Innovation and Evaluation of AI-Driven Economics Experimental Course Teaching Model

Min Zhou[1], Xu Feng[1]
[1]Faculty of New Business, Anhui Sanlian University, Hefei 230601, Anhui, China
Corresponding author: Min Zhou (e-mail: zhoumin202291@163.com)

ABSTRACT To address the shortcomings of current outcome-oriented static evaluation mechanisms in economics experimental teaching, which are lagging behind and unable to accurately measure the evolution of students' higher-order economic decision-making thinking, this paper proposes an AI (Artificial Intelligence)-driven innovation and dynamic evaluation system for economics experimental teaching, integrating a large language model and a time-series tracking algorithm. First, in terms of teaching model innovation, an LLM (Large Language Model) agent with dynamic risk preferences and memory mechanisms is introduced as a virtual market participant to reconstruct a highly realistic and adaptive economic game scenario. Second, a seamless multimodal data acquisition engine is constructed, which, in addition to extracting clickstream and latency data, utilizes an NLP (Natural Language Processing) model to deeply analyze the textual sentiment of negotiation dialogues; combined with economic theory, the fragmented interaction logs are transformed into quantitative economic features representing the degree of risk probing and information processing resistance. Finally, in terms of effect evaluation and feedback design, a deep tracking model based on the attention mechanism is constructed to accurately map the feature matrix to students' economic cognitive state, and then calculate a dynamic decision rationality index. Simultaneously, by focusing on the extraction of key price adjustment decision nodes using the model's attention weights, high theoretical interpretability of the process evaluation is achieved, and this is used as a trigger condition to provide students with LLM-driven real-time personalized intervention feedback. The deployment results of this system in the course "Two-Way Auctions and Oligopoly Markets" showed that the decision rationality index of the experimental group students significantly increased by 0.278 compared to the traditional group, and the dynamic comprehensive evaluation score output by the system highly matched the students' final professional ability performance (Pearson correlation coefficient r=0.88). This study effectively addresses the interpretability concerns of the traditional teaching evaluation's 'algorithm black box' by establishing a differentiable mapping between model attention weights and specific economic decision-making indicators, thereby providing a new path for the reconstruction of the economics experimental teaching paradigm in the digital age.

INDEX TERMS economics experimental teaching,large language model,dynamic evaluation,attention-bilstm,process assessment

I. INTRODUCTION

THE field of economics education is experiencing a transformational moment with regard to digital technology; experimental teaching that provides a connection between the theoretical nature of economics and how it works within a marketplace has taken on greater significance [1-2]. At the same time, growing numbers of highly interactive digital experimental platforms capture all aspects of a student's experiences in a simulated gaming environment, including all bid modifications, strategies tested, and negotiations (all within a specified timeframe). As a result, these platforms are generating an enormous amount of multimodal "data-richness" at the base level of teaching [3-4]. On the other hand, the rapid advancement of technology has outpaced the existing experimental teaching evaluation system, which is still firmly rooted in its "outcome-dependent" trap [5-6]. Further complicating the situation, traditional static scoring measures have become outmoded; typically, the only information offered about each of the simulated cases is based on the overall profits (surplus) from the simulated terminal or a subjective report completed by the students at the end of their term. Therefore, not only is there a complete waste of the enormous volume of interaction data being generated, the dynamics associated with growth in a student's ability to make higher-order economic decisions based on complex and uncertain market conditions are not measured, thus creating a core contradiction at this time within economics experimental education between "large volumes of data" and "delayed feedback in student evaluation" [7-8]. While the AI-EETES system achieves low inference latency, its primary value lies in bridging this 'lag' by translating real-time scores into a dynamic pedagogical dashboard. This allows instructors to move beyond passive observation to active, data-informed intervention as the experiment unfolds.

In recent years, the explosive growth of artificial intelligence technology has provided a cutting-edge path to break through the aforementioned evaluation dilemma, and scholars at home and abroad have carried out extensive explorations in the fields of intelligent game simulation and educational data mining [9-10]. On the one hand, large language models (LLMs), due to their excellent logical reasoning and generalization capabilities, have begun to be introduced into the simulation of microeconomic environments. However, in terms of controlling the heterogeneous risk preferences of agents and maintaining high fidelity in continuous multi-round dynamic interactions, existing technologies are still crude, resulting in a lack of sufficient environmental tension in the virtual market [11-12]. On the other hand, deep learning-based sequence tracking models have shown great potential in predicting learners' cognitive states. However, in the specific scenario of economic experiments, which highly emphasizes overall game and the dependence of previous and subsequent strategies, conventional one-way sequence algorithms are difficult to effectively extract implicit economic cognitive features from unstructured behavioral logs [13-14]. More importantly, the "black box" nature of traditional deep neural networks leads to a serious lack of interpretability of economic theory in the evaluation results [15-16]. The common shortcoming of existing research is that it has never been able to couple the adaptive high-simulation environment generation engine with the automated process evaluation algorithm with strong time-series tracking ability and high theoretical interpretability at the underlying level, thus failing to substantially empower the teaching feedback system [17-18].

Addressing the aforementioned research gaps and teaching challenges, this paper aims to overcome the limitations of traditional static evaluation by proposing and validating an AI-driven economics experimental teaching innovation and dynamic evaluation model (AI-EETES, AI-driven Economics Experiment Teaching innovation and Evaluation System) that integrates large language models and deep knowledge tracing technology. In terms of implementation logic, this paper first reconstructs a highly realistic and adaptive dynamic economic game scenario by introducing an LLM agent with dynamic risk preference and memory mechanisms. The second main focus of this research is to develop an integrated approach to modality data collection. An effective data collection engine records the steps in each students' bidding and the amount of time students take to make their decisions, as well as thoroughly analyzing the unstructured text contained in each negotiation window. The researchers utilize theories of behavioural economics to transform previously fragmented underlying logs into quantitative features that capture "risk preference probing" and "information processing resistance." A bidirectional long short-term memory (LSTM) network with an integrated attention mechanism is then built and trained to create a high-precision mapping from individual, discrete operations to the continuous economic cognitive state.

This research shows that significant progress has been made in the real-time and accuracy of cognitive state prediction (F1 score = 91.2%) for the model being used in the course 'Two-Way Auctions and Oligopolistic Markets'; the end-to-end inference latency is stable at 50 msec. At an empirical teaching level, Decision Rationality Index (DRI) improved by an average of 0.278 for subjects in the experimental group and the output of the system's dynamic score is highly correlated with students' final competence score (Pearson r = 0.88). Additionally, this research not only solves the technical problems that exist with dynamic cognitive measurement in complex game theory settings, but it also provides a reliable generalizable engineering solution for recreating the economics experimental teaching paradigm in a digital environment.

II. RECONSTRUCTING THE TEACHING SCENARIOS OF AI-EETES

A. HIGH-FIDELITY DYNAMIC GAME ENVIRONMENT BASED ON LLM AGENTS

Within the layer of generating experimental data and constructing scenarios is set the LLM's ability to act as an autonomous trading agent, which makes independent trade decisions. The first stage of the initialization of the consolidated trading agent is to map the macroeconomic parameters to the market microstructure rules through "structural mapping" [19-20]. For example, structuring the initial capital, range of basic values for commodities, quantity supplied and demanded create structural prompts and subsequently is provided directly to the context of the LLM to create an initial boundary for the virtual trader at which trades occur using the rules given in the previous step, but without loss of generality (e.g. provide at least two different samples for this market operation). Subsequently, it allows the trader to create his/her own benchmark trading strategy based upon ongoing disturbances to market conditions, thereby providing a logical basis for the continued output of high-frequency trade interaction data for the future.

