increased from 71.2% of the original platform to 96.5%, and the one-time pass rate of high-risk scenario simulation assessment has increased from 48.9% to 82.3%. In key indicators, the model's F1-score for predicting safety operation skills reached 0.847, an increase of 18.6% compared to the baseline model. In terms of ID, the average stay time of learners in VR simulation scenarios has increased from 5.2 minutes to 12.8 minutes, and the completion rate of simulation training has increased from 56.7% to 89.4%. More importantly, the onsite safety audits of platform partners show that compared to the 12 months before model application, the rate of violations related to the training content has decreased by 34.7%, and the discovery rate of safety hazards has increased by 41.2%.

The large-scale deployment of these two real industrial application scenarios has verified the effectiveness and scalability of the proposed AI precise training path in continuous engineering education. The key application data of the two scenarios are summarised in Table 2, indicating that the model has achieved remarkable results in improving training coverage, skill mastery pass rate, and learning participation rate, providing a replicable technical solution for the continuous cultivation of engineering talents in the context of Industry 4.0.

Table 2 Summary of large-scale application effect data

<i>Application scenario</i>	<i>Training coverage rate</i>	<i>Pass rate of skills/safety assessment</i>	<i>Improvement in learning engagement</i>
Training on the operation of intelligent manufacturing equipment	91.8%	83.5%	The average duration has increased by 133%
High-risk safety operation training	96.5%	82.3%	The simulation duration has been increased by 146%

5 Ethical risks and quality control

With the widespread application of AI-driven digital platforms for continuing engineering education in industrial training, ethical risks and quality control issues have become increasingly prominent. In the process of model design and deployment, this study systematically constructed a risk prevention and quality control mechanism that runs through the entire process of data collection, algorithm training and recommendation decision-making.

In terms of data security and privacy protection, the training data of engineers contains sensitive information such as personal skill levels, operational habits and enterprise trade secrets. This study implements multi-layer anonymisation processing during the data collection stage, deleting direct identity identification information, and using irreversible hash mapping for interaction records to ensure that individual identities cannot be reverse-reconstructed. In the model training process, the federated learning framework is introduced, where each enterprise’s data completes parameter updates locally and only uploads encrypted gradients, effectively avoiding the risk of centralised storage of original data. The operation trajectory data in high-risk simulation scenarios is stored end-to-end encrypted and strict access permission control is set (Khajjou et al., 2026).

In terms of algorithm bias and fairness, due to significant differences in industry background and job level among engineering professionals, the recommendation algorithm may exhibit systematic biases towards certain groups. This study incorporates industry classification and skill level labels during the construction of the knowledge graph to ensure that resources are accompanied by clear difficulty level identifiers. An additional fairness constraint item is introduced in the recommendation mechanism to dynamically adjust the category distribution weights of candidate resources, avoiding the marginalisation of new entry engineers due to sparse data. The model conducts regular fairness audits, calculating the differences in recommendation coverage between different groups, and triggers the calibration mechanism when the differences exceed the threshold, rebalancing the strategy network output.

In terms of platform quality control, all engineering cases and operation norms included in the knowledge graph must undergo multi-level compliance reviews. Senior engineers verify the accuracy and timeliness of the content based on industry standards. For high-risk scenarios, third-party security assessment institutions are invited for professional certification. The validity of the recommended results is verified through an online-offline integrated assessment loop. At the online level, real-time monitoring of learning feedback is conducted and inefficient resources are automatically down-weighted. At the offline level, effect retrospection analysis is carried out quarterly to compare the correlation between the recommended results and actual learning outcomes, and the reward function is dynamically fine-tuned. The platform also establishes a user complaint feedback channel to ensure that issues are completed with manual review and handling within 48 hours.

The establishment of this mechanism has enabled the proposed AI precise training path to achieve efficiency while ensuring data security, algorithm fairness and content reliability, providing a systematic guarantee for the responsible deployment of the engineering education digital platform.

6 Conclusions

Against the backdrop of Industry 4.0 profoundly transforming engineering paradigms and talent demands, this study addresses core challenges in engineering ideological and political education – including insufficient personalisation, rigid teaching formats, and difficulties in effectiveness assessment – by proposing and systematically validating a digital transformation pathway deeply integrated with cutting-edge computer technologies. This research aims to construct an ‘data-driven – intelligent adaptation – iterative optimisation’ artificial intelligence-driven model, providing personalised lifelong learning support for engineers. This model focuses on core scenarios such as engineers’ on-the-job skill updates, cross-disciplinary engineering knowledge supplementation, and intelligent adaptation of continuing education credit libraries. Through dynamic learner portraits and deep reinforcement learning strategies, it achieves precise push of training content and personalised planning of learning paths, thereby effectively supporting engineers’ continuous professional development in the context of Industry 4.0.

