International Journal of Advanced AI Applications

Volume2,Issue12,Dec.2026

Hong Kong Dawn Clarity Press Limited

http://www.dawnclarity.press/index.php/ijaaa

Online ISSN 3104-9338

Print ISSN 3104-932X

TABLE OF CONTENTS

An Adaptive Local Trajectory Planning Method for DWA Based on the ACT-R Cognitive Mechanism

Liwen Wang , Aijuan Li , Máté Zöldy, Xibo Wang, Chuanyan Xu, Yuansheng Cheng…… (1-23)

Constructing a Generative AI Framework for Personalised Instructional Resource Development: A Quasi-Experimental Evaluation in a Micro-Lecture Design and Production Course

Guihua Gao*, Hanyan Cai, Shiru Zhu…… (24-47)

Research on Factual Hallucination Detection and Mitigation in Large Language Models Based on Retrieval-Augmented Generation

Zehua Zhou……(48-65)

Microwave-Assisted Direct Fluorination of Polymers in Vibro-Fluidized Beds: From Reaction Kinetics to Reactor Design

Yayuan Zheng, Yang Zhou……(66-90)

Recursive Adversarial Engine: A Stability-Oriented Architecture for Reliable Large Language Model Multi-Agent Systems

Jianhua Fang……(91-110)

Impressum …… (111-112)

Online ISSN: 3104-9338
Print ISSN: 3104-932X

An Adaptive Local Trajectory Planning Method for DWA Based on the ACT-R Cognitive Mechanism

Liwen Wang [1], Aijuan Li [1]*, Máté Zöldy[2], Xibo Wang[1], Chuanyan Xu[1], Yuansheng Cheng[1]

[1] School of Automotive Engineering, Shandong Jiaotong University, Jinan 250357, China;

[2] Department of Environmental Economics and Sustainability, Faculty of Economics and Social Sciences,Budapest University of Technology and Economics, M″uegyetem rkp 3-5, H-1111 Budapest, Hungary

Received: September 21, 2026

Revised: September 30, 2026

Accepted: October 7, 2026

Published online: October 20, 2026

To appear in: *International Journal of Advanced AI Applications,* Vol. 2, No. 11 (November 2026)

* Corresponding Author: Aijuan Li (liaijuan@sdjtu.edu.cn)

Abstract. To address the issues of fixed evaluation weights and the lack of historical experience utilization in the traditional Dynamic Window Approach (DWA) for autonomous vehicles in complex dynamic environments, this paper proposes an adaptive local trajectory planning method based on the ACT-R cognitive mechanism. First, the system constructs an environmental perception based on obstacle clearance, predicted risks from dynamic obstacles, and target progress, and establishes three cognitive planning strategies: Normal, Avoid, and Recover. Second, the memory activation mechanism of ACT-R is utilized to describe the historical usage characteristics of different strategies, and utility feedback learning is employed to reflect the long-term effectiveness of strategies under different environmental states. This allows the results of strategic decisions to be mapped in real time to the multi-objective evaluation function of DWA, thereby achieving adaptive and dynamic adjustment of planning weights. Finally, simulation results show that in mixed static and dynamic obstacle scenarios, while ensuring 100% collision-free arrival, this method achieves an average speed 30.17% higher than traditional DWA, reduces travel time by 18.48%, and improves trajectory smoothness by 46.34%, effectively balancing safety and travel efficiency.

Keywords: *Dynamic Window Approach (DWA); ACT-R Cognitive Mechanisms; Local Path Planning; Unmanned Ground Vehicle*

1. Introduction

With the rapid development of intelligent perception, automatic control, and artificial

intelligence technologies, autonomous vehicles—equipped with autonomous perception, environmental understanding, and motion control capabilities—are increasingly being applied in scenarios such as logistics and transportation, field inspections, agricultural operations, and autonomous operations in complex environments[1]. During actual operation, autonomous vehicles must adjust their speed and direction in real time based on the distribution of surrounding obstacles, the vehicle's own motion state, and the target location. As the core link between environmental perception and motion control, local trajectory planning directly determines the safety, efficiency, and trajectory smoothness of unmanned vehicle operation[2]. However, traditional local trajectory planning algorithms struggle to simultaneously address real-time obstacle avoidance, motion constraints, and environmental adaptability when dealing with complex obstacles[3]; therefore, targeted improvements to these algorithms are of great significance.

The Dynamic Window Approach (DWA) predicts short-term motion trajectories by sampling candidate velocities within the dynamically feasible velocity space and evaluating their corresponding trajectories under motion and safety constraints. By directly incorporating the vehicle's kinematic and dynamic limitations into the local planning process, DWA can generate executable control commands while maintaining relatively low computational complexity and good real-time performance[4]. Therefore, it has been widely applied to local trajectory planning for unmanned vehicles operating in environments with obstacles. However, conventional DWA relies heavily on predefined velocity sampling strategies and fixed evaluation-function weights, which may limit its adaptability when the surrounding obstacle distribution and risk level vary significantly. In particular, in complex obstacle environments, fixed evaluation preferences can make it difficult to simultaneously balance trajectory safety, obstacle avoidance efficiency, motion smoothness, and forward progress.

To address these limitations, previous studies have introduced improvements to DWA from several aspects, including velocity sampling, evaluation functions, attraction-repulsion mechanisms, and dynamic obstacle modeling [5]. These studies aim to improve the adaptability of local trajectory planning by incorporating more environmental information into the trajectory generation and evaluation process. For example, Xu et al. dynamically adjusted the evaluation weights according to environmental complexity and the distance between the vehicle and obstacles, thereby improving the adaptability of DWA under different obstacle conditions [6]. Gong et al. introduced fuzzy control to realize adaptive adjustment of evaluation-function weights, allowing the relative importance of different evaluation factors to vary with the current

environmental state [7]. Liu et al. further developed a fuzzy-control-based local planning method, enabling DWA to dynamically regulate the influence of different evaluation factors according to changes in the surrounding environment [8]. These studies demonstrate that introducing adaptive decision mechanisms into the DWA evaluation process can effectively alleviate the limitations associated with fixed parameters and improve its adaptability to complex obstacle environments.

In environments with dynamic obstacles, local path planning must also take into account the future motion of obstacles, rather than relying solely on the current instantaneous distance. Hossain et al. combined an improved DWA with a gap-following method to enhance the local obstacle avoidance capabilities of a mobile platform when facing both static and moving obstacles [9]; Kobayashi and Motoi improved the safety of DWA candidate trajectories by introducing a virtual manipulator to handle dynamic obstacles [10]; Yang et al. combined an improved particle swarm optimization with DWA and introduced the concept of velocity obstacles to enhance trajectory planning capabilities in dynamic obstacle environments [11]; Sarkis et al. further employed online particle swarm optimization to dynamically adjust the weights of the DWA objective function, enabling the planner to modify the trade-offs between different evaluation metrics based on environmental feedback [12].

Existing research indicates that adaptively adjusting the DWA evaluation weights can effectively enhance environmental adaptability; however, most methods rely solely on the current state to adjust parameters, failing to account for historical experience and long-term benefits. For example, methods based on fuzzy rules or online optimization tend to focus on numerical solutions, making it difficult to capture the continuous evolutionary process of “ current environment — historical experience — strategy selection — feedback learning. ” Therefore, determining how to enable a local planner to incorporate historical experience into a dependency mechanism is a key issue for further enhancing the adaptability of DWA.

Cognitive architectures offer new approaches to addressing the aforementioned issues. ACT-R (Adaptive Control of Thought-Rational), as a cognitive architecture equipped with memory, learning, and decision-making mechanisms, can simulate human behavioral choices through activation, retrieval, and utility mechanisms. In recent years, it has been widely applied in cognitive decision-making and the modeling of autonomous driving behavior. For example, Cina and Rad noted that ACT-R can provide a unified framework for vehicle environmental understanding and decision-making [13]; Oh et al. used it to establish a model for predicting driving behavior and cognitive load in takeover scenarios, validating its ability to describe the

dynamic response of cognitive states in driving decision-making [14].

The memory and learning mechanisms of the ACT-R are equally applicable to describing decision-making processes that involve historical dependencies. van der Velde et al. studied the dynamic changes in the declarative memory parameters of the ACT-R, demonstrating that memory activation can reflect the influence of experiential information as it changes over time [15]; Yang et al. incorporated motivational mechanisms into the ACT-R, using cost-benefit relationships to influence behavioral choices, thereby providing a new modeling approach for utility-based dynamic decision-making [16]; Fisher et al. further utilized system factor techniques to analyze the core architectural assumptions of ACT-R, verifying that its memory, perception, and processing mechanisms have a theoretically grounded basis that can be quantitatively analyzed [17].

To address the issues of fixed weights and high-frequency oscillations in traditional DWA strategies, this paper proposes an adaptive local path planning method for DWA based on the ACT-R cognitive mechanism. By deeply integrating the human cognitive decision-making framework with path planning, this method comprehensively enhances the adaptive obstacle avoidance capabilities, safety, and motion continuity of autonomous vehicles in environments with complex dynamic and static obstacles. The specific research content is as follows:

(1) Develop a foundational DWA trajectory planning system based on kinematic constraints and multi-objective evaluation: To address the issues of insufficient dynamic response to moving and stationary obstacles and a tendency to get stuck in local optima in complex environments — common in traditional DWA — we have established a candidate trajectory prediction mechanism based on kinematic constraints and velocity-space sampling; We introduced spatiotemporal collision prediction and spatial short-term memory revisit penalties to establish a multi-objective comprehensive evaluation function that integrates heading, clearance, speed, progress, smoothness, and safety risks; supplemented by hard constraint arbitration and a safety escape mechanism, we achieved reliable trajectory generation.

(2) Designing a strategy decision-making and adaptive weight adjustment algorithm based on the ACT-R cognitive mechanism: To overcome the limitation of traditional adaptive parameter tuning methods — namely, their lack of historical dependence — we construct an ACT-R cognitive decision-making architecture comprising three planning strategies: Normal, Avoid, and Recover. We utilize environmental risk context and memory activation to represent historical features, and employ feedback-driven utility learning to drive the continuous evolution of strategies. By leveraging cooling and hysteresis mechanisms to suppress high-

frequency oscillations in strategies, we achieve an adaptive dynamic mapping from cognitive decisions to DWA evaluation weights.

(3) Carry out algorithm simulation comparisons and ablation analyses: Under a unified test scenario, conduct horizontal comparison simulations of the traditional DWA, Fuzzy-DWA, and the Full ACT-R-DWA algorithm proposed in this paper. Through ablation experiments, separate the core modules related to memory retrieval and utility learning, quantitatively analyze the impact of each module on the algorithm performance, and clarify the intrinsic evolutionary mechanism of memory retrieval and utility learning within the ACT-R framework.

The remainder of this paper is organized as follows: Chapter 2 introduces the DWA local trajectory planning method, including the principles of DWA trajectory generation and multi-objective trajectory evaluation methods; Chapter 3 presents DWA local trajectory planning based on ACT-R cognitive decision-making, covering the overall architecture and strategy definition, the construction of the ACT-R cognitive decision-making model, cognitive decision-making based on memory retrieval and utility, and ACT-R-driven dynamic weight adjustment and feedback updates for DWA; Chapter 4 validates the effectiveness of the proposed method through simulation comparison experiments; Chapter 5 summarizes the research findings.

2. Methodology

2.1. DWA Local Path Planning Method

2.1.1. Principles of DWA Trajectory Generation

The Dynamic Window Approach (DWA) is a real-time planning method that performs local trajectory search in velocity space. Unlike methods that search for geometric paths directly in Cartesian space, the DWA samples candidate control inputs within a finite velocity space based on the autonomous vehicle's current motion state, velocity constraints, and acceleration constraints, and predicts its motion trajectory over a short time horizon.

In this paper, the state of an autonomous vehicle is defined as:

Here, x and y represent the planar position of the autonomous vehicle, θ represents the

$$\mathbf{S} = [x, y, \theta, v, \omega]^T \tag{1}$$

heading angle, v represents the linear velocity, and ω represents the angular velocity. Considering the kinematic characteristics of the autonomous vehicle, this paper adopts a single-wheeled vehicle kinematic model to describe its local motion state, and the discrete simulation time step is set to:

$$\Delta t = 0.10s \tag{2}$$

For a given control input $\mathbf{u} = [v, \omega]^T$, the state of the autonomous vehicle is updated according to the following equation:

$$\begin{aligned} x_{k+1} &= x_k + v_k \cos(\theta_k)\,\Delta t \\ y_{k+1} &= y_k + v_k \sin(\theta_k)\,\Delta t \\ \theta_{k+1} &= \theta_k + \omega_k \Delta t \end{aligned} \tag{3}$$

Within a single prediction time step, the candidate linear and angular velocities remain constant, and local trajectories are generated through discrete integration. In the experiments in this paper, the prediction time step is set to:

$$T_{\text{pred}} = 2.0\text{s} \tag{4}$$

Therefore, each candidate trajectory consists of $N_{\text{pred}} = \frac{T_{\text{pred}}}{\Delta t} + 1$ discrete state points.

To ensure that the candidate control commands comply with the physical motion constraints of the autonomous vehicle, the system's maximum linear velocity, maximum angular velocity, maximum linear acceleration, and maximum angular acceleration are set as follows:

$$\begin{aligned} v_{\max} &= 1.45\text{m/s} \\ \omega_{\max} &= 1.35\text{rad/s} \\ a_{\text{v}} &= 1.0\text{m/s}^2 \\ a_{\omega} &= 1.8\text{rad/s}^2 \end{aligned} \tag{5}$$

During the current control cycle, the linear velocity dynamic window and the angular velocity dynamic window are expressed as follows:

$$\begin{aligned} V_d &= \{v | max(0, v_t - a_v \Delta t) \le v \le \min(v_{\max}, v_t + a_v \Delta t)\} \\ W_d &= \{\varpi | \max(-\varpi_{\max}, \varpi_t - a_\varpi \Delta t) \le \varpi \le \min(a_{\max}, \varpi_t + a_\varpi \Delta t)\} \end{aligned} \tag{6}$$

Within the dynamic window, linear velocity and angular velocity are sampled discretely, with sampling resolutions set at 0.08 m/s and 0.12 rad/s, respectively. At the same time, zero-velocity samples and exploration velocities constrained by safety clearance are introduced to increase the diversity of the candidate velocity space. The exploration velocity is defined as:

$$v_{explore} = \min\left(v_{high}, \max\left(v_{low}, 0.55 v_{safe}\right)\right) \tag{7}$$

In this context, v_{safe} is the maximum safe speed calculated based on the current obstacle clearance.

After the candidate trajectories are generated, calculate the minimum clearance distance between each trajectory point and the static rectangular obstacles, static circular obstacles, and dynamic circular obstacles, and determine feasibility based on the hard safety boundaries. If the clearance distance to any obstacle corresponding to a trajectory point does not exceed the hard safety boundary, the candidate trajectory is deemed infeasible.

$$d_{hard} = 0.05m \tag{8}$$

A time-varying motion model for dynamic obstacles is constructed based on their initial position, radius, velocity, and active time period. Spatio-temporal collision detection is performed by calculating the predicted positions at each discrete time step, thereby enabling the assessment of the spatio-temporal feasibility of local trajectories.

When there is no feasible trajectory that meets the safety constraints within the conventional dynamic window, it indicates that it is difficult for the conventional speed sampling to generate an effective trajectory that simultaneously satisfies the motion constraints and the safety distance requirements for obstacles in the current local environment. In response to this situation, this paper introduces a safety escape mechanism in the local planning execution stage. By constructing low-speed left and right turning trajectories and candidate trajectories for small radius turns, the search range for local feasible trajectories in the restricted state is expanded. For the generated candidate trajectories, safety checks are conducted point by point within the short-term prediction range, and the minimum obstacle clearance of each candidate trajectory is compared. When there is an escape trajectory that meets the safety boundary constraints, the candidate trajectory with a larger safety margin is prioritized for execution, enabling the unmanned vehicle to promptly leave the current restricted area and reducing the risk of stalling or falling into a locally unsolvable state due to the absence of feasible trajectories in the local area. This mechanism is mainly used to handle the special situation where no feasible trajectory can be obtained within the conventional dynamic window, thereby improving the continuity and executability of the local planning process.

2.1.2. Multi-Objective Trajectory Evaluation Methods

In order to precisely select the optimal motion trajectory that is safe, stable and suitable for the task, a multi-dimensional comprehensive evaluation of the feasible candidate trajectories is required. This evaluation focuses on the core characteristics of trajectory operation, the environment and task constraints, covering seven key evaluation dimensions such as heading consistency, obstacle clearance, movement speed, target progress, control smoothness, historical revisit and safety risks. It comprehensively considers the safety, stability, efficiency

and rationality of the trajectory in terms of passage, movement, task execution and path. The evaluation function is defined as:

$$J(\mathbf{u}) = w_H H_n + w_C C_n + w_V V_n + w_G G_n - w_{SM} SM - w_{loop} LP - w_{safe} SAFE \quad (9)$$

Among them, H_n、C_n、V_n、C_n represent the normalized heading evaluation, normalized clearance evaluation, normalized speed evaluation and normalized target progress evaluation respectively; SM is the angular velocity change penalty term, which is used to characterize the loss of trajectory control smoothness, LP is the historical revisit penalty term, which is used to constrain the problem of repeated circling of the path, and SAFE is the safety risk penalty term, which is used to quantify the potential hazards of trajectory passage. The heading evaluation item is calculated based on the angle between the heading at the predicted trajectory end and the direction from the end to the target.

$$H_n = 1 - \frac{\left|\text{wrapToPi}\left(\theta_{\text{goal}} - \theta_{\text{end}}\right)\right|}{\pi} \quad (10)$$

The heading evaluation item is used to measure the degree of closeness between the predicted trajectory's destination heading and the target direction. When the destination heading of the predicted trajectory points towards the target direction, it indicates that the robot's movement direction at the trajectory's end is consistent with the target direction, and a larger value can be taken; if the destination heading has a significant deviation from the target direction, this item should be appropriately set to a smaller value to guide the planner to preferentially select candidate trajectories with the final state facing the target. The clearance evaluation item is determined by the minimum obstacle clearance distance on the candidate trajectory:

$$C_n = clip\left(\frac{d_{\text{min}}}{d_{\text{clrCap}}}, 0,1\right) \quad (11)$$

When calculating obstacle distances, the safety cutoff distance is set to $d_{\text{clrCap}} = 4.0\text{m}$, and both dynamic and static obstacles are included in the calculation of the minimum obstacle distance d_{min}.