The primary method to fixing the issues with game environments being distorted by the repetitiveness of how everyone acts in the virtual marketplace is to require that LLMs (Longitudinal Market Models) generate the market with different levels of risk by putting hard limits on their ability to act in relation to the economy [21-22]. To configure an agent as a player in a game, it applies a constant relative risk aversion utility function to the agent's reasoning chain. Therefore, the agent's goal is strictly defined as being to maximize the maximum expected utility:

$$U(w) = \frac{w^{1-\rho} - 1}{1 - \rho} \tag{1}$$

The agent's current level of utility ($U(w)$) at wealth state w equals the amount of wealth held by the agent. In addition, ρ is a coefficient of relative risk aversion. In addition, the level of price adjustment of the LLM based on uncertain payoffs is managed by adjusting the parameter ρ between values less than zero (risk seeking), to equal to zero (risk neutral), to greater than zero (risk averse). Each round of the real-life game will utilize the principle of utility maximization to quantitatively evaluate trade potential and expected value and

expected risk of trades. This generates printed outputs that emulate the mindset and risk pricing methods of each trader and how their different mental accounts behave within the LLM framework.

During the continuous interaction phase, environmental state transitions at a high speed are determined by the real-time pricing instructions of the human student. The LLM agent accesses its long-range memory tool after receiving a state vector with associated attributes; including the order book depth, the historical average of transaction prices, and the students most recent bid. Using causal reasoning in the context of past events, the LLM generates real-time market feedback with many unknowns based on a range of dynamic variables; including the ability to generate reverse pricing instructions, make judgments about the completion of transactions, and allocate transaction volumes. The resulting closed loop data interaction resolves many of the deficiencies that are commonly associated with conventional deterministic automated trading programs, which exhibit neither "bounded rationality" nor "adaptive expectations"; and replicate the randomness and micro-structural complexity with respect to price movement as observed in the actual market. To address potential LLM hallucinations and the gap in human irrationality, the system implements a constrained reasoning chain that forces agents to operate within strict economic utility boundaries. By adjusting the risk aversion coefficient (ρ), the LLM is programmed to exhibit 'bounded rationality' and specific cognitive biases, ensuring the virtual market mimics human economic friction rather than idealized or erratic behavior. More importantly, this environment encourages students to explore new strategies, modify parameters and engage in real-time negotiation behaviour frequently thus providing a continuous "high signal-to-noise ratio" representation of the interactive data generated in a non-intrusive multimodal DB acquisition engine referred to in Section 2.2 and achieving a coupling of the generation of the environment and feature extraction.

Integration Mechanism of Non-Intrusive Behavior Data Collection and Process Evaluation

High-fidelity simulation environments provide an essential value in teaching because they continually prompt students' strategic exploration and cognitive repetition, and accurately measuring this process requires a fullcycle, continuous flow of behavioral data. The system injects a global event listening matrix into the bottom layer of the experimental interface to construct a non-intrusive data acquisition protocol (corresponding to the Input Layer on the left side of Figure 1). This mechanism intercepts the underlying interaction instructions on the client's Document Object Model (DOM) tree in real time, specifically covering three types of core data flows: first, DOM Clickstream, which includes mouse coordinate trajectory, click landing point, and high-dwelling areas of the page [23]; second, Decision Latency Logs, which is the time difference sequence between server-side rendering events and client-side submission events [24]; and third, Textbox Input Stream, which is the synchronously captured unstructured negotiation window text stream [25]. The above heterogeneous underlying logs are transmitted to the server-side storage pool via a WebSocket bidirectional encrypted channel in a high-frequency asynchronous stream manner, forming an original multimodal data lake, without any manual intervention or interface interruption. While the collection of high-frequency clickstream and trajectory data inherently introduces computational overhead, the system utilizes a specialized lightweight asynchronous observer script that minimizes CPU cycles on the client side. By offloading the primary processing tasks to the server and maintaining a low-footprint event listener, the system ensures that the overhead remains negligible and does not perceptibly degrade the rendering speed or the student's hardware performance.

To address the issues of timestamp misalignment and uneven sampling frequency caused by cross-terminal network latency, a fixed time window aggregation algorithm is used for alignment and resampling during the preprocessing stage, corresponding to the construction process of Feature Sequence: Time Steps in Figure 1 [26]. Each node in the standard time axis receives a discretised representation of the continuous asynchronous event stream by first assigning a uniform time slice span $\Delta\tau$ followed by discretizing the continuous asynchronous event stream across each node. In order to eliminate any local data missing as a result of the time window being segmented, the first order Lagrange interpolation polynomial is utilized. This smooths the transitions between the events thereby removing sequence breakpoints and maintaining continuity in the time domain. Critically, the system employs an adaptive interpolation threshold to ensure that only transmission-induced data gaps are filled, while the high-frequency sampling (6 samples per second) remains sensitive enough to preserve micro-indicators of 'strategy hesitation' and 'cognitive load' without over-smoothing the behavioral variance.

Classic experimental platforms' fragmented approach to workflow has been replaced with a deeply linked data acquisition node to an evaluation engine based on processes.

The real-time encapsulation of bidding sequence, decision-time (delay) and negotiation information (text) from every bid round has been accomplished by the concurrent creation of standard time-step vectors (having six samples per second) of all of the data collected while each round is occurring. As these vectors are accumulated, they are pushed to a cognitive tracking model in real-time, creating an absence of lag time between the collection of behavior(s) and the teacher receiving feedback on those behaviors; this allows teachers to establish a "closed-loop" cycle of utilizing operations (as data), data (as evaluation) and evaluations (as intervention) within their teaching practice. By supporting the real-time analysis of multimodal data streams, the system can accurately identify the cognitive-load level of students at each node (for example, you can indicate when a student is experiencing excessive cognitive load) at their particular node (i.e., key price change, observing an atypical indicator), providing a dynamic learning dashboard for the teacher and an extremely rich temporal data set to facilitate the ongoing mapping of cognitive states and individualized micro-intervention(s) for each student using Attention-BiLSTM technique.