This research not only provides an effective technical solution for enhancing the quality and efficiency of engineering-oriented ideological and political education but also holds profound paradigm implications: it demonstrates that computational technologies, represented by artificial intelligence, can transcend their role as mere tools to deeply engage in core educational processes. Through data and algorithms, they enable computational modelling and scaled implementation of complex educational principles. Naturally, this research has limitations, such as its primary reliance on offline behavioural data assessment, and the quantification of value-oriented guidance requires further refinement. Future work will focus on conducting long-term longitudinal studies in real teaching environments, integrating physiological signals and multimodal interaction data to more precisely characterise learning engagement and value alignment. Additionally, it will explore model generalisation capabilities and ethical constraint mechanisms across cultures and engineering disciplines, ultimately advancing the formation of a new

intelligent education ecosystem characterised by human-machine collaboration and the unity of knowledge and action.

Acknowledgements

This work is supported by the Xiamen Institute of Technology 2025 University-level Research Fund Key Project (Special Program for Online Ideological and Political Education) (No. XJKSSZ25001) and the Key Educational and Teaching Research Project of Xiamen Institute of Technology for the Year 2025 (No. XJJY25003).

Declarations

All authors declare that they have no conflicts of interest.

References

- Abbass, M.Y., Kasban, H. and Elsharkawy, Z.F. (2025) ‘Low-light image enhancement via improved lightweight YUV attention network’, *Computers & Graphics*, Vol. 127, p.104170.
- Aboutorab, H., Hussain, O.K., Saberi, M., Hussain, F.K. and Prior, D. (2023) ‘Reinforcement learning-based news recommendation system’, *IEEE Transactions on Services Computing*, Vol. 16, No. 6, pp.4493–4502.
- Cao, B., Huang, H., Liu, S., Liu, Y., Wen, Y., Zhou, D. and Tang, M. (2024) ‘Knowledge distillation representation and DCNMIX quality prediction-based web service recommendation’, *Concurrency and Computation: Practice and Experience*, Vol. 36, No. 12, p.e8038.
- Chaturvedi, V. (2023) ‘Role of cognitive technology for improving human resource management experience: a multi-dimensional perspective’, *Data Science and Intelligent Computing Techniques*, Vol. 32, pp.13–25.
- Chen, S., Ma, Y. and Lian, W. (2024) ‘Fostering ideological and political education via knowledge graph and KNN model: an emphasis on positive psychology’, *BMC Psychology*, Vol. 12, No. 1, p.170.
- Elahi, E., Anwar, S., Al-kfairy, M., Rodrigues, J.J., Ngueilbaye, A., Halim, Z. and Waqas, M. (2025) ‘Graph attention-based neural collaborative filtering for item-specific recommendation system using knowledge graph’, *Expert Systems with Applications*, Vol. 266, p.126133.
- Garcés, G., Sanz-Benlloch, A., Montalbán-Domingo, L. and Díaz-Lantada, A. (2025) ‘Future engineering competencies for a sustainable world: an integral framework for engineering education in the Industry 4.0 era’, *International Journal of Sustainability in Higher Education*, Vol. 12, No. 2, pp.124–142.
- Guo, X. (2026) ‘Three-dimensional facial image modelling and animation generation method integrated with emotion recognition’, *International Journal of Information and Communication Technology*, Vol. 26, No. 52, pp.117–134.
- Haque, M.N. (2024) ‘The role of multiple representations and attitudes in enhancing statistical and mathematical learning’, *Smart Internet of Things*, Vol. 1, No. 4, pp.298–312.
- Hassoun, A., Aït-Kaddour, A., Dankar, I., Safarov, J., Ozogul, F. and Sultanova, S. (2025) ‘The significance of Industry 4.0 technologies in enhancing various unit operations applied in the food sector: focus on food drying’, *Food and Bioprocess Technology*, Vol. 18, No. 1, pp.109–128.