The overall evaluation function of DWA is composed of the speed evaluation item, the smoothness penalty item, the target progress item, the revisit penalty item, and the safety risk penalty item. Each evaluation item comprehensively scores the candidate trajectory from perspectives such as driving efficiency, motion smoothness, task achievement degree, state diversity, and safety. Then, the weighted sum is calculated based on the weights, and the candidate with the highest total score is the optimal trajectory for the current cycle. Among

them, the speed evaluation item $V_n$is used to encourage the vehicle to maintain a higher driving speed: under the premise of meeting the safety constraints, the higher the speed corresponding to the candidate trajectory, the larger the value of this item, which is more conducive to shortening the task duration and improving the traffic efficiency. In practical applications, the weight of this item can be dynamically adjusted according to the road speed limit and task timeliness requirements. The smoothness penalty item SM is used to suppress the sudden change in angular velocity between adjacent cycles.

A target progress item is defined as:

$$G = \left\|\mathbf{p}_t - \mathbf{p}_{goal}\right\| - \left\|\mathbf{p}_{end} - \mathbf{p}_{goal}\right\| \tag{12}$$

Here, P_t represents the current position of the autonomous vehicle, and P_{end} represents the end point of the candidate trajectory. The larger the value of G, the more significant the reduction in distance to the target that the candidate trajectory can achieve.

Furthermore, to prevent the unmanned vehicle from getting stuck locally or repeatedly circling in complex obstacle areas, a revisit penalty term LP based on spatial short-term memory is introduced into the evaluation function. The environment is rasterized at a certain resolution, and the vehicle's recent driving trajectory is recorded using a sliding time window. The frequency of access to each grid cell and its time decay weight are statistically calculated in real time. When the end point of the candidate trajectory coincides with a spatial area with high recent frequency of access, it indicates that this trajectory is likely to guide the vehicle to re-enter the traversed area. At this time, the LP penalty value is dynamically increased to reduce the comprehensive score of this trajectory, prompting the planner to preferentially select areas with low access frequency and high freshness of time, thereby effectively avoiding getting stuck in a local optimal solution. At the same time, the access memory decays exponentially over time, taking into account the accessibility requirements of narrow passages and other necessary visiting areas, avoiding blocking necessary revisit paths due to excessive penalties, and achieving a balance between suppressing repetitive circling and ensuring continuous passage.

A security risk penalty is defined as:

$$SAFE = \left[\frac{\max\left(0, d_{\text{safe}} - d_{\min}\right)}{d_{\text{safe}}}\right]^2 \tag{13}$$

In particular, the safety distance is used to characterize the minimum distance that should be maintained between an autonomous vehicle and an obstacle, and is set as follows:

$$d_{\text{safe}} = 0.90\text{m} \tag{14}$$

This measure expands the trajectory evaluation from a simple assessment of collision feasibility to a comprehensive evaluation of safety margin: The evaluation is no longer limited to determining whether the candidate trajectory will collide with obstacles, but further measures the minimum clearance distance from obstacles during the entire driving process and the change in safety margin over time, thereby distinguishing the safer ones among the feasible trajectory set, providing a continuous and comparable safety basis for subsequent comprehensive decision-making. Among the candidate control inputs that meet the safety constraints, the control input with the highest comprehensive evaluation value of $\mathrm{J(u)}$ is selected as the local planning result for the current control cycle. This control input is obtained by weighted and comprehensive calculation of the speed evaluation item, smooth penalty item, task progress item, revisit penalty item and safety risk penalty item, taking into account driving efficiency, motion smoothness, task achievement and obstacle avoidance safety, and is updated in real time based on the latest environmental information within each control cycle, enabling timely response to the dynamic obstacle environment.

2.2. DWA Local Trajectory Planning Based on ACT-R Cognitive Decision-Making

2.2.1. Overall Architecture and Strategy Definition

Traditional DWA uses fixed evaluation function weights, with each evaluation metric having a constant weight parameter throughout the planning process. To address the varying requirements of autonomous vehicles for path following, obstacle avoidance safety, and dynamic performance under different operating conditions, a dynamic weight adjustment mechanism is introduced. This mechanism adaptively adjusts the weights of each evaluation metric based on the vehicle's current motion state, thereby enhancing the adaptability of local trajectory planning to different scenarios.

The "Normal" strategy is primarily used in open environments and emphasizes goal progress and speed; the 'Avoid' strategy is primarily used in situations where the risk of obstacles is elevated, and it increases the clearance from obstacles, dynamic risk, and safety weight; the "Recover" strategy is primarily used in high-risk situations or when goal progress is obstructed, and it increases heading, clearance, smoothness, and safety weight, while reducing the speed weight.

The mapping from strategies to DWA weights can be written as:

$$w_t = F(s_t, w_{base}) \tag{15}$$

Here, s_t represents the current policy, and w_{base} represents the base weight vector.

To achieve multi-objective coordinated optimization, the basic weight vector of the evaluation model is defined as follows:

$$\mathbf{w}_{base} = [0.25,0.48,0.24,0.42,0.16,0.12] \tag{16}$$

These correspond to heading, clearance, speed, progress, smoothness, and safety, respectively.

It is important to emphasize that the fallback mechanism used by DWA when no feasible trajectory is available and the "Recover" strategy generated by ACT-R are two distinct decision-making mechanisms. The former is used to handle exceptional situations where no feasible solution is found during local trajectory planning, while the latter involves the ACT-R cognitive decision-making mechanism selecting a strategy based on the autonomous vehicle's current operating state. Therefore, the "Recover" state refers solely to the recovery strategy generated by the ACT-R cognitive decision-making layer; it does not indicate a failure of the underlying DWA local planning. The two should not be equated or confused.

During actual operation, the system employs an arbitration mechanism that prioritizes low-level hard constraints. Regardless of the strategy currently output by the ACT-R cognitive layer, once the underlying DWA determines that the set of feasible trajectories is empty, the system will immediately halt the transmission of top-level weights and forcibly trigger the low-level safety escape mechanism (low-speed steering or in-place rotation). In contrast, the Recover strategy generated by ACT-R constitutes a proactive risk intervention. Its purpose is to guide the vehicle away from high-risk areas by preemptively tightening the safety and clearance weights before the DWA reaches a completely unsolvable state, thereby reducing the probability of triggering the lower-level forced escape mechanism.

2.2.2. Development of the ACT-R Cognitive Decision-Making Model

This paper introduces the ACT-R cognitive architecture into the DWA local trajectory planner to enable environmental risk perception, retrieval of historical experience, policy utility learning, and adaptive adjustment of planning weights. The system employs three cognitive strategies: "Normal" indicates a relatively open environment, where the autonomous vehicle prioritizes progress toward its goal and operating speed; "Avoid" indicates that an obstacle has entered the local risk zone, prompting the vehicle to increase clearance and safety; and "Recover" indicates a high level of dynamic risk or a significant impediment to the vehicle's

progress toward its goal, requiring the implementation of stronger safety and recovery strategies.

Each strategy corresponds to a memory block. A memory block contains features associated with the distance to obstacles and the target state, and also records the history of when that strategy was used.

In the ACT-R mechanism, the baseline activation level of a memory chunk is determined jointly by its historical usage frequency and the corresponding time interval, and can be expressed as:

$$B_i(t) = \ln\left[\max\left(0.15, \sum_j \exp[-d(t - t_{ij})]\right)\right] \tag{17}$$

Here, $d = 0.45$ is the memory decay parameter, and t_{ij} represents the time at which strategy i is triggered or maintained for the jth time beyond the time threshold. Given that the local planner operates at a high frequency, this paper adopts an event-driven approach to record the timing of strategy usage. Specifically, an event is recorded as a distinct historical experience only when a cognitive strategy switches, or when the same strategy remains active continuously for longer than a set time threshold (e.g., 1.0 s). This approach prevents abnormal saturation of activation levels caused by high-frequency sampling.

Compared to the initial constant activation model, the current implementation accounts for real-time decay and the influence of usage history, resulting in a dynamically changing activation curve. Environmental risk is primarily determined by the current minimum obstacle clearance and the predicted dynamic obstacle clearance. Static risk is defined as:

$$r_o = clip\left(\frac{1.9 - \max(0, d_{obs})}{1.9}, 0, 1\right) \tag{18}$$

Dynamic risk is defined as:

$$r_d = clip\left(\frac{1.45 - \max(0, d_{dyn}^{pred})}{1.45}, 0, 1\right) \tag{19}$$

Here, d_{obs} represents the current clearance around all obstacles, and d_{dyn}^{pred} represents the minimum predicted clearance between the autonomous vehicle and dynamic obstacles within the short-term prediction window.

2.2.3. Cognitive Decision-Making Based on Memory Retrieval and Utility

Building upon the basic activation model, we further consider the degree of alignment between the current environmental state and strategy memory, as well as the impact of different

risk scenarios on strategy selection. We construct a strategy activation model that incorporates environmental alignment information and risk context information, which is expressed mathematically as：

$$A_i = 0.28B_i + 0.18M_i + 0.42C_i \tag{20}$$

Here, M_i represents the match between the current obstacle distance and the target state and the policy memory block, and C_i represents the risk context term.

In the full ACT-R model, the system integrates memory activation and decision utility to select a strategy. The basic form of decision utility is defined as:

$$U_i^d = 0.55(U_i - U) + 0.65A_i + U_i^{context} \tag{21}$$

Here, U_i represents the long-term utility of a strategy, $\overline{U}$ is the average utility of the three strategies, and $U_i^{context}$ is the environmental context term calculated based on the current static risk, dynamic risk, and progress risk.

Use target progress, clearance changes, and dynamic risks as feedback signals. Let the change in target distance be:

$$\Delta g = g_t - g_{t+1} \tag{22}$$

The changes in obstacle clearance are as follows:

$$\Delta d = d_{t+1} - d_t \tag{23}$$

The reward vectors corresponding to the three strategies are:

$$\begin{aligned} R_1 &= 3.5\Delta g - 1.5r_o \\ R_2 &= 2.4\Delta d + 1.0r_o - 0.7(1 - r_o) + 0.4\Delta g \\ R_3 &= 1.5\Delta d + 2.8r_d - 1.1(1 - r_d) + 0.5\Delta g \end{aligned} \tag{24}$$

To curb the unbounded growth in utility caused by abnormal feedback, the reward function incorporates a bound that limits it to:

$$R_i \in [-2.5, 2.5] \tag{25}$$

This paper proposes a feedback-driven mechanism for incrementally updating the utility of a strategy:

$$U_i^{new} = (1 - \alpha_i)U_i + \alpha_i(U_i^0 + R_i) \tag{26}$$

$$\alpha_i = 0.10(0.35 + 0.65I[i = s]) \tag{27}$$

Furthermore, to ensure the stability of utility values during the dynamic update process, we

introduce a weak constraint from the prior utility on the current utility, causing the current utility to regress toward the prior utility at a modest rate. Specifically, this is expressed as:

$$U_i \leftarrow 0.975U_i + 0.025U_i^0 \tag{28}$$

This design prevents any single strategy from becoming permanently locked in and ensures that the accumulation of Recover utility does not permanently suppress Normal and Avoid.

Given that environmental risks exhibit a certain degree of temporal continuity, to prevent sudden changes in risk states from causing significant fluctuations in Activation and Utility within a single control cycle, a first-order finite-response filtering mechanism is introduced to smooth the updates of relevant state variables, thereby reducing the impact of instantaneous risk changes on policy decisions. The computational form is as follows:

$$\begin{aligned} A_{i,t}^f &= 0.82A_{i,t-1}^f + 0.18A_{i,t} \\ U_{i,t}^f &= 0.78U_{i,t-1}^f + 0.22U_{i,t} \end{aligned} \tag{29}$$

Therefore, Activation and Decision Utility still respond to changes in dynamic obstacles and target progress, but the curves exhibit continuous transitions, which better align with the nature of limited responses in cognitive processes.

2.2.4. ACT-R-Driven Dynamic Weight Adjustment and Feedback Updates in DWA

Strategy selection follows the principle of maximum decision utility:

$$s_t = \arg\max_i U_{i,t}^f \tag{30}$$

To prevent the strategy from oscillating frequently between adjacent control cycles, a cooling time and a hysteresis threshold are set. When the utility difference between the candidate strategy and the current strategy is less than the hysteresis threshold, or when the duration of the current strategy is less than the cooling time, the system maintains the original strategy:

$$s_t = s_{t-1}, \quad \text{if} U_{\max} - U_{\text{current}} \le h \text{or} \tau_{\text{current}} < T_{\text{cool}} \tag{31}$$

Specifically, when $\begin{matrix} h = 0.20 \\ T_{\text{cool}} = 0.70\text{s} \end{matrix}$, the system allows a strategy switch only if the utility advantage exceeds the hysteresis threshold and the cooling time meets the requirements.

After the strategy switch is complete, the weights of the DWA evaluation function are adaptively adjusted based on the current cognitive strategy state. Specifically, the “Normal” state emphasizes progress and speed; the “Avoid” state emphasizes clearance and risk; and the “Recover” state emphasizes heading, clearance, smoothness, and safety, while appropriately reducing the weight assigned to speed.

The complete ACT-R model further couples the patterns based on the strategy mapping:

$$\begin{aligned} w_C &\leftarrow 1.04 w_C \\ w_{SM} &\leftarrow 1.06 w_{SM} \\ w_G &\leftarrow 1.24 w_G \\ w_V &\leftarrow 1.18 w_V \\ w_{safe} &\leftarrow 1.03 w_{safe} \end{aligned} \tag{32}$$

This setting reflects the comprehensive balance of Full ACT-R-DWA in terms of safety, trajectory smoothness, and target progress, rather than simply focusing on improving any single evaluation item. The weights of each item and the intensity of penalties are adaptively adjusted according to the current environmental conditions and the vehicle's operating status. In open areas, a greater emphasis is placed on speed and target progress to improve traffic efficiency. In complex obstacle areas, safety risks and recall penalties are strengthened to ensure the margin for obstacle avoidance and path diversity. Thus, the algorithm can maintain the coordination and unity of safety, smoothness, and task completion in various scenarios. Therefore, the method proposed in this paper forms the following closed loop. The algorithm flowchart is shown in Figure 1:

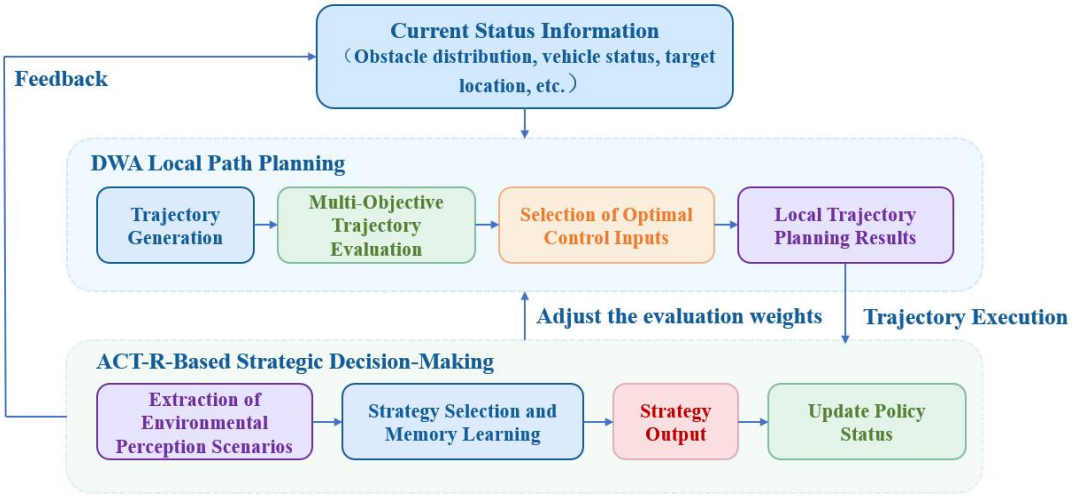


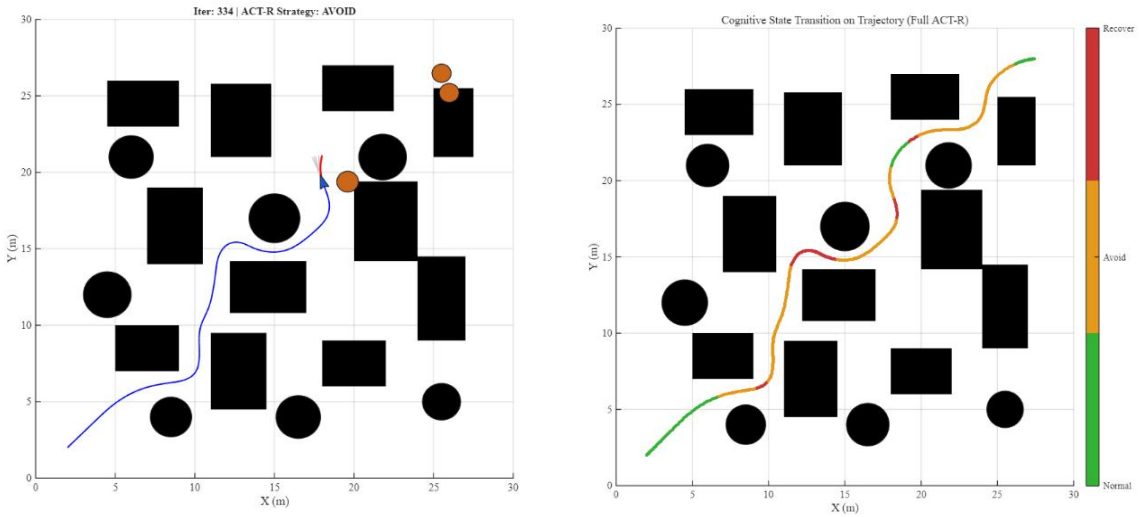
Figure 1. Algorithm Flowchart

3. Results

3.1. Comparative Analysis of Local Trajectory Planning Results

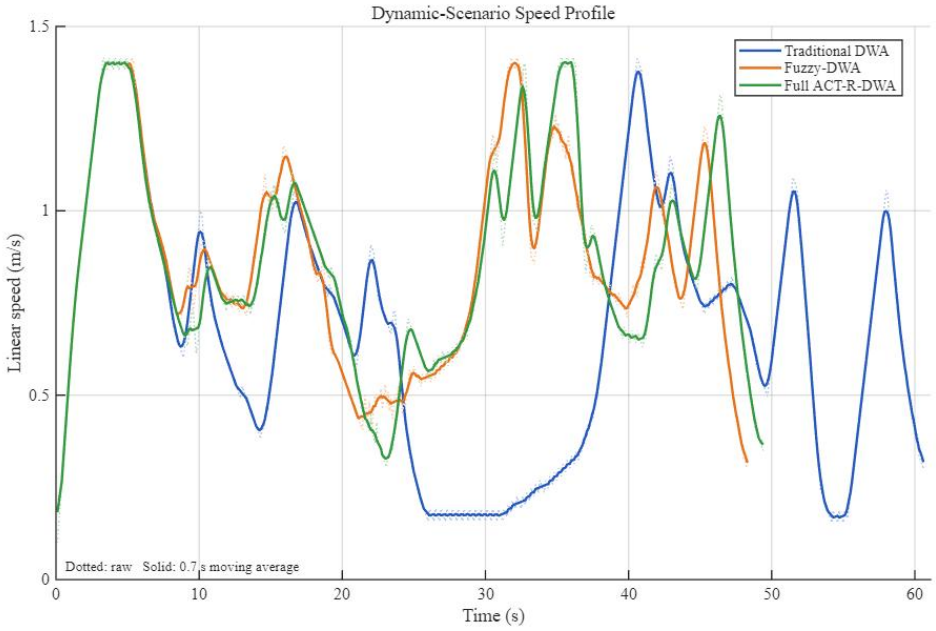
The experiment was conducted in a two-dimensional plane of 30m×30m. The starting point and the ending point of the unmanned vehicle were set as (2.0, 2.0) and (28.0, 28.0), respectively. A total of two scenarios were set for the test: the static scenario (Static Only), which only contained stationary obstacles, and the dynamic scenario (Static + Dynamic), which additionally introduced three dynamic obstacles. The positions of the dynamic obstacles were

updated in real time according to their movement attributes and were incorporated into the collision prediction model of the candidate trajectories. To verify the performance of the algorithm, this paper compared three traditional algorithms: DWA, Fuzzy-DWA, and Full ACT-R-DWA. The planning process and state transition of Full ACT-R-DWA are shown in Figure 2. The motion responses and safety clearances of different local planning algorithms in the dynamic obstacle scenario are compared in Figure 3, and the performance comparison table of the three algorithms is shown in Table 1.

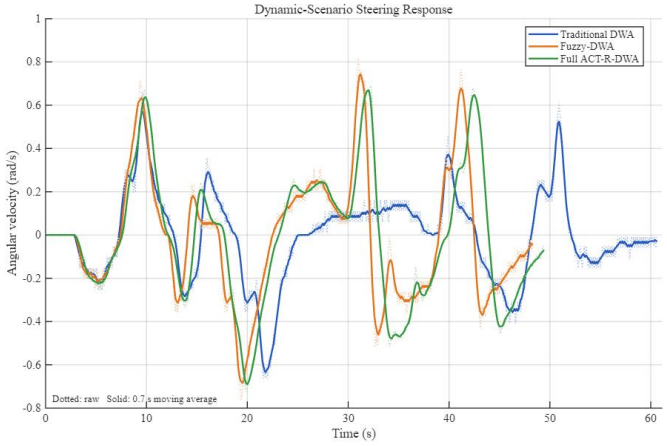


(a) Diagram of the Full ACT-R-DWA Planning Process (b) Cognitive strategy state transitions along the full ACT-R-DWA trajectory

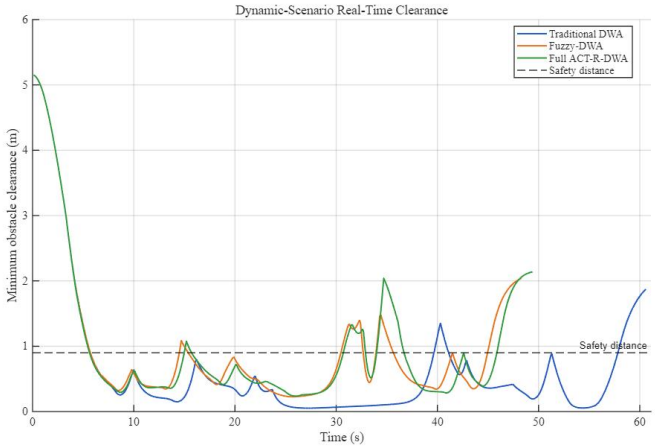
Figure 2: The Full ACT-R-DWA Planning Process and State Transitions in Cognitive Strategies



(a) Linear velocity response in scenarios with moving obstacles



(b) Angular Velocity Response in Scenarios with Moving Obstacles



(c) Real-time minimum obstacle distance in dynamic obstacle scenarios

Figure 3: Comparison of Motion Responses and Safety Margins for Different Local Planning Algorithms in a Dynamic Obstacle Scenario

Table 1: Performance Comparison of the Three Algorithms

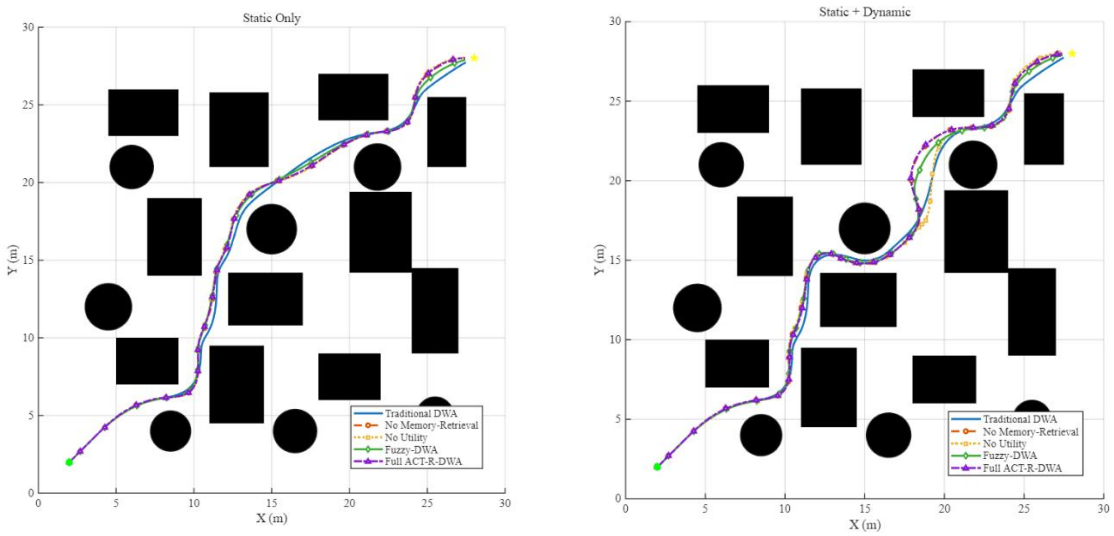
Algorithm	**Success Rate (%)**	**Number of collisions**	**Minimum Distance (m)**	**Path Length (m)**	**Average Speed (m/s)**	**Arrival Time (s)**	**Smoothness**	**Scheduled Time (ms)**
Traditional DWA	100	0	0.052	40.180	0.663	60.6	0.451	4.295
Fuzzy-DWA	100	0	0.226	41.764	0.865	48.3	0.496	4.634
Full ACT-R-DWA	100	0	0.237	42.610	0.863	49.4	0.242	4.586

Combining the simulation trajectories and kinematic responses shown in Figures 2 and 3, it can be seen that Full ACT-R-DWA is capable of smoothly switching between three cognitive

strategies—"Normal," "Obstacle Avoidance," and "Recovery"—based on local environmental risks, demonstrating dynamic decision-making and situational awareness capabilities. At the same time, it maintains a high average operating speed, smooth angular velocity changes, and sufficient real-time safety clearance, thereby avoiding the issue of approaching collision limits inherent in traditional methods. As shown in Table 1, while both methods achieved 100% zero collisions, Full ACT-R-DWA demonstrated significant advantages in terms of safety, operational efficiency, and trajectory quality. Compared to Traditional DWA: the minimum obstacle distance increased from 0.052 m to 0.237 m; the average speed increased from 0.663 m/s to 0.863 m/s, representing a 30.17% improvement; the arrival time decreased from 60.6 s to 49.4 s, a reduction of 18.48%; and the smoothness decreased from 0.451 to 0.242, representing a 46.34% improvement. Compared to Fuzzy-DWA: the minimum obstacle distance increased from 0.226 m to 0.237 m, a 4.87% improvement; smoothness decreased from 0.496 to 0.242, a significant improvement of 51.21%; and the average single-step planning time decreased from 4.634 ms to 4.586 ms, a reduction of 1.04%.

3.2. Experiment on the Ablation of ACT-R Cognitive Mechanisms and Performance Analysis

To validate the roles of each component module of ACT-R, four ablation scenarios were set up: Traditional DWA, No Memory-Retrieval, No Utility, and Full ACT-R-DWA. Additionally, Fuzzy-DWA—which does not rely on ACT-R—was included as a side-by-side baseline.

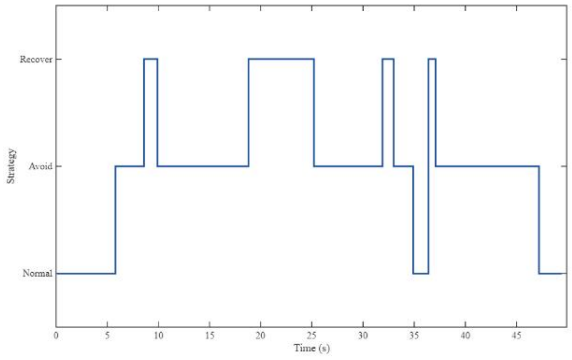


(a) Trajectory Comparison in Scenarios with Static Obstacles (b) Trajectory Comparison in Static–Dynamic Mixed Obstacle Scenarios

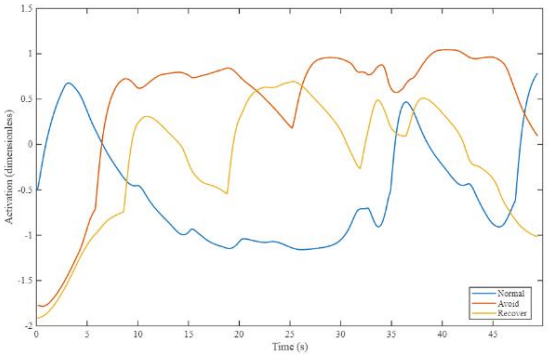
Figure 4: Comparison of Planned Trajectories Under Different Ablation Strategies

A comparison of planning trajectories under different strategies is shown in Figure 4. Based

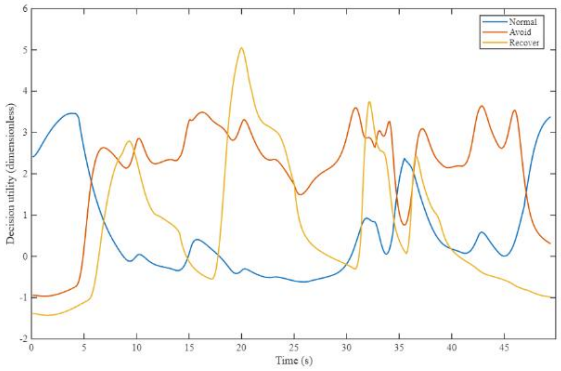
on this, a hybrid static-dynamic scenario was selected to analyze the operation of Full ACT-R-DWA. Figure 5 illustrates the process of ACT-R cognitive decision-making and the dynamic adjustment of DWA evaluation weights. The system achieves strategy competition and switching through memory retrieval activation and utility feedback, which are ultimately mapped to DWA's multiple evaluation weights to enable adaptive adjustment.



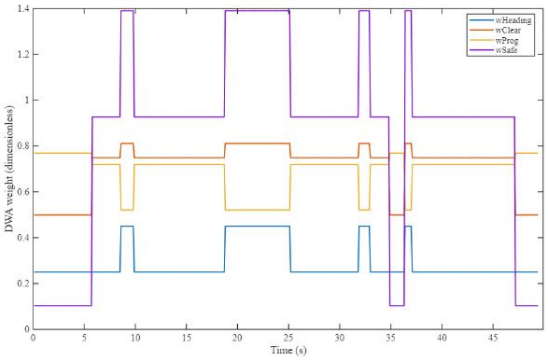
(a) The ACT-R Strategy State Transition Process



(b) Changes in memory retrieval activation across the three types of strategies



(c) The Evolution of Utility in ACT-R Strategy Decision-Making



(d) The process of dynamically adjusting DWA evaluation weights driven by ACT-R

Figure 5: Dynamic Adjustment Process of Evaluation Weights in ACT-R Cognitive Decision-Making and DWA

As shown in Figure 4, compared with the trajectory deviation and stagnation phenomenon that occur in traditional DWA in complex mixed scenarios, Full ACT-R-DWA can generate smoother obstacle avoidance trajectories and effectively suppress the speed changes in adjacent cycles, indicating that the cognitive-driven evaluation adjustment significantly improves the trajectory continuity and driving smoothness while ensuring the safety margin. Figure 5 (a) shows that the system strategy can dynamically switch according to the environmental state, and the hysteresis mechanism effectively suppresses the high-frequency oscillations during the strategy switching process, avoiding the control command jitter caused by environmental

perception disturbances. Figure 5 (b) shows that the activation level of memory retrieval evolves continuously based on historical experience and environmental matching degree, accurately reflecting the degree of correlation between the environmental state and historical calls, providing empirical support corresponding to the scenarios for strategy selection. Figure 5 (c) demonstrates the utility competition process, where the strategy switching is determined by the maximum value of the utilities of each cognitive module rather than a preset fixed rule, reflecting the system's adaptive decision-making ability for different scenarios. Figure 5 (d) indicates that the cognitive strategy can adjust the weights of each evaluation item of DWA in real time, verifying the effective driving of the ACT-R cognitive layer for local planning, and thereby achieving the comprehensive achievement of safety, smoothness and efficiency goals.

4. Discussion

In dynamic local trajectory planning, the traditional DWA usually adopts fixed evaluation weights and relies only on current perceptual information, which makes it difficult to adapt to time-varying obstacle distributions, cannot utilize historical decision-making experience, and tends to cause strategy oscillation when the environment changes rapidly. To address issues such as fixed evaluation weights, lack of historical experience utilization, and strategy oscillation in dynamic local trajectory planning, this study proposes an adaptive dynamic window approach (DWA) based on the ACT-R cognitive mechanism. By introducing a cognitive decision layer into the local planner, the proposed method transforms the static weight configuration of the conventional DWA into a cognition-driven adaptive adjustment process, enabling the evaluation function to be dynamically matched with the current environmental state.

The proposed "memory-utility-mapping" architecture overcomes traditional rule-based or numerical optimization methods by integrating historical dependence into real-time decision-making. Specifically, the event-driven memory activation model retrieves the most relevant historical experience according to the similarity between the current environment and past scenarios, while the feedback-driven utility learning model updates the preference strength of each candidate strategy based on the planning results. By combining the event-driven memory activation model with the feedback-driven utility learning, the system can capture the "environment-experience-strategy-feedback" loop, forming a closed loop in which historical experience continuously informs strategy selection and the selection results in turn revise the accumulated experience. Additionally, introducing cooling time and lag threshold, this system suppresses high-frequency strategy switching, thereby achieving a smooth real-time mapping

from cognitive decision-making to multi-objective DWA weights, which effectively avoids the control jitter and trajectory oscillation caused by frequent switching.