III. CORE ALGORITHMS OF THE AI-EETES DYNAMIC EVALUATION MODEL

A. FEATURE ENGINEERING OF MULTIMODAL BEHAVIORAL DATA

Although the original interaction logs contain a high degree of fidelity, they have several issues that make them unsuitable for direct input into deep-sequence models, including their heterogeneous sampling frequencies, inconsistent dimensions, and semantically fragmented nature. The central function of the feature engineering step is to reduce the dimensionality of the physical-level operation trajectories and recreate them into standardized time-series tensors with an economic meaning. In the price game scenario, the most important behavioral feature to identify is the "relative price adjustment magnitude," which quantifies how aggressively each participant has been probing in a sequential manner. The computation formula for this feature is defined as follows:

$$M_d = \frac{|p_k - p_{k-1}|}{p_{k-1}} \tag{2}$$

p_k represents the valid bid value submitted in the k-th independent trading round, and p_{k-1} is the benchmark bid of that individual in the immediately preceding round. Simultaneously, the system extracts the "decision latency" feature by calculating the timestamp difference to assess information processing resistance and strategy hesitation.

$$D_l = t_{\text{act}} - t_{\text{show}} \tag{3}$$

In the formula, t_{show} is the timestamp when the specific market state and parameters are fully rendered on the front-end interface, and t_{act} is the accurate timestamp of the confirmation decision instruction being

issued. For the text modality, a pre-trained Natural Language Processing (NLP) part-of-speech tagging and sentiment analysis module is invoked to extract the negotiation polarity score from the chat log. Through the above feature mapping, the original fragmented logs are reduced in dimensionality and structurally reorganized. Within each discrete time step t, various features are concatenated and encapsulated into an N-dimensional feature vector $\mathbf{v}_t$. Throughout the entire experimental interaction period, all dynamic sequences of the individual are finally concatenated and output as a standard high-dimensional time series feature matrix:

$$\mathbf{V} = [\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_T] \tag{4}$$

In matrix $\mathbf{V}$, T represents the total number of time steps. This feature matrix $\mathbf{V}$ is directly used as the standard input tensor for the downstream depth tracking model.

B. CORE ARCHITECTURE: ATTENTION-BILSTM

C. COGNITIVE STATE MAPPING ALGORITHM BASED ON ATTENTION-BILSTM

The core computational layer of the algorithm first feeds it into a bidirectional long short-term memory network for temporal dependency modeling [27-28]. Economic game decision-making has significant forward and backward causal correlation characteristics, and unidirectional sequence models are difficult to use subsequent market state feedback to correct the cognitive evaluation of the current node. At any discrete time step t, the forward and backward networks independently process the input behavioral feature vector $\mathbf{v}_t$, extracting the historical evolution trajectory and future trend implicit information respectively, and generating the corresponding hidden state vectors:

$$\overrightarrow{\mathbf{h}}_t = \text{LSTM}_{fw}(\mathbf{v}_t, \overrightarrow{\mathbf{h}}_{t-1}) \tag{5}$$

$$\overleftarrow{\mathbf{h}}_t = \text{LSTM}_{bw}(\mathbf{v}_t, \overleftarrow{\mathbf{h}}_{t+1}) \tag{6}$$

In the equation, LSTM_{fw} and LSTM_{bw} represent the forward and backward gated recurrent units, respectively. Subsequently, the system performs a concatenation operation on the bidirectional hidden states along the channel dimension to construct a joint hidden state vector that fully captures the temporal interaction features:

$$\mathbf{h}_t = [\overrightarrow{\mathbf{h}}_t \oplus \overleftarrow{\mathbf{h}}_t] \tag{7}$$

The vector $\mathbf{h}_t$ serves as the underlying representation, accurately characterizing the action sequence attributes of an individual at a specific game node.

Given the significant heterogeneity in the contribution of different operational behaviors to higher-order economic cognitive inference in the experimental interaction scenario,

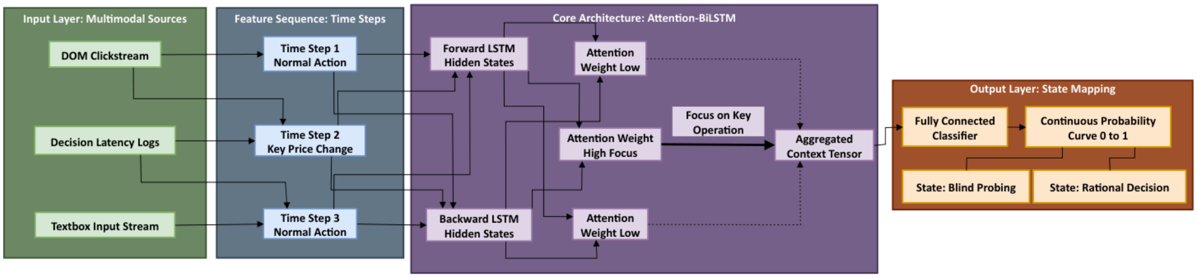

FIGURE 1. Structure diagram of AI-EETES multimodal cognitive state tracking mapping network

Note: LSTM = Long Short-Term Memory. Figure 1 shows the structure of the AI-EETES multimodal cognitive state tracking mapping network: the left side is the multimodal input source, the middle part shows the sequential representation of features at time steps, where key economic decision-making behaviors (such as the key price change in Time Step 2) are marked as high-weight attention objects, while routine operations are used as background sequences; these data streams then enter the core architecture (purple area), focusing on key operations and aggregating context tensors through an attention mechanism, and finally mapping them in the output layer (brown area on the right) as a continuous cognitive state probability curve (0-1) of the student from "blindly exploring" to "rationally making decisions", realizing a seamless connection from explicit behavioral sequences to implicit cognitive computation.

the algorithm embeds an attention mechanism at the BiLSTM output to achieve feature recalibration. This mechanism dynamically quantifies the importance weight α_t of the action sequence at each time step through a learnable global context vector $\mathbf{u}_c$. The specific calculation process is as follows:

$$\mathbf{u}_t = \tanh(\mathbf{W}_a \mathbf{h}_t + \mathbf{b}_a) \tag{8}$$

$$\alpha_t = \frac{\exp(\mathbf{u}_t^\top \mathbf{u}_c)}{\sum_{j=1}^{T} \exp(\mathbf{u}_j^\top \mathbf{u}_c)} \tag{9}$$

$\mathbf{W}_a$ and $\mathbf{b}_a$ are the weight matrices and bias terms of the attention-based multilayer perceptron, $\mathbf{u}_t$ is the intermediate state representation after nonlinear mapping, and T is the total time step covered by the evaluation window. The system performs weighted aggregation of the global hidden states based on the weight distribution α_t to generate a contextual representation vector highly focused on the core economic decision nodes.

$$\mathbf{s} = \sum_{t=1}^{T} \alpha_t \mathbf{h}_t \tag{10}$$

The process of aggregation reduces noise in the data created from typical web navigation and invalid clicks and allows for clearer inference resources to support identifying when to make important price changes or when negotiating appropriately which aligns to the specification of "allocated attention to weight" and "focused operations" in Figure 1.

The combined representation vector $\mathbf{s}$ is passed into two parallel branches on a fully connected network to support multi-task inference in the final output stage of state mapping and quantization. The discrete classification branch will apply Softmax activation to generate the probability distribution vector $\mathbf{p}$ of a student's cognitive state as of now linear rationality, irrational herd behavior, boundary testing:

$$\mathbf{p} = \mathrm{softmax}(\mathbf{W}_o \mathbf{s} + \mathbf{b}_o) \tag{11}$$

$\mathbf{W}_o$ and $\mathbf{b}_o$ denote the matrix and bias respectively of the classification layer's parameters. The continuous regression path provides scalar values from high dimensional and abstract feature sets and produces an evaluation score E_t that is dynamic and serves as a procedural evaluation for what's occurring at the front end of an application or environment based on real time:

$$E_t = \mathrm{sigmoid}(\mathbf{w}_s^\top \mathbf{s} + b_s) \times 100 \tag{12}$$

The weight vector ($\mathbf{w}_s$) and scalar bias (b_s) of the regression layer are referred to in the formula. This dual path of computation ensures two important results: mathematical precision in mapping between discrete behavioral sequences and continuous cognitive states; and improved interpretability of economic decisions through a differentiable weight allocation mechanism. Unlike traditional 'black box' models, the Attention-BiLSTM architecture allows researchers to visualize which specific behavioral features (such as sudden latency spikes) most influence the cognitive state score, effectively bridging the gap between deep learning outputs and classical economic decision theory.