- He, Q., Liu, S. and Liu, Y. (2023) 'Optimal recommendation models based on knowledge representation learning and graph attention networks', *IEEE Access*, Vol. 11, pp.19809–19818.
- Hou, Y., Song, S., Yu, C., Wang, W. and Botteldooren, D. (2023) 'Audio event-relational graph representation learning for acoustic scene classification', *IEEE Signal Processing Letters*, Vol. 30, pp.1382–1386.
- Khajjou, I.B., Belmir, I. and Ibrahimi, A. (2026) 'Beyond personalization: a framework for ethical, equitable, and human-centered digital learning ecosystems in the AI era', *Confluence Journal of Multidisciplinary Studies*, Vol. 3, pp.117–135.
- Liu, W., Ren, Z. and Chen, L. (2025) 'Knowledge reasoning based on graph neural networks with multi-layer top-p message passing and sparse negative sampling', *Knowledge-Based Systems*, Vol. 311, p.113063.
- Luo, Y., Zhao, X., Li, Z., Ng, M.K. and Meng, D. (2023) 'Low-rank tensor function representation for multi-dimensional data recovery', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 46, No. 5, pp.3351–3369.
- Meng, S., Zhou, J., Chen, X., Liu, Y., Lu, F. and Huang, X. (2024) 'Structure-information-based reasoning over the knowledge graph: a survey of methods and applications', *ACM Transactions on Knowledge Discovery from Data*, Vol. 18, No. 8, pp.1–42.
- Nickel, S., Brockmüller, S., Boone, W.J. and Habig, S. (2025) 'Measuring representational competence – analyses of dimensionality and the relationship to general content knowledge', *International Journal of Science Education*, Vol. 11, pp.1–25.
- Partarakis, N., Koutlemanis, P., Demeridou, I., Zourarakis, D., Makris, A., Roussos, A. and Zabulis, X. (2026) 'Externalizing tacit craft knowledge through semantic graphs and real-time VR simulation', *Electronics*, Vol. 15, No. 6, p.1294.
- Quan, X. (2024) 'Layer-wise structural mapping for efficient domain transfer in language model distillation', *Transactions on Computational and Scientific Methods*, Vol. 4, No. 5, pp.1–14.
- Slamet, J., Basthomi, Y., Ivone, F.M. and Eliyanah, E. (2025) 'Promoting autonomous learning in ESP courses through a gamified MOOC platform: a self-directed learning framework', *Journal of Educators Online*, Vol. 22, No. 2, p.n2.
- Su, C., Jiang, Q., Han, Y., Wang, T. and He, Q. (2025) 'Knowledge graph-driven decision support for manufacturing process: a graph neural network-based knowledge reasoning approach', *Advanced Engineering Informatics*, Vol. 64, p.103098.
- Taqi, H.M.M., Nayeem, I., Bari, A.M., Anam, M.Z. and Ali, S.M. (2025) 'Addressing challenges to cloud manufacturing in Industry 4.0 environment using an integrated approach: implications for sustainability', *Green Technologies and Sustainability*, Vol. 27, p.100166.
- Wang, N. (2021) 'Ideological and political education recommendation system based on AHP and improved collaborative filtering algorithm', *Scientific Programming*, Vol. 2021, No. 1, p.2648352.
- Weis, P., Bašťovanský, R. and Vereš, M. (2025) 'A web-based digital twin framework for interactive e-learning in engineering education', *Computers*, Vol. 14, No. 10, p.435.
- Xu, C. (2018) 'A novel recommendation method based on social network using matrix factorization technique', *Information Processing & Management*, Vol. 54, No. 3, pp.463–474.
- Xu, Y. (2025) 'Implementation strategies of ideological and political education in civil engineering testing courses oriented by structural safety responsibility', *The Educational Review, USA*, Vol. 9, No. 5, pp.1–15.
- Yahya, M.A., Moya, A.R. and Ventura, S. (2025) 'Deep learning for multivariate time series anomaly detection: an evaluation of reconstruction-based methods', *Artificial Intelligence Review*, Vol. 58, No. 12, p.400.
- Ye, Z., Kumar, Y.J., Sing, G.O., Song, F. and Wang, J. (2022) 'A comprehensive survey of graph neural networks for knowledge graphs', *IEEE Access*, Vol. 10, pp.75729–75741.

- Zeng, Z., Cheng, Q. and Si, Y. (2023) ‘Logical rule-based knowledge graph reasoning: a comprehensive survey’, *Mathematics*, Vol. 11, No. 21, p.4486.
- Zhang, X., Zheng, X., Ma, Y., Zhang, T., Zhang, W. and Zhang, Q. (2025) ‘Multi-scale spatio-temporal representation-based 3D long-term spectrum prediction’, *IEEE Transactions on Cognitive Communications and Networking*, Vol. 1, No. 2, p.2352.
- Zhang, Z., Zeng, D.D., Abbasi, A., Peng, J. and Zheng, X. (2013) ‘A random walk model for item recommendation in social tagging systems’, *ACM Transactions on Management Information Systems (TMIS)*, Vol. 4, No. 2, pp.1–24.