The comparison of simulation results shows that even when the success rate is as high as 100%, the complete ACT-R-DWA still outperforms the traditional DWA and the fuzzy DWA.Compared to traditional DWA, its minimum obstacle spacing increases from 0.052 meters to 0.237 meters, the average speed increases by 30.17% (0.863 m/s), and the trajectory smoothness improves by 46.34% (0.242), indicating that the proposed method substantially enlarges the safety margin while improving both driving efficiency and trajectory quality. Compared to Fuzzy-DWA, this method improves the smoothness by 51.21% while maintaining the average single-step planning time of only 4.586 milliseconds, meeting the real-time requirements. These results demonstrate that the ACT-R-driven adaptive DWA achieves a favorable trade-off among safety, efficiency, and smoothness in complex dynamic environments.

5. Conclusion

This paper addresses the issues of fixed evaluation weights, insufficient utilization of historical experience, and frequent oscillations in the strategy of the traditional Dynamic Window Approach (DWA) for unmanned vehicles in complex dynamic environments. To address these problems, a DWA adaptive local trajectory planning method based on the ACT-R cognitive mechanism is proposed. The main research conclusions are as follows:

(1) At the levels of trajectory generation and multi-objective evaluation: The trajectory generation and evaluation system based on DWA has been improved. By introducing spatial-temporal collision prediction and revisit penalty mechanisms, a multi-objective comprehensive evaluation function and a safety escape arbitration mechanism have been designed, providing a reliable underlying execution framework for the trajectory optimization in cognitive decision-making.

(2) At the cognitive decision-making and adaptive mapping level: An integrated ACT-R cognitive architecture of "memory - utility - mapping" was constructed. Through an event-driven mechanism, historical memories were activated, and combined with feedback-driven utility updates, the smooth evolution of three types of cognitive strategies was achieved; at the same time, cooling time and hysteresis mechanisms were introduced to inhibit frequent switching of strategies, ultimately realizing the adaptive real-time mapping of cognitive decisions to DWA evaluation weights.

(3) At the levels of simulation comparison : System simulation indicates that, under the condition of 100% zero collision arrival, Full ACT-R-DWA will increase the minimum obstacle distance from 0.052 m to 0.237 m, shorten the arrival time by 18.48%, and improve the smoothness index by 46.34% compared with Traditional DWA; compared with Fuzzy-DWA, it improves the minimum obstacle distance by 4.87% and the smoothness index by 51.21%, and the average single-step planning time is only 4.586 ms, fully meeting the real-time control requirements.

In conclusion, the DWA adaptive local trajectory planning method based on the ACT-R cognitive mechanism proposed in this paper takes into account the safety of local obstacle avoidance, motion efficiency, control smoothness, and the ability to utilize historical experience, providing a feasible solution for the autonomous smooth maneuver of unmanned vehicles in complex dynamic environments.

Acknowledgements

This project is supported by This project is supported by Jinan 2023 talent development special fund research leader studio project (Grant No. 202333067). Shandong Province Science and Technology oriented Small and Medium Enterprises Enhancement Project (Grant No. 2023TSGC0288). Shandong Provincial Higher Educational Youth Innovation Science and Technology Program (Grant No. 2019KJB019).

References

[1] Han, Y., Liu, Y., Guan, T., Lan, L., & Tang, N. (2025). Improved A* and dynamic window approach for unmanned vehicle path planning. Journal of Shandong University (Engineering Science Edition), 55(3), 16-24. DOI:10.6040/j.issn.1672-3961.0.2024.023

[2] Zhu, M., Hu, F., Li, N., Zhu, S., & Wu, Q. (2023). Summary of path planning algorithms for driverless vehicle. Agricultural Equipment & Vehicle Engineering, 61(11), 18-22. DOI:10.3969/j.issn.1673-3142.2023.11.004

[3] Zhang, R., Zhou, L., & Liu, Z. (2024). Dynamic path planning for mobile robot based on RRT* and dynamic window approach. Journal of System Simulation, 36(4), 957-968. DOI: 10.16182/j.issn1004731x.joss.22-1543

[4] Zhao, L., Li, C., Zhang, W., Chen, D., & Li, Z. (2024). Optimization of robot static dynamic local path planning method based on integrating attraction-repulsion and dynamic window approach. Information and Control, 53(2), 226-237. DOI: 10.13976/j.cnki.xk.2023.2578

[5] Wang, X., Qi, Z., Yang, Z., Wang, D., Huang, H., & Lu, M. (2025). Local path planning for agricultural robots based on DAV_DWA. Transactions of the Chinese Society for Agricultural Machinery, 56(2), 105-114.

[6] Xu, W., Zhang, Y., Yu, L., Zhang, T., & Cheng, Z. (2023). A local path planning algorithm based on improved dynamic window approach. Journal of Intelligent & Fuzzy Systems, 45(3), 4917-4933.

[7] Gong, X., Gao, Y., Wang, F., Zhu, D., Zhao, W., Wang, F., & Liu, Y. (2024). A local path planning algorithm for robots based on improved DWA. Electronics, 13(15), 2965.
[8] Liu, A., Liu, C., Li, L., Wang, R., & Lu, Z. (2025). An improved fuzzy-controlled local path planning algorithm based on dynamic window approach. Journal of Field Robotics, 42(2), 430-454. DOI: 10.1002/rob.22419
[9] Hossain, T., Habibullah, H., Islam, R., & Padilla, R. V. (2022). Local path planning for autonomous mobile robots by integrating modified dynamic-window approach and improved follow the gap method. Journal of Field Robotics, 39, 371-386.
[10] Kobayashi, M., & Motoi, N. (2022). Local path planning: Dynamic window approach with virtual manipulators considering dynamic obstacles. IEEE Access, 10, 17018-17029.
[11] Yang, Z., Li, N., Zhang, Y., & Li, J. (2023). Mobile robot path planning based on improved particle swarm optimization and improved dynamic window approach. Journal of Robotics, 2023, Article ID 6619841, 1-16.
[12] Sarkis, S. S., Elfayoumi, O. S., Jaradat, M. A., & Emran, B. J. (2025). Adaptive dynamic window approach for mobile robot motion control using particle swarm optimization. IEEE Access, 13, 212104-212120. DOI: 10.1109/ACCESS.2025.3624130
[13] Cina, M., & Rad, A. B. (2023). Categorized review of drive simulators and driver behavior analysis focusing on ACT-R architecture in autonomous vehicles. Sustainable Energy Technologies and Assessments, 56, 103044.
[14] Oh, H., Yun, Y., & Myung, R. (2024). Driver behavior and mental workload for takeover safety in automated driving: ACT-R prediction modeling approach. Traffic Injury Prevention, 25(3), 381-389.
[15] van der Velde, M., Sense, F., Borst, J. P., & van Maanen, L. (2022). Capturing dynamic performance in a cognitive model: Estimating ACT-R memory parameters with the linear ballistic accumulator. Topics in Cognitive Science, 14(4), 889-903.
[16] Yang, Y. C., & Stocco, A. (2024). Allocating mental effort in cognitive tasks: A model of motivation in the ACT-R cognitive architecture. Topics in Cognitive Science, 16(1), 74-91.
[17] Fisher, C. R., Houpt, J., Larue, O., & Schmidt, K. (2025). Using systems factorial technology for global model analysis of ACT-R's core architectural assumptions. Journal of Mathematical Psychology, 125, 102924.

Print ISSN: 3104-932X
Online ISSN: 3104-9338

Constructing a Generative AI Framework for Personalised Instructional Resource Development: A Quasi-Experimental Evaluation in a Micro-Lecture Design and Production Course

Guihua Gao*, Hanyan Cai, Shiru Zhu

Yiyang Normal College, Yiyang 413000, Hunan, China

Received: October 8, 2026

Revised:October 17, 2026

Accepted: October 21, 2026

Published online: October 22, 2026

To appear in: *International Journal of Advanced AI Applications,* Vol. 2, No. 11(November 2026)

Corresponding Author: Guihua Gao (2260540996@qq.com)

Abstract. Generative artificial intelligence is promoted as a way of tailoring instruction to learners, yet most reported applications stay within language and general-education settings, and few measure outcomes in courses whose value lies in design and production skill. This paper describes a four-module generative AI framework for resource development that links learner and content modelling, generation, assessment and community iteration to a course-grounded knowledge base, with the teacher reviewing before release. It was evaluated in a twelve-week quasi-experiment in a professional core course, Micro-lecture Design and Production, at a Chinese vocational college. Two intact classes (n = 45, 44) were taught by the same instructor, who designed the framework; assignment was not random. Covariance-adjusted posttest means differed by 5.6 points (82.3 versus 76.7, $p < .001$, partial $\eta^2 = .24$), and the framework-supported class reported higher motivation and satisfaction; self-directed learning differed modestly. Preparation time per unit fell by half, though the work became more concentrated and the knowledge base carried a large one-off cost. Because the design cannot separate the framework from the additional resources it produced, the results are framed as associations rather than effects.

Keywords: *generative artificial intelligence; personalised learning; instructional resource design; micro-lecture production; teacher education*

1. Introduction

Large language models now draft explanations, summarise readings, compose practice questions and rewrite passages for different reading levels, a capability documented from the transformer architecture [1] to the scale at which instruction-following behaviour emerges [2].

Differentiated instructional material is among the most labour-intensive parts of teaching and is now one of the tasks generative AI supports most directly [3,4].

Institutional responses have been cautious but constructive. UNESCO treats generative AI as a tool whose value depends on the governance and pedagogy built around it rather than on the model [5], and reviews reach a similar conclusion about the potential and the risks [6]. What determines whether such a tool helps is not what it generates but whether its output fits the curriculum, the learners and the assessment already in place [7].

The empirical literature has not kept pace. Reviews keep finding the same profile — small perception surveys, a concentration on writing and language learning, general-education settings [8,9] — with three recurring weaknesses named in recent scoping work [10]. Many studies report what students think about generative AI rather than what they learn with it; where outcomes are measured, the intervention is usually the model's general capability rather than a designed workflow another teacher could reproduce; and courses assessed on what students make rather than what they recall are almost absent.

The gap is most consequential in teacher education, where using these tools well sits at the junction of technology, pedagogy and subject matter [11], and studies of pre-service teachers report positive attitudes but seldom test what those teachers can actually design [12]. Micro-lectures make the question concrete: short videos are routine in blended, vocational and distance teaching [13], and production quality rests on design judgements — how far to segment a topic, how narration should sit against imagery — that experienced teachers apply tacitly and that govern how long students watch and how much they retain [14].

We designed and evaluated a generative-AI-supported framework for personalised instructional resource development (PIRD) in the professional core course Micro-lecture Design and Production, taken by second-year Modern Educational Technology students at a Chinese higher vocational college. The evaluation is a twelve-week quasi-experiment in two intact classes combining an examination with validated scales, platform logs and interviews; it asked whether achievement differs in a course whose outcomes are design and production skill, how the affective outcomes compare, and what changes for the teacher. Two boundaries apply from the outset: classes were not randomly assigned and the instructor was also a researcher, so the design supports statements about association rather than effect, and the framework was deployed as a package, so no difference can be attributed to a single component (Section 5.2).

2. Related Work

2.1. Generative AI in Education: From Capability to Application

Research on generative AI in education has grown quickly, with an uneven centre of gravity. A systematic review catalogued largely pedagogical opportunities and challenges — adaptive pacing, timely feedback, reduced administrative burden — along with the persistent difficulty of aligning automated output with learning outcomes [4], while a companion review of higher education found empirical work clustered around writing support, language learning and student perceptions, with little attention to how instructors build the technology into a course [8]. Recent scoping work adds a warning: without deliberate design, LLM output favours fluent but shallow text over the disciplinary reasoning courses exist to develop [10].

Studies that do look at instruction tend to examine how students use a general-purpose chatbot, not how a purpose-built system performs. Surveys across four universities found benefits framed as efficiency and quality of explanation, and concerns framed as accuracy and over-dependence [15]; a study in science education found strong performance on conceptual questions and unreliable performance on procedural ones [16]. What remains unsettled is not whether a language model can produce instructional text, but how to organise that capability so the text is accurate, well pitched and tied to assessment.

2.2. Personalised and Adaptive Learning: The Older Evidence Base

Personalised instruction is an old ambition. Bloom framed the problem four decades ago: one-to-one tutoring produces outcomes about two standard deviations above conventional group instruction [17]. Reviews of adaptive learning document steady growth alongside persistent methodological problems [18], and a multi-year implementation study found modest average effects that varied across sites, with implementation strength predicting outcomes better than the technology chosen [19].

A synthesis of this literature distinguishes what gets personalised: pace, content, sequence or modality [20]. Most adaptive systems personalise pace and sequence through item-response estimation, which helps little when the difficulty is representational — when a student cannot connect a step she can already perform to the reason it works. Generative AI is attractive exactly here, because it can produce alternative explanations and representations of the same concept on demand. Two theories constrain how those alternatives should be graded: cognitive load theory holds that worked examples lose their advantage as expertise grows, so the useful form of a resource depends on the learner's current state rather than a fixed template [21], while the

zone of proximal development adds that support should sit just beyond what the learner can already do alone [22]. The framework operationalises both crudely, as a difficulty band with a target success probability of .60 to .85 on the practice items attached to a resource (Section 4.5); a three-level band is not a continuous estimate of what a student can do with help.

2.3. Resource Design and Assessment in Teacher-Education Courses

Instructional design for complex professional skills has its own tradition. The four-component model organises learning tasks by whole-task progression [23], and Merrill's first principles make the same point from the learner's side: new knowledge is best demonstrated in the context of a task the student will perform [24].

Courses that produce instructional media add a second body of design knowledge. Multimedia learning research has established that the way verbal and pictorial information is combined determines how much of it learners can process, and that a small set of principles — segmenting a presentation into learner-paced parts, presenting narration and imagery together instead of duplicating one in the other, and signalling what matters — has consistent and substantial effects [25]. These principles specify what a script, a storyboard and a screen layout should look like, so a resource can be judged against criteria that do not depend on taste; Section 4.5 audits all 36 released resources against five of them. Assessment research points the same way: formative assessment produces gains when it gives students information they can act on [26], the effect depends on the quality of feedback rather than its quantity [27], and students benefit most from feedback that helps them judge their own work against a standard [28]. For generative AI this is also a hazard, since the model that supplies an immediate explanation can also supply a fluent but wrong one: hallucination is a documented behaviour of language generation rather than an occasional defect [29], which is why grounding output in a retrieved corpus has become standard wherever factual reliability matters [30].

Teacher knowledge is the third element, and it is where this study sits. TPACK separates what a tool can supply from what only a teacher can decide [11]: the framework supplies content knowledge (a concept graph and an item bank), technological content knowledge (tiered explanations and annotated storyboards) and technological pedagogical knowledge (automated assessment and the infrastructure for interaction), but no module supplies pedagogical content knowledge — the judgement about what counts as a good micro-lecture on this topic, for these students. That judgement sits at the review stage, which is why the stage is obligatory rather than optional, and why the study measured student outcomes and teacher time rather than the teacher's TPACK.

3. Methodology

3.1. Research Design

The study used a quasi-experimental pretest-posttest design with two intact classes; random assignment was not possible within the institution's timetable. The class that received framework-generated resources is referred to as the experimental class and the other as the control class. A pretest was administered in week one, the intervention ran for twelve teaching weeks, and a posttest followed in week thirteen, with platform logs, interviews and teacher-side time logs as supplementary data.

Three features of the design bound what can be concluded. Both classes were taught by the same instructor, who had designed the framework and knew the group assignment, so instructional behaviour cannot be assumed identical across classes. With 45 and 44 students the study had 80% power to detect a difference of about d = 0.60, so smaller effects, the common case in education research, would have gone undetected. The intervention was deployed as a complete package, so the analysis estimates the association between the package and the outcomes rather than the contribution of its components.

3.2. Setting, Participants and Ethics

The study was conducted at a public higher vocational college in Hunan Province, China, in the professional core course Micro-lecture Design and Production, which runs for sixteen weeks at four class hours per week in a digital media production workshop and covers the full production cycle, from script writing to narration recording, editing and publishing. Two second-year classes took part, giving 89 students after three were excluded for missing more than two of the twelve weekly quizzes; one instructor taught both, with 45 and 44 students respectively, using the same syllabus, textbook, workshop sequence and examination paper.

Students were told that declining or withdrawing would not affect their grades, that the two classes would be taught with different supporting materials, and that the comparison would be disclosed to both classes at the end of the semester; written consent was obtained from all 89. Platform data were anonymised and students were referenced by pseudonymous identifiers.

Two measures were taken against the most obvious sources of bias. The examination was marked against a scheme finalised before the intervention began, and its four design questions were scored independently by two raters from a different department who had no teaching contact with either class, were not told which class was which, and rated de-identified submissions (agreement, ICC(2,2) = .81). Neither addresses instructional bias: the author

designed the framework and taught both classes, and its screening rules were fixed before the semester rather than adjusted during it (Section 5.4).

Because the materials supplemented rather than replaced instruction, the control class received the full curriculum; delayed provision was judged the least restrictive way to permit a comparison, and the package was released to control-class students after the posttest (Section 4.7).

3.3. The PIRD Framework

The design rests on one principle: generative AI should be embedded in a defined instructional function, with a human decision at the point where material reaches students. It has four modules supported by a shared service layer and is bounded above by a teacher review stage that no resource bypasses (Figure 1); Table 1 summarises what each module does and which teacher-side task it replaces.