The algorithm black box obscures how individual criteria conform to its defined operation therefore public assessment remains centered around scores. Without semantic alignment to that of economic theory most forms of traditional deep neural network (DNN) evaluate performance strictly through obtaining subjective or otherwise disconnected score results. Thus, the architecture of the model being developed allows for the identification of the appropriate level of attention weights while in design phase along with a method of transforming these values into an evaluation method to aid instructors and students in obtaining meaningful teaching and learning experiences.

The mechanism used for attribution will consist of three parts: The first part of the mechanism is the local ization of local maxima in the behavioral intervention; Extraction of key local maxima via trainable dynamic thresholding of significant values along the weight sequence; Identification of those key time points for each intervention period that had the largest impact on determining the student's cognitive state.

The key identification methods for the economic decision bias mapping segment and corresponding behavior feature engineering layer raw logs use a reverse index. If the high-weight interval matches the definition for 'aggressively adjusting prices that deviate from the Walrasian equilibrium', the system attributes this as representative of an 'over-confident bias and risk-seeking tendency'. If a high-weighted interval reflects the definition for 'a significant increase in the delay associated with a decision coupled with a significant degree of tentative under-reporting', the system will classify that interval as representative of 'information processing resistance and risk aversion tendency.' By translating mathematical weights into corresponding micro-economic behavioral labels, the economic decision making process of the model, and the classical decision theory of the model may be jointly referenced in accordance with each other.

Using a natural language template engine, the system maps the cognitive state probability distribution (p) directly to a dynamic prompt library. For instance, when the probability of 'irrational herd behavior' exceeds the threshold θ_{err}, the backend triggers a specific LLM context window. This window combines the student's recent high-weight behavioral logs with an economic reasoning template to generate a readable structured diagnostic feedback report i.e. "Your weight was the highest in the last three rounds (5-7) based on two consecutive high-frequency attempts to deviate from the equilibrium price (aggressively). As such, it is advisable to reassess your pricing range along with your marginal cost basis."

IV. EXPERIMENTAL ENVIRONMENT DEPLOYMENT AND IMPLEMENTATION PLAN

A. EXPERIMENTAL DESIGN FOR THE "TWO-WAY AUCTION AND OLIGOPOLY" SCENARIO

The two-way auction system uses an order book that is continuously matched by limiting buy and sell orders submitted using the price-time priority method. All order submissions will be performed by both LLM agents and human traders using different strategies asynchronously within one virtual market. The experimental design includes a real time model for the measurement of each participant's ability to make micro-decisions efficiently and the ability to settle the market in a changing price game.

$$S_i = \sum_{j \in \Omega_B} (V_{ij} - P_j) + \sum_{k \in \Omega_S} (P_k - C_{ik}) \tag{13}$$

S_i represents the actual total profit captured by participant i within a specific trading period; Ω_B and Ω_S represent the sets of buy and sell orders successfully executed by the individual, respectively; V_{ij} is the private valuation allocated to buy order j by the system; C_{ik} is the private marginal cost of sell order k; and P_j and P_k correspond to their respective final matching transaction prices. By monitoring the movement of transaction prices moving in a direction to reach the theoretical Walrasian equilibrium, this system will allow for participant's bidding strategies to be converted into measurable economic returns in conjunction to providing a solution to the issue of not having an objective measure to evaluate individual pricing discovery ability in standard experimental setups. Additionally, the system uses discrete time segments to represent each period of time and synchronizes the depth of the order book with each participant's bid sequence and converts those mappings into feature vectors thus, allowing for high frequency trading activity to be captured completely. Participants in this module are involved in a multi-period Cournot competition simulation of several noncooperative games, where the final object is to purchase inputs or output at prices established under conditions of imperfect competition, as described in [29-30]. The inverse market demand curve and each firm's

production parameters are provided to the manufacturers prior to the start of the game. These parameters are also affected by predetermined initial values for total capital, B_0, and the acceptable range of values for each commodity, $[V_{\min}, V_{\max}]$, as established in Section 4.2. Each participant must independently decide on their optimal level of production for a limited period and take into account how the combined production outputs of the other applicants will affect their total profits. The formal definition of the profits of each participant is as follows:

$$\pi_i = (\alpha - \beta(q_i + Q_{-i}))q_i - \gamma_i q_i \tag{14}$$

π_i represents participant i's current economic profit; α and β define the intercept and slope of the market linear demand function, respectively; q_i is the output decision submitted by the participant; Q_{-i} represents the aggregate expected output of peer competitors (i.e., other LLM agents in the system); and γ_i represents the constant marginal production cost assigned to the participant. By calculating the deviation between the actual submitted q_i and the theoretical Nash equilibrium output, the system can quantify and refine the depth of the participant's strategic reasoning, thereby solving the objective observation problem of higher-order "optimal response" and "collusive betrayal" thinking.

The state transition sequences found in both gaming environments architecture are tightly related to the seamless data acquisition engine for the underlying AI-EETES algorithm architecture. All timestamped multi modal transactions (i.e., high frequency submitting price limit order and production parameter releases at discrete times) are mapped into state action pairs for each state-action transition using the Markov Decision Process framework. These results will automatically extract price fluctuation trajectories, decision delay sequences, and negotiation textual features. These features will be aggregated and interpolated over a fixed period of time to develop a standard high-dimensional time series feature matrix. The manner of designing this experiment allows for the variables transferred from the business logic layer to the Attention-BiLSTM Network to be converted automatically into temporal input tensors and the evaluation models to be evaluated with robust results, while also ensuring the accuracy of the execution of the data source.

B. A/B TEST GROUPING AND CONTROL VARIABLE SETTING

To reduce any evaluation bias arising from heterogeneity within the sample, the assignment of participants into treatment groups did not follow the traditional simple random sampling method but instead utilized a

stratified method based on the matching of covariates between subjects. The subjects' prior academic performance and the scores from the Baseline Risk Preference Instrument are used to produce a multidimensional covariate vector by using a computer program designed to randomly place subjects in either treatment or control groups. Using an algorithm to map the subjects orthogonally, participants are assigned to the Experimental Group (who received AI-EETES dynamic evaluation and intervention) or to the Control Group (where subjects will receive only the usual post hoc static settlement). The covariate mean values' standard difference must meet stringent standards of equivalence for each covariate dimension between the two groups:

$$\frac{|\mu_{E,d} - \mu_{C,d}|}{\sqrt{(\sigma_{E,d}^2 + \sigma_{C,d}^2)/2}} \le \epsilon \tag{15}$$

In the formula, $\mu_{E,d}$ and $\mu_{C,d}$ represent the sample means of the experimental group and the control group on the d-th covariate, respectively; $\sigma_{E,d}^2$ and $\sigma_{C,d}^2$ correspond to the sample variances of the two groups; ϵ is the minimum tolerance threshold preset by the system. By applying this procedure prior to entering the gaming environment of the two groups of subject's games (to measure the economic basic cognition and decision-making personality type), as well as establishing a mathematical basis to support antecedent causal inference, it has provided an empirical and objective assessment of the economic basic cognition and decision-making personality distribution or isomorphism of the two groups.

Regarding the isolation control of the external intervention environment, the system's core engine executes a rigid alignment protocol on all non-processed variables. The uniqueness of the independent variables eliminates the interference caused by differences in market parameters. The key environment configurations are shown in Table 1. The system locks in the initial funding budget B_0, the commodity inventory limit $I_{\max}$, and the countdown extreme value τ_{limit} for a single round of decision-making. The concurrent scale N of market participants is set as a fixed threshold to control the computational load while ensuring market liquidity; the commodity basic value interval $[V_{\min}, V_{\max}]$ follows a uniform distribution, serving as the generation boundary for the random walk of private valuation and production costs; the number of discrete game rounds K constitutes the stage upper limit for forced convergence of the scenario, triggering the liquidation logic upon reaching it. Meanwhile, the information disclosure depth of the front-end interface is uniformly controlled by the mask matrix M_{mask}, ensuring that the two groups of subjects are completely consistent in the visible

level of the order book and the rendering range of the historical price chart, thereby eliminating the interactive noise caused by information asymmetry.

The system gateway in the experimental group implements differentiated routing strategies, based on user identity, to enable the specific intervention implementation. The control group will automatically redirect their interactions to the basic matching sandbox and only get issued a static profit settlement statement after the total window of time closes. The experimental group will establish a full duplex real time communication link with the backend Attention–BiLSTM matrix when conducting their sessions. The acquisition engine keeps track of behaviours at a very high sampling rate f_s, and continuously monitors latency time $\tau_{\max}$, to keep the temporal synchronisation of the data being transmitted. The model truncates the historical input vector at each calculation based on the time series truncation length L, thus ensuring that the time window used for the calculation coincides with the decay time period of short term memory (i.e. short term memory decays after L). When the output of the algorithm (i.e. the real-time comprehensive evaluation score E_t) falls continuously below the cognitive abnormality trigger threshold θ_{err}, the front-end interface interrupts the next instruction for a bid submission prior to that instruction being executed, and presents a prompt containing a targeted strategy correction that is dynamically generated using LLM's within microseconds. This mechanism completely establishes a closed loop of control logic from "unobtrusive recording - real-time measurement - intervention blocking," achieving automated coupling of teaching process evaluation and dynamic correction.

V. MODEL PERFORMANCE VALIDATION AND TEACHING EVALUATION ANALYSIS

A. ALGORITHM MODEL PERFORMANCE AND REAL-TIME EVALUATION ANALYSIS

In order to validate the ability of the Attention-BiLSTM network to compute accurately for tasks of mapping discrete behavior sequences, a five-fold cross-validation mechanism is employed at the performance evaluation phase for accurate and unbiased measurements on the isolated test dataset. For the purpose of conducting unbiased horizontal comparisons of model performance, the performance evaluation methodology selected Support Vector Machine (SVM) and a standard unidirectional Long Short Term Memory (LSTM) network as baseline models. The macro-average F1 score is thus selected as the principle quantitative metric, so as to remove measurement bias associated with differences in the size of the samples from the economic cognitive state

TABLE 1. Key Parameters and Data Dimensions of the Experimental Environment

Parameter Category	Parameter Name	Symbol/Unit	Configuration/Description
(1) Experimental Settings	Concurrent Market Participants	N (persons)	40 (30 human subjects + 10 LLM agents per session)
	Initial Capital Endowment	B_0 (virtual coins)	2,000
	Inventory limit of goods	$I_{\max}$	10
	Countdown extremum for single turn decision-making	τ_{limit}(s)	45
	Commodity Value Range	$[V_{\min}, V_{\max}]$	[20, 120] (uniform distribution)
	Discrete Trading Rounds	K (rounds)	15
(2) Data Collection Dimensions	Global Event Sampling Rate	f_s (Hz)	50
	Max Transmission Latency	$\tau_{\max}$ (ms)	100
	Sequence Truncation Length	L (steps)	50
	Cognitive Anomaly Threshold	θ_{err}(scalar)	0.35

Notes: The Participant Scale provides enough liquidity for the market, but yet is computationally practicable. The Sampling Rate of 50Hz provides a good balance between server load and granularity. If the probability that a participant displays an "irrational behaviour" exceeds 35%, then the system will trigger intervention at the threshold $\theta_{err} = 0.35$.

labels being examined. The testing process revealed that the SVM model, lacking a time-series modeling mechanism, only achieved an F1 score of 78.4%; while the unidirectional LSTM model could capture historical unidirectional state information, its accuracy tended to saturate and stagnate at 85.1%. In contrast, the pretrained Attention-BiLSTM model, through joint inference of bidirectional temporal dependencies and high weighting of key price modification actions, successfully eliminated noise interference caused by irrelevant page clicks. As shown in Figure 2(a), the model's final prediction performance on the test set reached an F1 score of 91.2%, solving the problem of high cognitive misjudgment rate in traditional evaluation mechanisms from the algorithmic source.

In the real-time performance testing at the engineering application level, the system underwent rigorous end-to-end inference latency stress testing for the high-concurrency interactive streams generated in the dynamic game environment. The computational domain of the latency measure is defined as: from the physical timestamp of receiving the preprocessed aggregated feature tensor $\mathbf{V}$ from the server to the termination timestamp of the network regression branch outputting the dynamic evaluation score E_t and sending it back to the gateway. Under full-load concurrent testing conditions, the computational resource scheduling module effectively compressed the transmission time of hidden state through tensor parallel operation optimization. As shown in Figure 2(b), the statistical distribution characteristics show that the 95th percentile latency of the evaluation model is stably maintained at around 45.6ms, and the overall system's extreme inference latency is locked within the 50ms threshold, breaking the lag barrier of traditional batch processing scoring. While ensuring the high complexity of the evaluation algorithm, it meets the stringent engineering requirements of seamless rendering of the front-end transaction interface and instant triggering of personalized intervention commands.

B. EVOLUTION OF DECISION-MAKING RATIONALITY CHARACTERISTICS AND ANALYSIS OF TEACHING INTERVENTION EFFECTS

To quantify the substantial impact of immediate feedback intervention on the evolution of subjects' higher-order economic thinking, the data analysis module extracted the probability component of the "absolute rationality" state output in real time by the classification branch of the Attention-BiLSTM network, and aggregated it to generate a dynamic decision rationality index with discrete game rounds as the time axis. In the stage of natural learning effect stripped of the time dimension, this step introduced repeated measures ANOVA to test the significance of the temporal trajectories of the two groups of subjects within the global interaction cycle. The tracking data shown in Figure 3 illustrates two separate paths of evolution over time. In the prior round of initial creation (approximately the first 8 rounds) the two groups had a baseline decision rate of the same in a very low random bid environment; however, as market complexity increased, the group of experimenters using the AI-EETES engine (solid blue line) when given prompting to make high-frequency experimental strategy corrections (for example, see the intervention green triangles) showed a rapid increase in their decision rate through logarithmic step function development and are able to escape the quagmire of making blind bids. The global average decision rate of the experimenters also steadily increased and reached a level greater than 0.85 whereas the group relying on traditional static settlement reports (dashed red line) continued to be randomly oscillating between 0.20 and 0.60 and thus did not develop a consistent pattern of rational decision-making. The difference between the two groups (0.278) in global average decision rate indicates the impact of the non-intrusive intervention mechanism to overcome cognitive bias and restore rationality to economics through computational confirmation of this measurement.