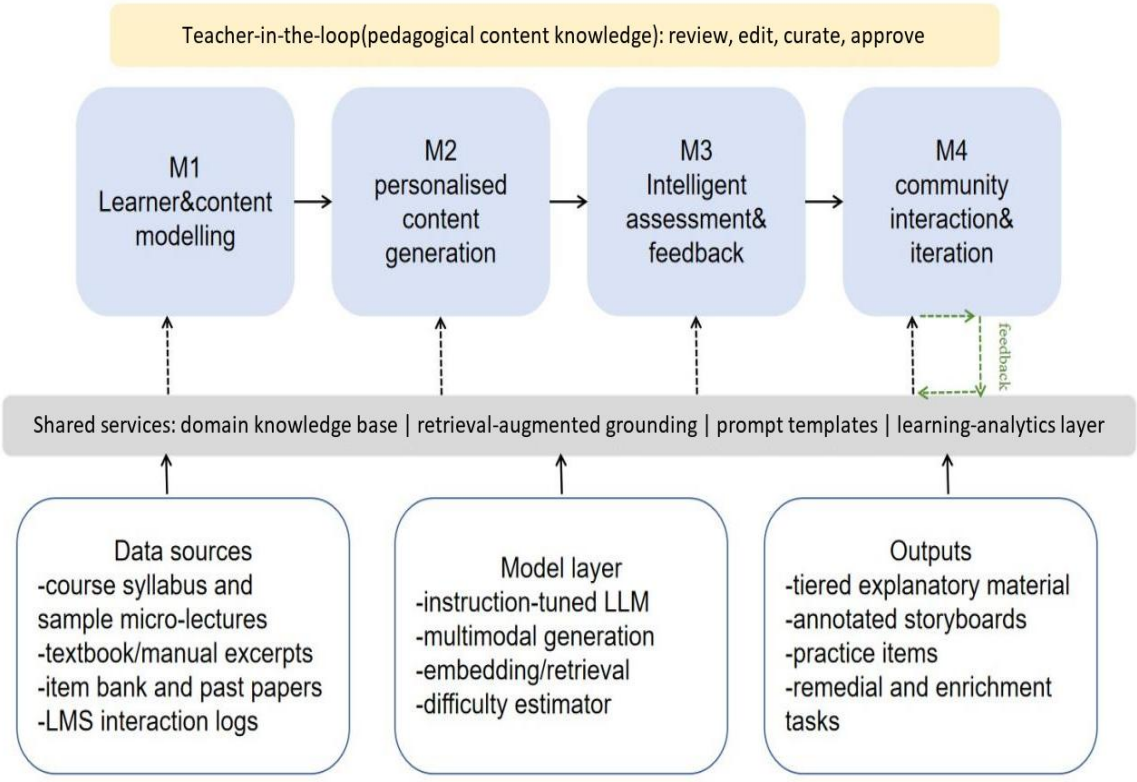


Figure 1. Architecture of the proposed PIRD framework, showing the four modules, the shared service layer, the feedback path from assessment back to learner modelling, and the teacher review stage at which pedagogical content knowledge is exercised.

Two parts of the knowledge base were built rather than collected, and their provenance bears on transfer. The concept graph was constructed by the instructor together with experienced

teachers of the same course and revised over two rounds of independent review, in which a node was retained only where the reviewers judged it a prerequisite for a topic the course assesses and relations were kept only where they agreed on direction; the resulting graph was then checked against the previous cohort's error patterns on the weekly quizzes. The item bank was calibrated before release rather than written to a target in the abstract. Items inherited from earlier cohorts carry difficulty estimates derived from their p-values in those cohorts' quizzes, and items that were newly generated were reviewed by hand before entering the bank, so that every item is stored with the topic and the difficulty band it was assigned. This is what allows a generation request to specify a target difficulty band instead of leaving that judgement to the model.

Table 1. Functional design of the four framework modules.

Module	Primary function	Key technologies	Teacher-side task replaced
M1. Learner and content modelling	Per-student mastery and representation profiles; course concept graph	Diagnostic scoring; log analysis; prerequisite mapping	Manual review of quiz and homework records
M2. Personalised content generation	Tiered explanations, worked examples, guided practice items	Instruction-tuned LLM; retrieval-augmented generation; storyboard rendering	Drafting of differentiated handouts and exercise sets
M3. Intelligent assessment	Calibrated items, automated scoring, error profiles	Difficulty estimation; structured marking schemes; concept tagging	Item writing; marking of objective items; error analysis
M4. Community interaction and iteration	Annotated discussion; approval-gated replies; resource regeneration	Forum integration; request logging; weekly analytics	Handling repetitive questions; revising materials between cohorts

The system runs on an instruction-tuned LLM accessed through a commercial API, with an off-the-shelf embedding model for retrieval [2]. No model was fine-tuned: discipline-specific behaviour came from retrieval grounding and prompt design, which keeps the system portable across courses and puts the burden of accuracy on the knowledge base. Module 3 reuses the same embedding model for tagging instead of training a separate classifier.

3.4. Knowledge Grounding and Prompt Design

Generation follows a fixed seven-step workflow (Figure 2): define the objective, the difficulty band and the output format; retrieve grounded passages; assemble a structured prompt from a versioned template; produce three candidate resources; screen them automatically; obtain the teacher's review and release; and collect performance data, which updates the learner

profile and closes the loop.

The screening in step 5 is where an unchecked model could release a plausible error. Each candidate is checked on alignment with the objective, on unsupported factual claims relative to the retrieved passages, on readability, and on consistency between the declared difficulty band and the estimated item difficulty. From week six, after the audits reported in Section 4.5 showed that design-level properties were being missed, signalling markers and temporal contiguity between narration and imagery were added as further checks. A candidate that fails is regenerated once, and a second failure flags it for manual attention.

Prompts were reusable templates with four slots — role, task, constraint and exemplar. One constraint-slot template asks the model to explain to second-year pre-service teachers how narration and on-screen imagery should be synchronised in a micro-lecture, in at most 400 characters and drawing only on the retrieved passages, while the exemplar slot holds one instructor-approved resource that fixes tone and structure. Structured prompts of this kind, and explicit step-by-step reasoning cues, change output quality substantially [31]. The templates used here are the fourth revision, iterated through a two-week pilot.

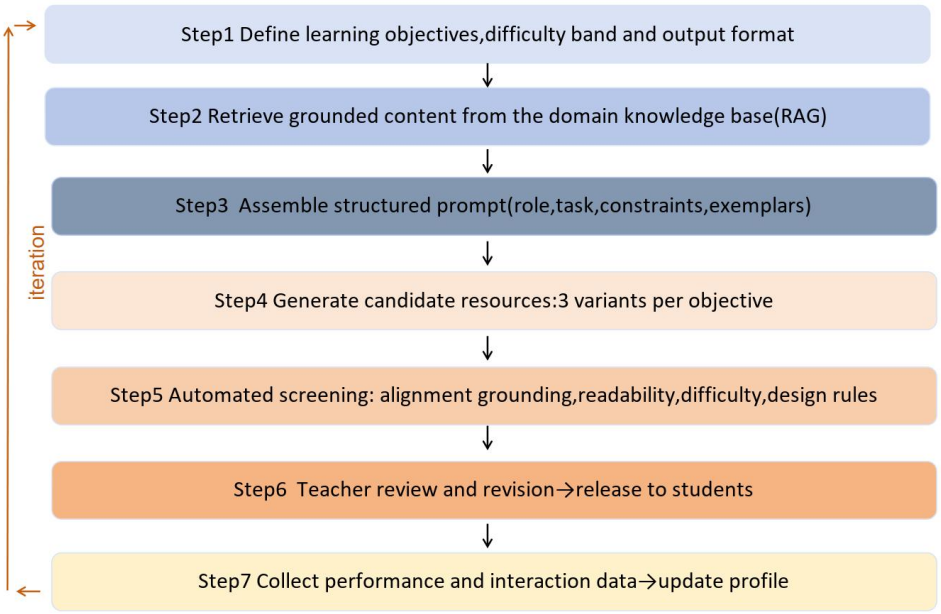


Figure 2. Seven-step resource generation workflow with automated screening and teacher review.

Three design decisions shaped the results. Every request carries an explicit difficulty band instead of leaving that judgement to the model. The grounding requirement is enforced in the prompt and then checked automatically, which reduces unsupported claims without eliminating them. And the framework offers several variants at once, on the reasoning that choosing between options is easier than specifying the answer in advance — a decision that had a

consequence for students we had not anticipated (Section 4.7).

Retrieval is fixed rather than tuned. The corpus described in Section 3.3 is chunked and indexed offline, and a request is answered from a bounded set of the highest-ranked passages for that query rather than from the whole corpus; the grounding check in step 5 is then carried out against exactly those passages, so the retrieved set places a hard ceiling on what can be released. The same retriever serves generation in Module 2 and concept tagging in Module 3, so a released resource and the error profile derived from it rest on the same representation of the course content. Two properties of this layer matter for replication. It is a replaceable component rather than part of the intervention: the embedding model is a general-purpose, off-the-shelf one and, like the generator, was not fine-tuned, so the framework's behaviour was pinned to the knowledge base and the prompt templates rather than to a particular vector space, and a re-implementation can substitute another embedding model without altering the framework's logic. And no retrieval setting was tuned on outcome data: the values in force during the semester were those configured before it began and were left unchanged while it ran.

3.5. Instruments

Four instruments were used. Achievement was measured by a 100-point examination with 40 objective items worth 40 marks and four design questions worth 60 marks, written by the teaching team and piloted on a previous cohort. Learning motivation used an adapted 14-item MSLQ subscale covering intrinsic goal orientation and task value [32]; satisfaction used a 10-item scale adapted from earlier work on technology-supported courses; and self-directed learning used a 12-item scale covering goal setting, strategy use and self-monitoring [28].

All self-report scales used a five-point Likert format, and reliability was estimated separately at pretest and posttest (Table 2); the motivation measure was the weakest at pretest, and results involving it are read with that in mind. Semi-structured interviews were held with 12 students from the experimental class, chosen to span the achievement distribution, lasting between 22 and 38 minutes; intercoder agreement reached Cohen's $\kappa = .74$.

Because the three self-report scales are adaptations rather than new instruments, their adaptation is described here. Each was translated into Chinese and then back-translated to check that the items retained their meaning, with discrepancies resolved by discussion rather than by dropping items, so that the scales can be read as the instruments they are adapted from. Two of the three also inherit a validated factor structure: the motivation measure is the MSLQ subscale whose intrinsic goal orientation and task value dimensions were established in the original

validation [32], and the self-directed learning measure follows the goal setting, strategy use and self-monitoring structure reported by Zimmerman [28] and used since. The present study therefore treats the factor structure of the scales as inherited rather than re-established, and reports reliability rather than a new factor solution. With 89 respondents and 36 items across the three scales, the case-to-item ratio is roughly 2.5 to 1, well below the 5 to 1 usually advised for exploratory factor analysis and far below what a confirmatory model would need; a solution estimated on this sample would be unstable, and none is reported. Scale scores are used as observed composites, and the coefficients in Table 2 should be read alongside the original validations rather than in place of them.

Table 2. Instruments, constructs, and reliability coefficients.

Instrument	Construct measured	Items	Response format	Source	α (pretest) / α (posttest)
Achievement test	Course knowledge and micro-lecture design reasoning	40 objective + 4 design	Score 0-100	Written by the teaching team, externally reviewed	n/a / .87
Motivation scale	Intrinsic goal orientation and task value	14	5-point Likert	Adapted from Pintrich et al. [32]	.79 / .83
Satisfaction scale	Perceived usefulness, material clarity, continuance intention	10	5-point Likert	Adapted from prior technology-supported course studies	.84 / .86
Self-directed learning scale	Goal setting, strategy use, self-monitoring	12	5-point Likert	Adapted from Zimmerman [28]	.81 / .85

3.6. Intervention Procedure

The intervention ran for twelve consecutive teaching weeks. In the experimental class, each topic was supported by three framework-generated resources released at the start of the week alongside the standard textbook reading, and students could request alternative explanations and extra practice items through the platform, with every request logged. The control class covered the same topics with the same syllabus and workshop sequence, using the textbook, the instructor's slides and a fixed exercise set from previous cohorts, with no additional materials and no access to the assistant. The asymmetry in material volume is treated as a confound rather than a design feature: experimental students received roughly three times as many resources and, through the request mechanism, substantially more practice and feedback

(Section 5.2).

Several features were held constant: the same contact hours, workshop exercises and weekly ten-item quiz. The quizzes differed in origin but not in coverage, since experimental quizzes came from the approved pool of generated items and control quizzes from the historical bank, with each item matched on topic and estimated difficulty beforehand. What could not be held constant is how much additional material each student worked through, since time on task was not recorded for the control class, whose practice material was distributed on paper.

3.7. Data Analysis

Analysis was carried out in SPSS 26. Baseline equivalence was tested with independent-samples t tests and χ^2 tests, and we report 95% confidence intervals for the differences rather than relying on the significance test alone, since at this sample size a non-significant difference is not evidence of equivalence.

The primary outcome, posttest achievement, was analysed with analysis of covariance using pretest score as the covariate. Four assumptions were checked and are reported with the result: normality of the residuals within groups (Shapiro-Wilk), homogeneity of error variance (Levene's test), a linear relationship between covariate and outcome within each group, and homogeneity of regression slopes. The between-group difference is reported as an adjusted mean difference with a 95% confidence interval, next to the raw difference.

Affective and skill outcomes were analysed with a mixed analysis of variance, group between subjects and time within. The group × time interaction is the confirmatory test; where it was significant, the posttest between-group comparison is reported as a difference with a 95% confidence interval and as a standardised mean difference with its own interval. Because four confirmatory outcomes were tested, the corresponding p values are also reported after Holm correction [33] within that family. Interview transcripts were analysed thematically following Braun and Clarke's six-phase procedure [34], with two coders working independently and disagreements resolved by discussion.

4. Results

4.1. Baseline Equivalence

The two classes did not differ significantly on any baseline measure. As Table 3 shows, they were comparable in gender composition, age, prior grade point average and prior experience with AI tools, and their pretest achievement scores were statistically indistinguishable (68.4

versus 67.2, t(87) = 0.58, p = .56). Self-reported motivation, satisfaction and self-directed learning at pretest also showed no significant differences, with all p values above .60.

The intervals tell a more careful story than the tests. The 95% confidence interval for the pretest achievement difference was [-2.9, 5.3] points, so a real baseline gap of several points cannot be excluded, and the same holds for the other continuous measures. With 89 students, equivalence is not something this study can demonstrate, so baseline balance is a reason for caution rather than a licence to attribute the posttest difference to the intervention.

Table 3. Baseline characteristics of the two intact classes.

Characteristic	Experimental (n = 45)	Control (n = 44)	Test statistic	p	Difference [95% CI]
Female, n (%)	15 (33.3)	14 (31.8)	$\chi^2 = 0.02$	.88	
Age, M (SD)	19.4 (0.8)	19.5 (0.8)	t = -0.58	.56	-0.1 [-0.5, 0.3]
Prior GPA, M (SD)	76.8 (6.4)	77.3 (6.1)	t = -0.38	.71	-0.5 [-3.1, 2.1]
Prior use of AI tools, n (%)	27 (60.0)	24 (54.5)	$\chi^2 = 0.27$	.60	
Pretest achievement, M (SD)	68.4 (9.6)	67.2 (10.1)	t = 0.58	.56	1.2 [-2.9, 5.3]
Pretest motivation, M (SD)	3.71 (0.62)	3.66 (0.59)	t = 0.39	.70	0.05 [-0.20, 0.30]
Pretest satisfaction, M (SD)	3.55 (0.60)	3.58 (0.63)	t = -0.23	.82	-0.03 [-0.28, 0.22]
Pretest self-directed learning, M (SD)	3.42 (0.58)	3.45 (0.61)	t = -0.24	.81	-0.03 [-0.28, 0.22]

Note. CI = confidence interval. All tests are two-tailed with df = 87. Dashes indicate categorical variables, for which a difference in means is not defined.

4.2. Academic Achievement

The experimental class scored higher on the end-of-term examination. Raw posttest means were 82.6 (SD = 8.1) and 76.4 (SD = 9.3), a difference of 6.2 points, 95% CI [2.6, 9.8], d = 0.71 [0.28, 1.14]. Analysis of covariance with pretest score as the covariate gave adjusted means of 82.3 and 76.7, an adjusted difference of 5.6 points, 95% CI [2.9, 8.3], F(1, 86) = 27.4, p < .001, partial $\eta^2 = .24$. All four assumption checks were acceptable, including homogeneity of regression slopes. The full test statistics are given in Table 4.

The distribution of scores says something the means do not: 31 of 45 experimental students (68.9%) scored above 80, against 18 of 44 (40.9%) in the control class, $\chi^2(1) = 7.03$, p = .008, and 4 against 11 scored below 70, so the difference was not confined to the top of the distribution.

Table 4. Achievement test results with analysis of covariance.

Group	n	Pretest M (SD)	Posttest M (SD)	Adjusted posttest M (SE)	95% CI for adjusted mean
Experimental	45	68.4 (9.6)	82.6 (8.1)	82.3 (0.95)	[80.4, 84.2]
Control	44	67.2 (10.1)	76.4 (9.3)	76.7 (0.96)	[74.8, 78.6]

Note. Adjusted difference = 5.6 [2.9, 8.3]; $F(1, 86) = 27.4$, $p < .001$, partial $\eta^2 = .24$; raw difference = 6.2 [2.6, 9.8], Cohen's d = 0.71 [0.28, 1.14]. Assumption checks: Shapiro-Wilk $W = .98$, $p = .31$ (experimental) and $W = .97$, $p = .24$ (control); Levene $F(1, 87) = 1.42$, $p = .24$; covariate-outcome $r = .62$ and .58 within groups; covariate-by-group $F(1, 85) = 0.71$, $p = .40$. CI = confidence interval.