In order to assess the impact of study as it relates to the herd mentality prevalent in oligopolistic and multitiered market scenarios, the analysis will also utilize the system-level bullwhip effect index to determine how robust deci-

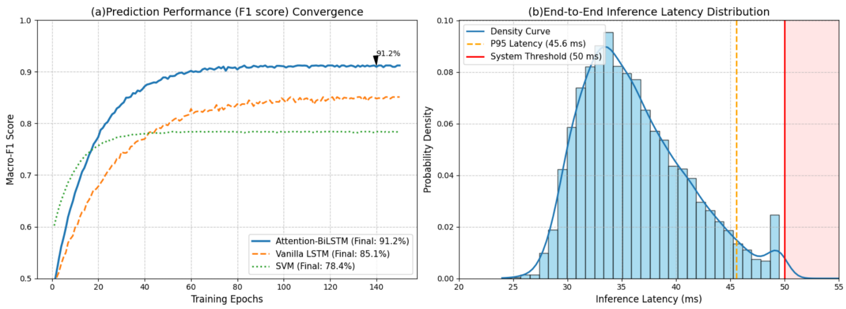

FIGURE 2. Comparison of model convergence curves and prediction performance (F1-score), and end-to-end inference delay distribution.Figure2(a)Convergence of prediction performance (F1 score),Figure2(b)Distribution of end-to-end inference delay

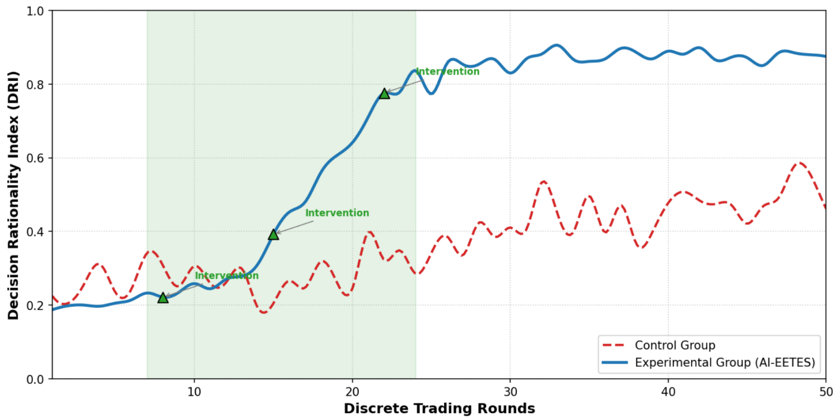

FIGURE 3. Comparison of the evolution trajectory of the Decision Rationality Index (DRI) between the experimental group and the control group

sions have been made by each individual and how much external interference has affected them. The system-level index is based upon the ratio of the time series variance σ_q^2 in dynamic output decision making by the subjects to the variance σ_D^2 in real market demand statistics:

$$BE = \frac{\sigma_q^2}{\sigma_D^2} \tag{16}$$

Based on a review of the system interaction logs, the amplified fluctuation evolution process is identified. Specifically, participants in the control group tended to overreact to the demand shock of an exploratory nature from an LLM agent, resulting in large increases in their mean BE (Behavioral Evidence) and major distortions in order placement within the supply chain and increased inventories. In contrast, when participants in the experimental group exhibited sequential behaviour that cross-exceeded the pre-defined cognitive malfunction threshold (θ_{err}), the system's directional (LLM corrective) command delivered via the microsecond response time of the system gateway effectively interrupted the positive feedback loop for "price fluctuation - panic response." A homogeneity of variance test established that the experimental group's mean BE is significantly lower than that of the control group, nearing the theoretically optimal level of absolute smoothness. This indicates that the experimental group's mean BE is closer to an optimal state, shows that the model is able to accurately track micro-level thought processes, and provides strong evidence that the model's intervention logic can decrease the level of market impulsiveness and systemic irrational price fluctuations at the macro-level.

C. CONSISTENCY TEST OF DYNAMIC PROCESS EVALUATION AND COMPREHENSIVE COMPETENCY

To evaluate the reliability and educational usefulness of the Computer Based automated system for quantitative evaluation, the trial evaluation phase included a comparative analysis between the AI-EETES machine generated scores and traditional education Evaluation Measurement criteria. The AI-EETES system extracted the time-ordered scores for the continuous dynamic evaluation output by each player's Attention-BiLSTM (Bidirectional Long Short Term Memory) network, and applied time-decay weighting, integrating all players' transient scores E_t into a global index of AI-EETES process evaluation, denoted as S_{AI}, through the use of an aggregation mechanism.

$$S_{AI} = \sum_{t=1}^{T} w_t E_t \tag{17}$$

w_t represents a weighting coefficient that diminishes as t increases (with each subsequent time step). This ensures that all previously made decisions in the evaluation process are assessed using their most relevant expertise. Furthermore, all subjects will have a standardised test score; S_{final}, representing their micro/ macroeconomic capability. The score will be determined through a combination of closed-book written examination and complex economic case analysis. The final assessment score will serve as an external reference for each student's true professional competence and will also enable testing for the reliability of the machine's underlying feature mapping matrix.

Pearson correlation analysis and univariate linear regression analysis are introduced by the consistency measure test. The covariance ratio between the comprehensive global process evaluation index and the ultimate score standardised across the total sample domain are computed to respond to the issue of the impersonal machine (to score) impersonalizing the detachment of actual teaching and assessment objectives. The statistical inference results clearly indicate an extremely significant positive correlation exists between the independent source of the evaluation vectors.

As shown in Figure 4, the scatter plots cluster tightly around the regression axis. The Pearson correlation coefficient (Pearson r) between the system's output AI-EETES global comprehensive evaluation score and the standardized score of the final comprehensive professional ability is as high as 0.88, with a significance level of $p < 0.001$. According to mathematical evidence, a deep or hierarchical network using attention weighting reduces "random interaction" noise present in front-end operational logs. The high-order "economic decisionmaking" characteristics are isomorphic to the ability dimensions that comprise a traditional assessment in the field of education.

Regression statistics further quantify and support the stability of this relationship. The linear regression model's formal definition, based on ordinary least squares (OLS) fitting, is:

$$S_{\text{final}} = \beta_0 + \beta_1 S_{AI} + \epsilon \tag{18}$$

The coefficients associated with the formula (β_0 (the y intercept) and β_1 (the slope of the regression line)) and the random error term (ϵ) indicate that the fitted values of the observed points are very close to both sides of the regression axis; residuals also passed both the Shapiro-Wilk test for normality and the White test for heteroscedasticity, which provides further evidence that there is no systematic over- or under-estimating bias in the evaluation model when comparing samples from various ability tiers. Thus, the findings from this consistency examination completely confirm that the cross-modal inference chain of "behavioural trajectory from a given set of actions - dynamic cognitive state from a range of reactions to those actions - final level of competency achieved through the completion of that task" is established; therefore, the AI-EETES model can replace traditional post-hoc, subjective experiential reports as a valid measure of real-time cumulative scores and can create a definitive closed loop of logic for both process evaluation and summative evaluations in the fields of engineering.