The two halves of the examination moved differently: on the 40 objective items the classes scored 34.1 and 32.4, a difference of 1.7 marks, and on the four design questions 48.5 and 44.0, a difference of 4.5 marks. The proportional gap is larger on the design questions, where a well-grounded worked example should help most, but a sample of this size cannot establish that the two gaps differ reliably, so the pattern is suggestive rather than evidence that the framework acted specifically on design reasoning.

Weekly quizzes tracked the same pattern (Figure 3): the classes began close together, at 71.2 and 70.4, and the gap widened during weeks four to seven, which cover script writing and storyboarding, then held at about six points. That timing is not accidental, since those are the two topics with the densest prerequisite structure in the concept graph, and they generated 214 of the 583 requests logged across the semester.

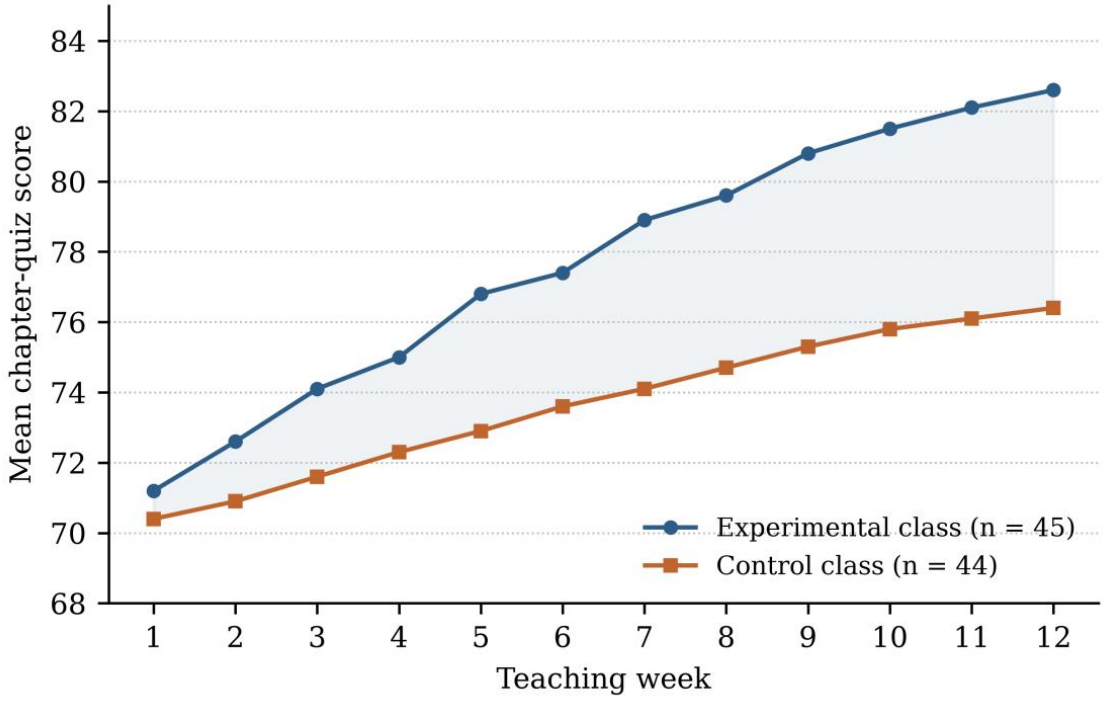


Figure 3. Mean weekly quiz scores for the two classes across twelve teaching weeks.

One further check bears on the volume confound discussed in Section 5.2. If the difference were driven mainly by students working through more material, platform use should track performance within the experimental class. It does not: request count correlated weakly with posttest score, r(43) = .14, p = .36, and students in the top quartile of requests scored 83.1 against 82.4 for the rest, t(43) = 0.36, p = .72. With 45 students this is a weak test, and a null association cannot exclude a volume effect that operated equally across the class.

4.3. Motivation and Satisfaction

Both classes became more motivated, and the framework-supported class did so to a greater degree: the group × time interaction was F(1, 87) = 11.6, p = .001, partial η^2 = .12, with posttest means of 4.19 (SD = 0.52) against 3.84 (SD = 0.57), a difference of 0.35, 95% CI [0.12, 0.58], d = 0.64 [0.21, 1.07]. Learning satisfaction followed a similar and slightly stronger pattern: 4.33 (SD = 0.48) against 3.97 (SD = 0.55), a difference of 0.36, 95% CI [0.15, 0.57], interaction F(1, 87) = 18.9, p < .001, partial η^2 = .18, d = 0.70 [0.27, 1.13]. Descriptive statistics appear in Table 5 and Figure 4; the posttest comparisons appear in Table 6.

The two kinds of statistic are not interchangeable: in Table 5 the interaction F and partial η^2 describe change over time, whereas the standardised differences in Table 6 describe the gap between classes at posttest. The multiplicity correction changes no conclusion, since every Holm-adjusted p value leads to the same decision as the uncorrected one.

Table 5. Pretest and posttest scores on the affective and skill outcomes.

Variable	Exp. pretest M (SD)	Exp. posttest M (SD)	Ctrl. pretest M (SD)	Ctrl. posttest M (SD)	F (group × time)	p	partial η^2
Learning motivation	3.71 (0.62)	4.19 (0.52)	3.66 (0.59)	3.84 (0.57)	11.61	.001	.12
Learning satisfaction	3.55 (0.60)	4.33 (0.48)	3.58 (0.63)	3.97 (0.55)	18.90	<.001	.18
Self-directed learning	3.42 (0.58)	4.05 (0.51)	3.45 (0.61)	3.71 (0.58)	5.41	.022	.06

Note. All F tests in this table are the mixed-analysis-of-variance group × time interaction, with df = 1, 87. The partial η^2 values describe that interaction; posttest between-group differences, including the covariate-adjusted achievement comparison (F(1, 86)), are reported in Table 6.

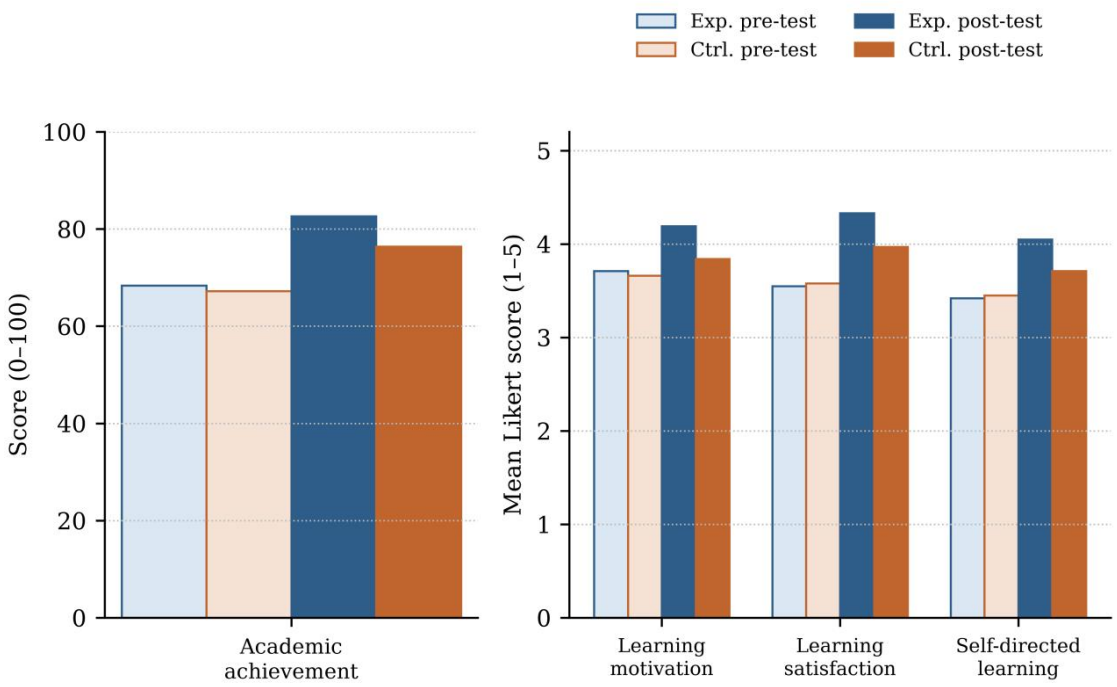


Figure 4. Pretest and posttest scores on achievement and on the three affective and skill measures.

Table 6. Posttest between-group comparisons with confidence intervals and multiplicity correction.

Variable	Difference [95% CI]	Cohen's d [95% CI]	p	p (Holm)
Academic achievement (adjusted)	5.6 [2.9, 8.3]	0.71 [0.28, 1.14]	<.001	<.001
Learning motivation	0.35 [0.12, 0.58]	0.64 [0.21, 1.07]	.001	.003
Learning satisfaction	0.36 [0.15, 0.57]	0.70 [0.27, 1.13]	<.001	<.001
Self-directed learning	0.34 [0.11, 0.57]	0.62 [0.19, 1.05]	.022	.044

Note. Achievement is the covariate-adjusted difference in examination marks; the remaining rows are posttest differences in scale points. Holm = Holm-Bonferroni correction within the family of four confirmatory comparisons. CI = confidence interval.

4.4. Self-Directed Learning and Creative Design Performance

Self-directed learning followed the affective measures with a smaller margin: posttest means of 4.05 (SD = 0.51) and 3.71 (SD = 0.58), a difference of 0.34, 95% CI [0.11, 0.57], d = 0.62 [0.19, 1.05], interaction $F(1, 87) = 5.4$, $p = .022$, partial $\eta^2 = .06$. This is the weakest of the significant comparisons and the one that would change if a stricter correction were used; Holm correction within the family leaves it significant at .044, and we read it as tentative rather than settled.

A fifth measure is noted only for completeness. A five-dimension creative design rubric

produced a nonsignificant difference of 0.07 scale points, 95% CI [-0.22, 0.36], p = .25, with a ceiling on three of the five dimensions. Because it measured a transfer the framework was not built to produce and could not discriminate among competent work, the result is uninformative in either direction and is not treated as a finding.

4.5. Resource Quality: Multimedia Design Principles and Difficulty Calibration

The principles cited in Section 2.3 are the criteria against which the micro-lectures students were asked to produce would be judged, so a framework that models them should satisfy them. All 36 released resources were audited against five of these principles by two raters, with agreement of Cohen's κ = .71 (Table 7).

Table 7. Audit of released resources against five multimedia design principles.

Multimedia principle	Operational check applied to the released resource	Resources passing, n = 36 (%)
Segmenting	Content divided into learner-paced parts of 120 seconds or less	33 (92)
Signalling	Key elements explicitly marked in the script or in the framing	30 (83)
Redundancy	No duplication of spoken narration as on-screen text	35 (97)
Temporal contiguity	Narration and the imagery it describes presented together	28 (78)
Modality	Spoken narration used in place of on-screen prose alongside imagery	34 (94)

Note. Rated by two raters from scripts and storyboards; inter-rater agreement, Cohen's κ = .71. From week six the screening step added checks for signalling and temporal contiguity, after which compliance with contiguity rose from 64% to 89% and with signalling from 72% to 91%.

Compliance was high on three principles and weak on the other two, and the weak two were exactly the properties the screening step did not check before week six: screening on textual properties did not transfer to design properties visible only when narration and image are considered together.

Difficulty calibration can be assessed in the same way, since Module 3 scores the practice items that accompany each resource. Mean realised accuracy was .72, inside the .60 to .85 target in 29 of 36 resources (81%), and misses were concentrated in the first three weeks, when 5 of 12 resources fell outside the band against 2 of 24 thereafter. Because realised accuracy is a post-hoc quantity while the band was set in advance, this checks calibration rather than tests it.

4.6. Resource Development Efficiency

The teacher-side effects rest on logs rather than self-report and are the most directly interpretable figures in the study. Table 8 compares the conventional workflow used with the previous cohort against the framework-supported workflow used here, per teaching unit of four class hours, across the 16 units of the course.

Table 8. Resource development efficiency per teaching unit of four class hours.

Indicator	Conventional workflow	Framework-supported workflow	Change
Drafting time per unit, h	4.5	1.1	-76%
Review and editing time per unit, h	1.0	1.7	+70%
Other preparation time per unit, h	1.3	0.4	-69%
Total preparation time per unit, h	6.8	3.2	-53%
Resources released per unit	1	3	+200%
Preparation time per released resource, h	6.8	1.07	-84%
Share of preparation time requiring disciplinary judgement, %	31	74	+43
Disciplinary-judgement time per unit, h	2.1	2.4	+14%
Global workload rating per unit (1-5), M (SD)	4.31 (0.28)	3.14 (0.41)	-1.17
First-cohort cost per unit including knowledge base, h	6.8	20.7	+204%

Note. Conventional figures come from the previous cohort of the same course (n = 47), taught by the same instructor. Unit = four class hours; with 16 units, paired comparisons have df = 15. Knowledge-base construction is a one-off cost of about 280 hours, excluded from the steady-state rows and included in the first-cohort row; the approach reaches break-even after roughly 78 teaching units of reuse.

Total preparation time fell from 6.8 to 3.2 hours per unit, a paired difference across the 16 units of -3.6 hours, $t(15) = -11.4$, $p < .001$, and its composition changed: drafting fell from 4.5 to 1.1 hours while review and editing rose from 1.0 to 1.7 hours. The task shifted from authoring to reviewing.

The saving needs three qualifications, and they matter more than the headline. Released resources per unit tripled, so on a per-resource basis the comparison is not 3.2 against 6.8 hours but 1.07 against 6.8: part of the gain was reinvested in more varied material rather than in

reduced workload. The work also became more concentrated. The instructor judged that 74% of the framework-workflow time required disciplinary judgement, against 31% before, so although total preparation time roughly halved, the absolute time spent on disciplinary verification changed little, from about 2.1 to 2.4 hours per unit. The global workload rating fell from 4.31 (SD = 0.28) to 3.14 (SD = 0.41), a paired difference of -1.17, t(15) = -9.7, p < .001, consistent with less total effort but a heavier demand per minute. Finally, these figures describe steady-state operation and exclude the cost of building the knowledge base: charged to the first cohort alone, that cost raises the figure from 6.8 to about 20.7 hours per unit, and the approach breaks even only after roughly 78 teaching units, or five cohorts.

Review effort rose by 70%: 68% of released resources were accepted without edit, 22% needed minor revision and 10% were discarded, and the discard rate fell from 19% in the first three weeks to 4% by week eight, suggesting that the templates and the knowledge base improved through use.

4.7. Student Perceptions

Thematic analysis of the twelve interviews produced four themes. Perceived relevance came up most often: students valued the fit between the generated material and the specific teaching points they were working on, and several contrasted it with general-purpose chatbots, which produced answers that sounded right but did not match the workshop's approach. Cognitive load was the second theme: the tiered resources made it easier to start a topic, but a smaller group said the abundance of material created a new problem, namely which of three variants to read first. Trust was the third, and it pointed in two directions at once: students trusted the material because the instructor had approved it, distrusted general chatbots for the mirror-image reason, and yet credited the discussion that followed a generated explanation, rather than the explanation itself, for their improvement. That suggests the teacher-in-the-loop stage works as a credibility signal and not only as a quality filter.

The fourth theme, raised by six students, was fairness. Two in the experimental class worried that access to the assistant gave them an advantage they had not earned, and four asked at the disclosure session whether the control class had been disadvantaged. Two control-class students had asked during the semester why the classes were taught differently and were told only that the difference would be explained later, which protected the comparison but left a legitimate question unanswered. After disclosure, 29 of 44 control-class students (66%) downloaded at least one unit's set within two weeks and nine attended an additional tutorial; delayed provision did not resolve the equity question, because half the cohort still spent a semester without a

support that the study suggests was useful.

5. Discussion

5.1. Interpreting the Achievement Difference

The central result is a difference of about 5.6 points between two intact classes in a course that would seem to resist automated support. Micro-lecture design is procedural, judgement-heavy and tool-specific: quality depends on decisions experienced teachers rarely put into words, and generic models produce scripts that read fluently while ignoring the principles that make a short video learnable. Because assignment was not random, and both classes were taught by the instructor who built the framework, the difference is an association and not an effect, and three features of the design are plausible candidates for it, none separable from the others.

Grounding is the first: restricting generation to passages retrieved from a course-specific knowledge base constrains the model to the terminology, design conventions and evaluation criteria students meet in the workshop, which addresses the mismatch they themselves described. Difficulty calibration is the second: because every request carried an explicit difficulty band and a representation preference, material was pitched at the level the learner profile indicated, and the check reported in Section 4.5 shows the target was met for most resources. This connects to the personalisation literature, addressing the representation and explanation dimensions that item-response systems handle poorly [20], and operationalising, crudely, the graded fading of support that cognitive load theory recommends [21]. Teacher review is the third: it did not remove the instructor but moved the effort from composing to verifying and, in TPACK terms, left pedagogical content knowledge with the person who had it, so the difference reflects generated resources, expert review and the additional material released together, which this study cannot separate.

5.2. What This Design Cannot Separate

Three rival explanations are worth naming, and one is more plausible than the framework-based account. The first is the bundle itself: the framework combined a knowledge base, retrieval grounding, difficulty calibration, teacher review, three variants per objective, automated screening and a request mechanism, and the study reports a single estimate for all of them together. A difference attributed to "the framework" is therefore a statement about a package that could be more or less effective than any of its parts.

The second explanation is the most plausible of the three and the least exotic: the

experimental class received more of everything — three resources per unit instead of one, a mechanism for requesting further explanations and further practice, and a teacher who was reviewing material rather than writing it. Increasing the quantity of practice with feedback is one of the better-established ways to raise achievement, and this study did not hold it constant. The platform logs offer a partial answer: request frequency was unrelated to posttest score within the experimental class, which argues against a simple dosage account for the more engaged students but cannot rule out an effect operating for everyone. Until the comparison is made, the conservative reading is that this study evaluated a resource-rich package, not an AI component within a fixed quantity of instruction.

The third is differential attention. A teacher who designed a framework and knows which class is receiving it is not a neutral instrument, and expectancy effects are hard to exclude when the same person teaches both groups. Measurement was constrained by a fixed marking scheme and blind scoring, but delivery was not, so the estimate is better read as an upper bound. A design that could decompose the package needs at least three arms beyond a control: the full framework, a grounded-only arm with calibration and the variant mechanism disabled, and a volume-matched arm in which the control class receives three conventional resources and an equivalent practice set with no generative component. The third arm matters most, because it separates the technology question from the quantity question.

5.3. Implications for Practice

The knowledge base is the binding constraint. The efficiency gain came from keeping generation on grounded content, so the quality of that grounding decides whether the output is useful or merely fluent. An institution considering a similar approach should budget for knowledge-base construction as the main cost, not for model access, which is cheap by comparison.

The cost structure is not what the headline number suggests. Preparation time per unit halved, but per released resource it fell much further, and the absolute time spent on disciplinary verification barely moved; the pattern to expect is a redistribution towards less drafting, more verification, more volume and more concentration. In TPACK terms a framework of this kind can supply content, representational and assessment knowledge, but the judgement about what counts as good work in a given course stays with the teacher, which is why the review stage is not a transitional safeguard that better models will make unnecessary.

There is a student-facing implication too: some students handled three variants well and others did not know where to start, so a short orientation in the first week, and possibly a

recommended reading order in the interface, would fix a problem the framework created rather than solved. Finally, on these figures the approach is more expensive than the conventional workflow for at least the first cohort and becomes cheaper only on reuse — a normal profile for an investment in reusable material, but not what "preparation time fell by half" conveys on its own.

5.4. Limitations

The most serious limitation is the non-random assignment of classes. The groups were equivalent on every baseline measure collected, but the confidence intervals in Table 3 show how little that establishes at this sample size, unmeasured differences between intact classes cannot be ruled out, and the design does not support causal claims. The instructor also designed the framework, taught both classes and knew the group assignment, so expectations cannot be entirely excluded even though the examination scheme was fixed in advance and the design questions were scored blind. Replication with instructors who did not build the system is needed before the results travel.

The second limitation is the bundle described in Section 5.2: the study cannot attribute the difference to grounding, calibration, teacher review or added volume, and the volume explanation is the most obvious rival account. Time on task was not measured for the control class, and the study could detect only effects of about $d = 0.60$ or larger, so the self-directed learning result would not have been detected had it been smaller. The scope is narrow for the same reason: one course, one institution, 89 students, twelve weeks, with the posttest collected immediately after the intervention, so nothing here speaks to retention.

The self-report scales carry a further limitation of the same kind: they were translated and back-translated but not re-validated in this cohort, so their construct validity rests on the original instrument validations rather than on evidence collected here.

Finally, the ethics of unequal access deserve to be stated rather than acknowledged in passing. Half the cohort spent a semester without a support that appears to have helped their peers, a deliberate design choice made to permit a comparison. Releasing the materials after the posttest mitigates but does not undo the asymmetry, and in an institutional deployment the question of who receives AI-supported instruction, and on what basis, would have to be settled as policy rather than as an experimental condition.

6. Conclusion

This study asked whether generative AI can support instruction in a course whose outcomes

are made rather than recalled, when it is organised around a defined framework instead of used as a general-purpose assistant. The evidence is consistent with a moderate positive association: after adjusting for pretest achievement, the framework-supported class scored 5.6 points higher, reported higher motivation and satisfaction, and showed a smaller gain in self-directed learning, though the design cannot separate the framework from the extra resources that came with it.

Two conclusions survive those qualifications. Grounding and calibration are not optional refinements: general models produce fluent material about teaching and learning, and fluent material that is subtly wrong is worse than nothing for students who cannot yet tell the difference, so restricting generation to a curated knowledge base and fixing the difficulty of output against a learner profile are what made generation usable as instruction here. And the teacher stays in the loop, which is where the pedagogy lives. The most defensible claim the study supports is not that generative AI teaches, but that a constrained generative workflow changes what a teacher can produce in the time available, at the price of a substantial upfront investment.

Three questions follow. Does the association hold when the volume of material is held constant, as in the ablation study described in Section 5.2? Does it persist beyond the semester and outside the course for which the knowledge base was built? And how should the benefits be distributed when access cannot be universal and a comparison requires that it is not?

Fund Project

This paper is the research result of the general project of the 2024 Scientific Research Project of the Hunan Provincial Department of Education, "Research on Personalized Teaching Resource Design Based on Generative AI Technology" (Project No.: 24C1265).

References

[1] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30, 5998-6008. https://arxiv.org/abs/1706.03762

[2] Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., & Amodei, D. (2020). Language models are few-shot learners. Advances in Neural Information Processing Systems, 33, 1877-1901. https://arxiv.org/abs/2005.14165

[3] Kasneci, E., Sessler, K., Kuechemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Guennemann, S., Huellermeier, E., Krusche, S., Kutyniok, G., Michaeli, T., Nerdel, C., Pfeffer, J., Poquet, O., Sailer, M., Schmidt, A., Seidel, T., Stadler, M., Weller, J., Kuhn, J., & Kasneci, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. Learning and Individual Differences, 103, 102274. https://doi.org/10.1016/j.lindif.2023.102274

[4] Chiu, T. K. F., Xia, Q., Zhou, X., Chai, C. S., & Cheng, M. (2023). Systematic literature review on opportunities, challenges, and future research recommendations of artificial intelligence in education. Computers and Education: Artificial Intelligence, 4, 100118. https://doi.org/10.1016/j.caeai.2022.100118

[5] UNESCO. (2023). Guidance for generative AI in education and research. Paris: UNESCO. https://unesdoc.unesco.org/ark:/48223/pf0000386693

[6] Holmes, W., Bialik, M., & Fadel, C. (2019). Artificial intelligence in education: Promises and implications for teaching and learning. Boston, MA: Center for Curriculum Redesign. https://curriculumredesign.org/our-work/artificial-intelligence-in-education/

[7] Celik, I., Dindar, M., Muukkonen, H., & Jaervelae, S. (2022). The promises and challenges of artificial intelligence for teachers: A systematic review of research. TechTrends, 66(4), 616-630. https://doi.org/10.1007/s11528-022-00715-y

[8] Crompton, H., & Burke, D. (2023). Artificial intelligence in higher education: The state of the field. International Journal of Educational Technology in Higher Education, 20(1), 22. https://doi.org/10.1186/s41239-023-00392-8

[9] Zawacki-Richter, O., Marin, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education: Where are the educators? International Journal of Educational Technology in Higher Education, 16(1), 39. https://doi.org/10.1186/s41239-019-0171-0

[10] Yan, L., Sha, L., Zhao, L., Li, Y., Martinez-Maldonado, R., Chen, G., Li, X., Jin, Y., & Gasevic, D. (2024). Practical and ethical challenges of large language models in education: A systematic scoping review. British Journal of Educational Technology, 55(1), 90-112. https://doi.org/10.1111/bjet.13370

[11] Mishra, P., & Koehler, M. J. (2006). Technological pedagogical content knowledge: A framework for teacher knowledge. Teachers College Record, 108(6), 1017-1054. https://doi.org/10.1111/j.1467-9620.2006.00684.x

[12] Celik, I. (2023). Towards intelligent-TPACK: An empirical study on teachers' professional knowledge to ethically integrate artificial intelligence (AI)-based tools into education. Computers in Human Behavior, 138, 107468. https://doi.org/10.1016/j.chb.2022.107468

[13] Shieh, D. (2009). These lectures are gone in 60 seconds. The Chronicle of Higher Education, 55(26), A1, A13.

[14] Guo, P. J., Kim, J., & Rubin, R. (2014). How video production affects student engagement: An empirical study of MOOC videos. In Proceedings of the First ACM Conference on Learning at Scale (pp. 41-50). New York, NY: ACM. https://doi.org/10.1145/2556325.2566239

[15] Chan, C. K. Y., & Hu, W. (2023). Students' voices on generative AI: Perceptions, benefits, and challenges in higher education. International Journal of Educational Technology in Higher Education, 20(1), 43. https://doi.org/10.1186/s41239-023-00411-8

[16] Cooper, G. (2023). Examining science education in ChatGPT: An exploratory study of generative artificial intelligence. Journal of Science Education and Technology, 32(3), 444-452. https://doi.org/10.1007/s10956-023-10039-y

[17] Bloom, B. S. (1984). The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. Educational Researcher, 13(6), 4-16. https://doi.org/10.3102/0013189X013006004

[18] Xie, H., Chu, H.-C., Hwang, G.-J., & Wang, C.-C. (2019). Trends and development in technology-enhanced adaptive/personalized learning: A systematic review of journal publications from 2007 to 2017. Computers & Education, 140, 103599. https://doi.org/10.1016/j.compedu.2019.103599

[19] Pane, J. F., Steiner, E. D., Baird, M. D., & Hamilton, L. S. (2017). Informing progress:

Insights on personalized learning implementation and effects. Santa Monica, CA: RAND Corporation. https://www.rand.org/pubs/research_reports/RR2042.html
[20] Zhang, L., Basham, J. D., & Yang, S. (2020). Understanding the implementation of personalized learning: A research synthesis. Educational Research Review, 31, 100339. https://doi.org/10.1016/j.edurev.2020.100339
[21] Sweller, J., van Merrienboer, J. J. G., & Paas, F. (2019). Cognitive architecture and instructional design: 20 years later. Educational Psychology Review, 31(2), 261-292. https://doi.org/10.1007/s10648-019-09465-5
[22] Vygotsky, L. S. (1978). Mind in society: The development of higher psychological processes. Cambridge, MA: Harvard University Press.
[23] van Merrienboer, J. J. G., & Kirschner, P. A. (2018). Ten steps to complex learning: A systematic approach to four-component instructional design (3rd ed.). New York, NY: Routledge. https://doi.org/10.4324/9781315113210
[24] Merrill, M. D. (2002). First principles of instruction. Educational Technology Research and Development, 50(3), 43-59. https://doi.org/10.1007/BF02505024
[25] Mayer, R. E., & Moreno, R. (2003). Nine ways to reduce cognitive load in multimedia learning. Educational Psychologist, 38(1), 43-52. https://doi.org/10.1207/S15326985EP3801_6
[26] Black, P., & Wiliam, D. (1998). Assessment and classroom learning. Assessment in Education: Principles, Policy & Practice, 5(1), 7-74. https://doi.org/10.1080/0969595980050102
[27] Hattie, J., & Timperley, H. (2007). The power of feedback. Review of Educational Research, 77(1), 81-112. https://doi.org/10.3102/003465430298487
[28] Zimmerman, B. J. (2002). Becoming a self-regulated learner: An overview. Theory into Practice, 41(2), 64-70. https://doi.org/10.1207/s15430421tip4102_2
[29] Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y. J., Madotto, A., & Fung, P. (2023). Survey of hallucination in natural language generation. ACM Computing Surveys, 55(12), 1-38. https://doi.org/10.1145/3571730
[30] Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Kuttler, H., Lewis, M., Yih, W., Rocktaeschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. Advances in Neural Information Processing Systems, 33, 9459-9474. https://arxiv.org/abs/2005.11401
[31] Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E., Le, Q., & Zhou, D. (2022). Chain-of-thought prompting elicits reasoning in large language models. Advances in Neural Information Processing Systems, 35, 24824-24837. https://arxiv.org/abs/2201.11903
[32] Pintrich, P. R., Smith, D. A. F., Garcia, T., & McKeachie, W. J. (1993). Reliability and predictive validity of the Motivated Strategies for Learning Questionnaire (MSLQ). Educational and Psychological Measurement, 53(3), 801-813. https://doi.org/10.1177/0013164493053003024
[33] Holm, S. (1979). A simple sequentially rejective multiple test procedure. Scandinavian Journal of Statistics, 6(2), 65-70.
[34] Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. Qualitative Research in Psychology, 3(2), 77-102. https://doi.org/10.1191/1478088706qp063oa

Print ISSN: 3104-932X
Online ISSN: 3104-9338

Research on Factual Hallucination Detection and Mitigation in Large Language Models Based on Retrieval-Augmented Generation

Zehua Zhou

Reedy High School, 3003 Stonebrook Parkway, Frisco, TX, USA

Received: October 5, 2026

Revised: October 9, 2026

Accepted: October 19, 2026

Published online: October 22, 2026

To appear in: *International Journal of Advanced AI Applications*, Vol. 2, No. 11 (November 2026)

Corresponding Author: Zehua Zhou (lai@cis.ac.cn)

Abstract. Large language models (LLMs) can generate fluent, persuasive responses that nevertheless contain unsupported, outdated, or contradictory factual claims. Retrieval-augmented generation (RAG) reduces dependence on model parameters by supplying external evidence, but retrieval alone does not guarantee factuality: the evidence may be incomplete, irrelevant, conflicting, poorly ranked, or ignored during generation. This paper develops a transparent, evidence-grounded framework for detecting and mitigating factual hallucinations in RAG-based LLM applications. The framework treats factuality as a claim-level decision problem rather than a single response-level label. It combines query decomposition, hybrid retrieval, evidence quality assessment, atomic-claim extraction, claim-to-evidence verification, uncertainty calibration, source-linked generation, and risk-sensitive response control. A proposed evaluation protocol connects claim precision, evidence coverage, retrieval recall, citation correctness, abstention quality, latency, and human review. Existing methods—including SelfCheckGPT, FActScore, RAGAS, and Chain-of-Verification—are mapped to complementary parts of the pipeline rather than treated as interchangeable solutions. The paper is a framework and evaluation-design study; it does not report new model experiments or claim measured performance improvements. Its main contribution is a reproducible blueprint for testing whether an integrated system can reduce unsupported factual claims while preserving answer usefulness, traceability, and computational efficiency.

Keywords: *Large Language Models; Factual Hallucination; Retrieval-Augmented Generation; Claim Verification; Uncertainty Estimation; Factuality Evaluation*

1. Introduction

1.1. Research Background

Large language models have become general-purpose interfaces for question answering, summarization, analytical writing, coding assistance, and information support. Their fluency can make generated statements appear authoritative even when a claim is inaccurate or not supported by available evidence. Research on hallucination in natural language generation distinguishes several failure modes, including content that conflicts with source material and content that is not supported by the source or reference context [1]. Recent surveys further emphasize that factuality is multidimensional: the same response may contain correct statements, unsupported details, and errors of attribution at once [2].

The issue is especially consequential in knowledge-intensive applications. A system may be asked to explain a policy, summarize a scientific article, compare product specifications, or answer questions about a changing body of documentation. In these settings, factual correctness depends not only on language-model parameters, but also on access to authoritative, relevant, and current information. Benchmarks such as TruthfulQA show why fluent generation cannot be equated with truthfulness: a model may reproduce misconceptions learned from training data rather than provide a reliable answer [3].

Retrieval-augmented generation (RAG) offers a practical path toward grounding answers in external information. Instead of relying entirely on parametric memory, a RAG system retrieves documents or passages and conditions generation on that material [4]. This design can provide traceable evidence and make updates possible without retraining the underlying model. Nevertheless, retrieval changes the failure surface rather than eliminating it. A search component may miss the best source, surface a stale page, rank a distractor highly, or return passages that agree only superficially with the query. A generator may then overstate what the retrieved text establishes, combine facts from different entities, or add unsupported names, dates, numbers, and causal explanations.

Accordingly, the central problem is not simply to retrieve more text. It is to determine which generated claims are supported by which pieces of evidence, detect when the evidence is insufficient or contradictory, and select a response action that reflects the strength of that evidence. This paper proposes a layered framework in which detection and mitigation are coupled: a factuality signal is useful only when it can trigger an appropriate action, such as targeted re-retrieval, claim revision, citation insertion, uncertainty qualification, or abstention.

1.2. Problem Statement and Research Questions

Existing approaches address different fragments of the problem. Sampling-based consistency methods can detect disagreement between independently generated answers without requiring an external database [5]. Atomic-fact evaluation can identify which propositions in a long answer have support from a knowledge source [6]. RAG evaluation frameworks separate retrieval quality from generation faithfulness [7], while verification-oriented generation methods instruct a model to draft, formulate verification questions, answer those questions, and revise its response [8]. These methods offer useful components, but practical systems require a coordinated policy for combining their signals under different evidence conditions and risk levels.

This paper addresses three research questions:

RQ1. How can factual hallucinations in RAG-based LLM responses be identified at the atomic-claim level using retrieved evidence and explicit support labels?

RQ2. How can detection outcomes be converted into mitigation actions that account for source quality, evidence conflict, uncertainty, and the consequences of error?