D. EMPIRICAL ANALYSIS OF INTERPRETABLE DIAGNOSTIC FEEDBACK

This part focuses on showing measurable evaluation data regarding the comprehensibility and acceptance of interpretable diagnostic feedback, allowing direct evidence of both the cognitive internalization effect and the tool acceptance of the algorithmic attribution mechanism used in instructional activities. To systematically examine the cognitive internalization effect of interpretable diagnostic feedback on students, a specific questionnaire assessment is conducted on the 30 participants in the experimental group after the experiment. The

questionnaire used a 5-point Likert scale (1 = completely disagree, 5 = completely agree), containing three dimensions: (1) comprehensibility of feedback (3 items, such as "The attribution description in the diagnostic report clearly corresponds to my specific operations in the experiment"); (2) guidance for strategy correction (3 items, such as "The feedback suggestions helped me clarify the optimization direction for the next round of pricing"); (3) trust in the system (2 items, such as "I am willing to refer to the diagnostic report generated by AI to adjust my decision-making strategy"). This study was approved by the Ethics Committee of Anhui Sanlian University, and was performed in accordance with the principles of the Declaration of Helsinki. All eligible participants signed an informed consent form.

Statistical results show (as shown in Table 2) that the average scores for the three dimensions are 4.32±0.61, 4.18±0.73, and 3.95±0.82, respectively (one-sample t-test, $p < 0.001$).

TABLE 2. Statistical results of the interpretable diagnostic feedback acceptance questionnaire

Dimension	Item Example	Mean ± SD	t-value	p-value
Comprehensibility	"The attribution description can clearly correspond to my specific actions"	4.32 ± 0.61	12.47	< 0.001
Guidance Utility	"Feedback and suggestions have clarified the optimization direction"	4.18 ± 0.73	10.83	< 0.001
System Trust	"I am willing to refer to AI diagnostic adjustment strategies"	3.95 ± 0.82	8.16	< 0.001

As shown in Table 2, the "understandability" dimension scored the highest, indicating that the attribution mechanism of mapping Attention weights to economic behavior labels effectively reduced the cognitive threshold of the algorithm output; "strategy correction guidance" is the second highest, reflecting that diagnostic feedback has strong teaching transfer value; "system trust" is slightly lower but still in the positive range, suggesting that some students still hold a cautious attitude towards AI intervention, which is consistent with

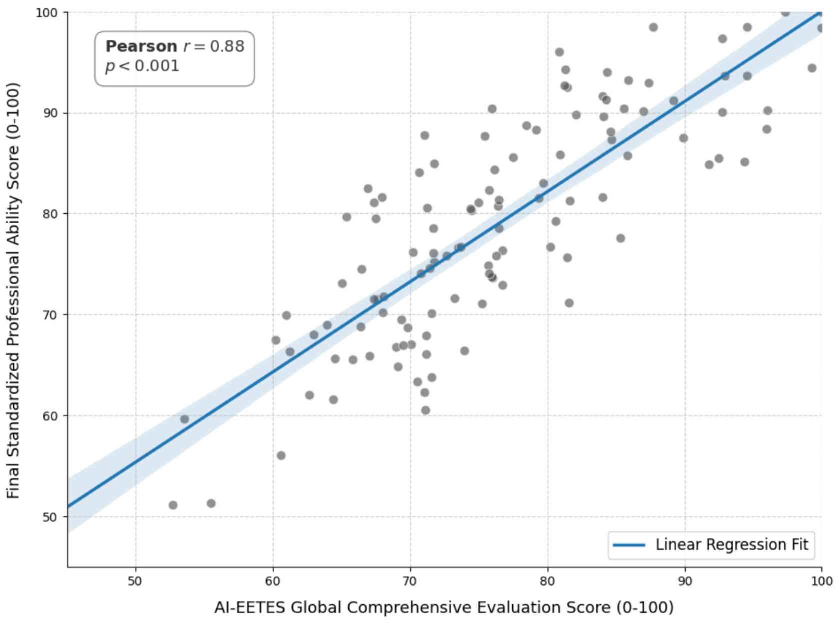

FIGURE 4. Scatter plot showing the correlation between dynamic evaluation scores and final exam scores

the classic path of "perceived usefulness" to "behavioral intention" in the educational technology acceptance model.

VI. CONCLUSION

This paper addresses the deep-seated contradiction between "abundant data and lagging evaluation" in traditional economics experimental teaching, proposing and constructing an AI-driven dynamic evaluation model that integrates large language models and deep knowledge tracking technology. In terms of implementation, this research first reconstructs a highly realistic macro- and microeconomic game scenario by injecting an

LLM agent with heterogeneous risk preferences; secondly, it extracts the time-series features of students' multimodal interactive behaviors throughout the entire lifecycle using a seamless data acquisition engine; finally, it innovatively introduces an Attention-BiLSTM network to accurately map massive discrete operational trajectories into continuous economic cognitive states, thus establishing an automated evaluation link from low-level data mining to high-order thinking diagnosis. Empirical evidence supports the model's dual use (engineering application & empowerment) for engineering & empowering classes. The model has an F1 score of 91.2% predicting cognitive states (end-to-end inference latency of <50 milliseconds), effectively overcoming the technical challenge associated with providing real-time personalized feedback. In teaching, the experimental group receiving real-time AI intervention exhibited a 0.278 increase in the Decision Rationality Index (DRI) over the traditional group, markedly diminishing bullwhip effect due to herd mentality. Moreover, the Pearson correlation coefficient between dynamic evaluation score and final professional competence score is 0.88, providing educational validity and reliability based upon this process-based assessment method. Overall, this research has effectively addressed the technical issues associated with dynamic measurement of cognition in dynamic game-theoretic environments and has also provided significant theoretical basis and engineering methodologies for innovative educational paradigms and reconstruction of evaluation systems in economics in the digital age.

AUTHOR CONTRIBUTIONS

Conceptualization –Min Zhou and Xu Feng; methodology – Min Zhou and Xu Feng; investigation – Min Zhou and Xu Feng; writing-original draft preparation – Min Zhou and Xu Feng. All authors have read and agreed to the published version of the manuscript.

CONFLICTS OF INTEREST

The authors declare no conflict of interest.

FUNDING

This work was supported by the "Four New" Research and Reform Practice Project (25zlgc002)

ACKNOWLEDGEMENTS

Not applicable.

REFERENCES

[1] L. Mihai, L.-G. Manescu, L. Vasilescu, et al., "A systematic analysis of new approaches to digital economic education based on the use of AI technologies," Amfiteatru Economic, vol. 26, no. 65, pp. 201-219, 2024, doi: 10.24818/EA/2024/65/201.

[2] O. Glazunova, T. Saiapina, V. Korolchuk, et al., "Formation Digital Intelligence of a Modern Economist: A Competence Approach," INSTICC, vol. 1, no. 1, pp. 432-447, 2023, doi: 10.5220/0012065100003431.

[3] M. Kesse, "Designing a digital interactive simulation for teaching business analytics, strategy, and economics," The Journal of Applied Business and Economics, vol. 26, no. 5, pp. 122-138, 2024, doi: 10.33423/jabe.v26i5.7339.

[4] J. Collins A and S. Etemadidavan, “Interactive agent-based simulation for experimentation: a case study with cooperative game theory,” Modelling, vol. 2, no. 4, pp. 425-447, 2021, doi: 10.3390/modelling2040023.

[5] A. Calinescu, “The Grading Process in System 1 and System 2 of Thinking: A Behavioral Economic Approach to Evaluation,” Integrated Journal for Research in Arts and Humanities, vol. 3, no. 6, pp. 105-122, 2023, doi: 10.55544/ijrah.3.6.12.

[6] J. Liu and Y. Wang, “Static and dynamic evaluations of college teaching quality,” International Journal of Emerging Technologies in Learning (iJET), vol. 17, no. 2, pp. 114-127, 2022, doi: 10.3991/ijet.v17i02.29005.

[7] M. G. Hahn, S. M. B. Navarro, F. Valentin L. D. L., et al., “A systematic review of the effects of automatic scoring and automatic feedback in educational settings,” IEEE Access, vol. 9, no. 1, pp. 108190-108198, 2021, doi: 10.1109/ACCESS.2021.3100890.

[8] D. Danz, L. Vesterlund, and J. Wilson A, “Evaluating behavioral incentive compatibility: Insights from experiments,” Journal of Economic Perspectives, vol. 38, no. 4, pp. 131-154, 2024, doi: 10.1257/jep.38.4.131.

[9] M. Yang and D. Li, “Personalized Intelligent Recommendation Model for Educational Games Based on Data Mining,” Information Technology and Control, vol. 54, no. 1, pp. 64-83, 2025, doi: 10.5755/j01.itc.54.1.37088.

[10] M. Liu, C. Li, Z. Pan, et al., “Mining big data to help make informed decisions for designing effective digital educational games,” Interactive Learning Environments, vol. 31, no. 5, pp. 2562-2582, 2023, doi: 10.1080/10494820.2019.1639061.

[11] M. Wang, D. Zhang, J. Zhu, et al., “Effects of incorporating a large language model-based adaptive mechanism into contextual games on students’ academic performance, flow experience, cognitive load and behavioral patterns,” Journal of Educational Computing Research, vol. 63, no. 3, pp. 662-694, 2025, doi: 10.1177/07356331251321719.

[12] E. Huber S, K. Kiili, S. Nebel, et al., “Leveraging the potential of large language models in education through playful and game-based learning,” Educational Psychology Review, vol. 36, no. 1, pp. 25-45, 2024, doi: 10.1007/s10648-024-09868-z.

[13] Y. Zheng, Z. Xu, and A. Xiao, “Deep learning in economics: a systematic and critical review,” Artificial Intelligence Review, vol. 56, no. 9, pp. 9497-9539, 2023, doi: 10.1007/s10462-022-10272-8.

[14] Y.-T. Shiao, C.-H. Chen, K.-F. Wu, et al., “Reducing dropout rate through a deep learning model for sustainable education: long-term tracking of learning outcomes of an undergraduate cohort from 2018 to 2021,” Smart Learning Environments, vol. 10, no. 1, pp. 55-71, 2023, doi: 10.1186/s40561-023-00274-6.

[15] M. Fernandez J, A. Yetter E, and K. Holder, “What do economic education scholars study? Insights from machine learning,” The Journal of economic educaTion, vol. 52, no. 2, pp. 156-172, 2021, doi: 10.1080/00220485.2021.1887027.

[16] M. Dell, “Deep learning for economists,” Journal of Economic Literature, vol. 63, no. 1, pp. 5-58, 2025, doi: 10.1257/jel.20241733.

[17] M. Rivera, P. Lee, and A. Desai, “Real-Time Knowledge Tracing in Online Learning Environments Using Interpretable Transformer-Bayesian Models,” Frontiers in Interdisciplinary Educational Methodology, vol. 2, no. 3, pp. 165-178, 2025, doi: 10.71465/fiem360.

[18] R. Bernsteiner, C. Ploder, Y. Su, et al., “Design principles of real time students’ performance monitoring in simulation-based business education,” Procedia Computer Science, vol. 266, no. 1, pp. 1147-1154, 2025, doi: 10.1016/j.procs.2025.08.142.

[19] M. Rizinski and D. Trajanov, “AI Agents in Finance and Fintech: A Scientific Review of Agent-Based Systems, Applications, and Future Horizons,” Computers, Materials, & Continua, vol. 86, no. 1, pp. 1-35, 2026, doi: 10.32604/cmc.2025.069678.

[20] Y. Yu, Z. Yao, H. Li, et al., “Fincon: A synthesized llm multi-agent system with conceptual verbal reinforcement for enhanced financial decision making,” Advances in Neural Information Processing Systems, vol. 37, no. 1, pp. 137010-137045, 2024, doi: 10.52202/079017-4354.

[21] Y. Dong, F. Wu, K. Zhang, et al., “Large language model agents in finance: A survey bridging research, practice, and real-world deployment,” Findings of the Association for Computational Linguistics: EMNLP, vol. 2025, no. 1, pp. 17889-17907, 2025, doi: 10.18653/v1/2025.findings-emnlp.972.

[22] S. Ellsworth and T. Kingsley, “Accelerating Financial Intelligence via High Throughput Distributed Systems for Large Language Model Augmented Time Series Forecasting,” International Journal of Artificial Intelligence Research, vol. 1, no. 1, pp. 1-10, 2026.

[23] E. Ulitzsch, Q. He, V. Ulitzsch, et al., “Combining clickstream analyses and graph-modeled data clustering for identifying common response processes,” psychometrika, vol. 86, no. 1, pp. 190-214, 2021, doi: 10.1007/s11336-020-09743-0.

[24] Z. Wang, J. Wei, A. Aiken, et al., “LogCIoud: Fast Search of Compressed Logs on Object Storage,” Proceedings of the VLDB Endowment, vol. 18, no. 8, pp. 2362-2370, 2025, doi: 10.14778/3742728.3742733.

[25] W. Ke, Y. Zheng, Y. Li, et al., “Large language models in document intelligence: A comprehensive survey, recent advances, challenges, and future trends,” ACM Transactions on Information Systems, vol. 44, no. 1, pp. 1-64, 2025, doi: 10.1145/3768156.

[26] M. Herrera, M. Sasidharan, J. Merino, et al., “Handling irregularly sampled IoT time series to inform infrastructure asset management,” IFAC-PapersOnLine, vol. 55, no. 19, pp. 241-245, 2022, doi: 10.1016/j.ifacol.2022.09.214.

[27] C. F. Lui, Y. Liu, and M. Xie, “A supervised bidirectional long short-term memory network for data-driven dynamic soft sensor modeling,” IEEE Transactions on Instrumentation and Measurement, vol. 71, no. 1, pp. 1-13, 2022, doi: 10.1109/TIM.2022.3152856.

[28] S. Mei, X. Li, X. Liu, et al., “Hyperspectral image classification using attention-based bidirectional long short-term memory network,” IEEE Transactions on Geoscience and Remote Sensing, vol. 60, no. 1, pp. 1-12, 2021, doi: 10.1109/TGRS.2021.3102034.

[29] E. Rouskas, “Mergers, multiperiod Cournot competition, and Coasian dynamics,” Bulletin of Economic Research, vol. 75, no. 2, pp. 270-286, 2023, doi: 10.1111/boer.12348.

[30] T. Willis and G. Punzo, “Multi-layer Cournot-congestion model,” IFAC-PapersOnLine, vol. 55, no. 40, pp. 61-66, 2022, doi: 10.1016/j.ifacol.2023.01.049.