RQ3. How should an integrated approach be evaluated so that factuality gains are measured alongside retrieval quality, citation validity, answer utility, abstention behavior, and inference cost?

1.3. Contributions and Scope

The paper makes three contributions. First, it defines a practical taxonomy of factual failure modes across retrieval, evidence interpretation, claim generation, and source attribution. Second, it proposes an integrated RAG control pipeline that connects retrieval diagnostics, claim decomposition, evidence verification, and risk-sensitive mitigation. Third, it specifies an evaluation protocol using established factuality and RAG assessment ideas, with explicit baselines, metrics, ablations, and reproducibility requirements. The analysis is grounded in published work [1-10], but it does not introduce an evaluated implementation, new dataset, or newly measured performance result. All experimental steps below are proposed for future empirical validation, not reported as completed tests.

2. Theoretical Background and Related Work

2.1. What Constitutes a Factual Hallucination?

For this study, a factual hallucination is a proposition in a model-generated response that is

false, unsupported by the evidence available to the system, or presented with greater certainty than the evidence justifies. This operational definition intentionally separates truth in the world from support in the system context. A statement may happen to be true but still be unsupported by retrieved evidence; conversely, a response may faithfully summarize a source that itself is outdated or erroneous. Evaluation should record these situations separately rather than merge them into one binary label.

Four categories are particularly relevant to RAG. (1) Retrieval omission occurs when the corpus contains suitable evidence but the retriever fails to return it. (2) Evidence misinterpretation occurs when the correct source is present but the generator draws an invalid inference, confuses scope, or misreads a condition. (3) Unsupported elaboration occurs when a response adds plausible details that are absent from the retrieved evidence. (4) Attribution failure occurs when a citation does not support the adjacent claim, refers to the wrong entity, or points to a source that cannot be independently checked. Contradiction and temporal failure are cross-cutting categories: different sources may disagree, or an otherwise relevant source may no longer describe the current state of affairs.

The distinction between intrinsic and extrinsic hallucination in natural language generation is useful here [1]. Intrinsic failures conflict with the supplied context; extrinsic failures introduce material that cannot be verified from that context. A RAG system may exhibit both. A response that reverses a retrieved date is intrinsically inconsistent with the evidence, while one that appends an unmentioned statistic is an unsupported elaboration even if it sounds plausible. Labeling these cases separately helps determine whether mitigation should focus on better retrieval, stricter evidence use, or external verification.

2.2. Retrieval-Augmented Generation

A conventional RAG pipeline includes document preparation, indexing, query processing, retrieval, context construction, and answer generation. Documents are divided into passages or chunks, transformed into searchable representations, and stored in lexical or vector indexes. At inference time, the question is processed and related passages are retrieved; the selected context is then included in the prompt used by the generator [4]. A hybrid retriever can combine lexical relevance with semantic similarity, while a reranker can reorder candidates using a deeper query–passage relevance model.

The architecture offers two fundamental advantages. It gives a model access to non-parametric knowledge that can be inspected and refreshed, and it creates the possibility of attaching a source to a generated claim. However, retrieval correctness and generation

faithfulness are distinct dimensions. The presence of a relevant passage does not guarantee that the generator will use it correctly. Conversely, a faithful generator cannot produce a well-supported answer if the necessary evidence was never retrieved. Evaluation frameworks such as RAGAS explicitly separate properties of the retrieved context and the generated answer, illustrating the need for multi-dimensional assessment [7].

Longer context is not automatically better. Liu et al. report that language models can use relevant information less reliably when its position changes within a long context, with degraded performance when evidence appears in the middle of the input [9]. Thus, simply increasing the number of retrieved passages may increase token cost and distractor exposure without proportionate factual benefit. A factuality-oriented RAG system should optimize the relevance, diversity, freshness, and sufficiency of evidence—not only the total amount of context retrieved.

2.3. Detection Approaches

2.3.1. Consistency-Based Detection

SelfCheckGPT uses the intuition that repeated samples from a model are more likely to agree on familiar, grounded information than on unstable fabricated details. By comparing sampled passages or sentence-level claims, it estimates the likelihood that content is non-factual without requiring a dedicated external database [5]. This approach is useful where model access is limited or where retrieval cannot be assumed. Its limitations are important: consistent outputs can still repeat the same misconception, and disagreement can arise from stylistic variation or ambiguity rather than a factual error. Consistency should therefore be interpreted as a signal, not a proof of truth.

2.3.2. Atomic-Claim Evaluation

FActScore addresses the mixed quality of long-form answers by decomposing a generation into atomic facts and measuring the proportion supported by a reliable knowledge source [6]. Claim-level analysis is suitable for RAG because an answer may contain both supported and unsupported propositions, making a single answer-level label too coarse. Atomic claims also allow targeted repair: the system can preserve verified statements while rewriting or removing only the unsupported portion. The quality of decomposition and the definition of an acceptable evidence match remain critical, especially for temporal, comparative, numerical, and causal claims.

2.3.3. Claim Verification and Entailment

Fact verification research treats a claim as an object to be checked against evidence, a task exemplified by the FEVER benchmark [10]. In an RAG setting, a claim verifier can classify the relationship between each atomic claim and retrieved passages as supported, contradicted, or not verifiable. The third category must not be conflated with contradiction: absence of evidence may indicate retrieval failure, incomplete documentation, or a genuinely unknowable answer. A verifier should also preserve the source passage and the reasoning trace used for its label so that errors can be audited.

2.4. Mitigation Approaches

2.4.1. Evidence-Grounded Generation

The first mitigation is to constrain generation to information supported by retrieved sources. The prompt can request source-linked claims, prohibit unsupported numbers or names, distinguish direct evidence from inference, and instruct the model to state that evidence is insufficient when necessary. Such instructions improve operational clarity but should not be relied on alone: a model can still overlook instructions or make unsupported inferences. Evidence-grounded generation works best when paired with post-generation verification and an action policy.

2.4.2. Retrieval Repair and Query Refinement

When a claim is unsupported, the system should diagnose whether the required fact was absent from the candidate set, buried in a long context, or present but misinterpreted. If retrieval recall is the likely issue, it can reformulate the query, add entity aliases, use a different retrieval channel, broaden the search, or retrieve from a more authoritative source. Where available, metadata filters should constrain retrieval by publication date, jurisdiction, document version, or source type. Retrieval repair should be bounded by a retry budget so that repeated searching does not cause unbounded latency or amplify low-quality evidence.

2.4.3. Verification Before Finalization

Chain-of-Verification (CoVe) proposes a multi-step procedure in which a model drafts an answer, develops verification questions, answers those questions independently, and then produces a revised response [8]. This pattern can be adapted to RAG by creating one or more check questions for every material claim and resolving them against retrieved passages before finalization. Independent verification is especially useful for dates, quantities, named entities, and multi-part answers, though the verifier itself can make mistakes. Critical claims should

therefore use evidence from sources separate from the initial draft when feasible.

2.4.4. Calibrated Abstention

A dependable system must sometimes decline to provide a direct answer. Abstention may be appropriate when the evidence is missing, sources conflict without a principled way to resolve them, source recency is unclear, or a high-impact claim remains below a chosen support threshold. A useful abstention is informative rather than opaque: it briefly states which part cannot be verified, identifies the evidence gap, and indicates what additional source or clarification would be needed. Because excessive abstention makes a system unhelpful, the evaluation must measure both erroneous answers and avoidable refusals.

2.5. Research Gap

Prior work provides strong building blocks for factuality measurement, retrieval evaluation, and multi-step verification, but several integration challenges remain. First, the retrieval layer and the answer-verification layer are often measured separately, making it difficult to attribute errors to the stage that caused them. Second, a single support score does not express the impact of an error: an unsupported phrase in a casual explanation is not equivalent to an invented legal requirement or medical contraindication. Third, source quality, recency, contradiction, and citation alignment are often treated as secondary diagnostics even though they directly affect whether evidence warrants a claim. Fourth, many evaluations emphasize factuality while leaving latency, token cost, and abstention utility under-specified. The framework below connects these concerns within one auditable workflow.

3. Proposed Evidence-Grounded Framework

3.1. Design Principles

The proposed framework is designed around five principles. Evidence first: important factual statements should be traceable to a source rather than justified only by model confidence. Granularity: checks should operate on atomic claims so that partial errors can be repaired without discarding an entire answer. Separation of roles: retrieval, generation, verification, and response policy should have explicit interfaces. Risk sensitivity: the threshold for answering should reflect the potential consequence of error. Auditability: the system should retain query, retrieved passages, source metadata, claim labels, decision reasons, and final citations in a reproducible record.

These principles do not imply that every sentence requires a costly multi-model verification

pass. Low-risk tasks may use lightweight checks, while high-impact or weakly supported answers receive stronger verification. The system can allocate compute adaptively based on evidence insufficiency, model disagreement, entity ambiguity, temporal sensitivity, and the estimated impact of the claim. This design treats factuality as a controlled workflow rather than a single classifier added after generation.

3.2. End-to-End Pipeline

The pipeline begins with a user question and ends with an answer whose factual claims have been checked, qualified, or withheld. The stages are: (1) question analysis; (2) evidence retrieval; (3) evidence filtering and ranking; (4) grounded draft generation; (5) atomic-claim extraction; (6) claim-to-evidence verification; (7) risk-sensitive mitigation; and (8) citation and response audit. The steps may run iteratively: verification findings can trigger another retrieval round or require the generator to rewrite a specific claim. The system should store stage outputs so that a failed response can be diagnosed after deployment. The stages are summarized in Table 1.

Table 1. Proposed RAG factuality-control pipeline

Stage	Primary output / control
Question analysis	Extract intent, entities, time frame, requested evidence type, and risk level.
Retrieval	Combine lexical and semantic search; preserve document IDs, timestamps, and passage boundaries.
Evidence filtering	Remove duplicates, down-rank stale or low-authority sources, and detect conflicts.
Grounded draft	Generate only source-linked claims; mark inference and missing information explicitly.
Claim verification	Split the draft into atomic claims and label each supported, contradicted, or unresolved.
Mitigation + audit	Re-retrieve, rewrite, qualify, cite, escalate or abstain; validate citations and record decisions/costs.

3.3. Query Analysis and Evidence Retrieval

Query analysis should identify the target entities, requested relation, temporal scope, and the level of specificity required. For example, a question about a school policy should include the institution and the relevant school year; a question about a regulation should include jurisdiction and effective date. When a query is ambiguous, the system may ask for clarification rather than retrieving a large set of loosely related documents. This step reduces entity substitution and accidental blending of similarly named people, organizations, or policies.

Retrieval should use a documented collection of sources appropriate to the task. A hybrid

search strategy combines keyword match—useful for exact names, identifiers, and phrases—with semantic search for paraphrased or conceptually related passages. A reranking stage can prioritize passages that directly address the requested relation. Source metadata should include title, publisher, publication or update date, document version, URL or stable identifier, and where possible an authority label. The system should not treat ranking score alone as proof of truth; ranking estimates relevance, not factual reliability.

Evidence sufficiency should be checked before generation. A retrieval module can ask whether the returned passages cover the question's required slots: who, what, when, where, quantity, condition, and exception. A query with multiple parts may require separate evidence for each part. If only some slots are supported, the generator should answer the supported portion and explicitly qualify the remainder, or trigger another search. Context compression must preserve the wording and qualifiers that determine a claim's meaning; dropping a "may," "except," or effective date can invert the source's implication.

3.4. Atomic-Claim Extraction and Verification

After a draft is generated, a claim extractor divides it into minimal propositions that can be independently checked. "The policy was introduced in 2022 and applies to all students" should become at least two claims: one about the introduction date and one about scope. Compound claims are difficult to verify as a unit because one component may be supported while another is not. The extractor should preserve relationships between claims and the source spans from which the draft was generated.

For each claim, the verifier retrieves or selects the most relevant evidence and assigns one of three primary labels: supported, contradicted, or unresolved. "Supported" means that the cited evidence entails or directly substantiates the claim within the source's stated scope. "Contradicted" means the evidence states an incompatible fact or condition. "Unresolved" means the current evidence neither establishes nor refutes the claim. This three-way decision avoids the false equivalence between absence of retrieved support and demonstrated falsity.

The verifier should evaluate entailment at a semantic level rather than rely solely on lexical overlap. Dates, units, negation, quantifiers, modal verbs, conditionals, comparisons, and entity identity deserve explicit checks. A statement that "some participants improved" does not support "all participants improved"; a source describing a proposal does not establish that it was enacted; and a historical statement may not support a present-tense claim. If the verifier relies on another language model, its output should include a short rationale tied to exact

evidence spans and should be calibrated against human annotations before being used as a hard gate.

A separate citation-alignment check then verifies that every citation points to a resolvable source and supports the nearby proposition. Source existence and claim support are different checks: a valid URL does not imply that the cited page substantiates the statement. The final answer should maintain a mapping from claim IDs to source IDs and evidence spans so that reviewers can trace exactly why each claim was retained, rewritten, or removed.

3.5. Mitigation Policy and Risk Tiers

The mitigation controller combines support labels, source quality, model uncertainty, and estimated impact. Its actions should be deterministic enough to audit while allowing task-specific configuration. A single low-risk unresolved detail can often be removed or qualified. A central claim with weak support should trigger targeted re-retrieval. A contradiction between high-quality sources may require a dated comparison or an explicit presentation of the disagreement. If a high-impact claim remains unresolved, the system should abstain from giving a definitive answer and explain the evidence limitation.

Action A — retain and cite. Use when one or more appropriate sources directly support the claim and no material contradiction is detected.

Action B — qualify. Use when evidence supports a narrower statement than the draft, or when the source is limited by date, population, jurisdiction, or uncertainty.

Action C — repair and verify. Use targeted retrieval or query decomposition when key evidence is missing, then regenerate only the affected claims and repeat verification.

Action D — disclose conflict. Use when credible sources disagree. State the relevant dates or source contexts instead of silently selecting a convenient answer.

Action E — abstain or escalate. Use when evidence is insufficient for a consequential claim, critical citations cannot be validated, or automated checks remain unstable after a bounded retry budget.

Risk tiers should be defined by deployment context rather than by the model alone. A low-stakes brainstorming task may tolerate a looser threshold if uncertainty is clearly signaled. Legal, medical, financial, educational-policy, and safety-related responses require stricter source criteria, more explicit temporal checks, and human oversight when the evidence is ambiguous. The framework is not a substitute for domain-specific governance; it is a technical mechanism that helps surface unsupported claims before the final response is delivered.

3.6. Operational Algorithm

The control logic can be summarized as follows. First, analyze the query and assign an initial risk tier. Second, retrieve a bounded candidate set using hybrid search and rerank the results. Third, assess evidence coverage and source metadata; if coverage is insufficient, refine the query within a fixed retry budget. Fourth, generate a draft that associates factual statements with source identifiers. Fifth, split the draft into atomic claims and verify each claim against the evidence. Sixth, route unsupported or contradicted claims to targeted retrieval, rewriting, qualification, conflict disclosure, abstention, or human review according to the risk tier. Seventh, validate citations and produce an audit record. This sequence should terminate when all material claims pass the configured support threshold or the retry budget is exhausted.

An implementation should distinguish hard constraints from model-generated recommendations. For example, a citation URL that fails to resolve can be treated as a hard failure, while a semantic entailment score may be a probabilistic signal subject to calibration. The policy layer must define minimum evidence requirements and prevent the generator from overriding them merely by restating a claim more confidently. Logs should include model and index versions, retrieval parameters, timestamps, support labels, retry count, and final action to allow reproducibility and incident review.

4. Proposed Evaluation Methodology

4.1. Study Design and Hypotheses

A controlled evaluation is proposed to test the framework against component baselines. The unit of analysis should be a question–answer pair with gold evidence and claim-level annotations where available. The experiment should compare the same underlying generator across different retrieval and verification configurations so that observed differences can be attributed to the pipeline rather than to unrelated model changes. This is a proposed protocol; the present manuscript does not claim that the experiment has been run.

H1. An integrated retrieval-plus-claim-verification pipeline is expected to reduce the proportion of unsupported factual claims relative to generation without post-generation verification.

H2. Evidence-aware mitigation is expected to improve citation support and claim-level factual precision compared with a RAG baseline that provides retrieved passages but does not validate the generated claims.

H3. Risk-sensitive routing is expected to produce a better trade-off between factuality and answer coverage than a single global abstention threshold, although the benefit may vary by task and domain.

H4. Increasing retrieval context size without evidence filtering is not expected to improve factuality monotonically, because long or noisy contexts can introduce distractors and make relevant evidence harder to use [9].

4.2. Data and Task Construction

The evaluation should combine established public benchmarks with a curated RAG test set. TruthfulQA can test whether responses reproduce common misconceptions rather than provide truthful answers [3]. FEVER offers a fact-verification setting with claims and supporting or refuting evidence [10]. FActScore provides a claim-level factuality framework for longer-form generations [6]. RAGAS metrics can be used to inspect dimensions such as retrieval relevance and faithfulness in a RAG pipeline [7]. These resources measure different aspects of behavior and should not be combined into a single score without documenting their definitions and limitations.

A domain-specific RAG set should contain questions whose answers can be determined from a frozen collection of dated, authoritative documents. It should include direct fact lookup, multi-hop questions, temporal questions, questions with conflicting sources, unanswerable questions, and queries containing ambiguous entities. For each question, annotators should record the gold answer or acceptable answer set, gold supporting passages, source dates, required claim slots, and whether the available corpus is sufficient to answer. The corpus and test questions must be split to prevent leakage between tuning and evaluation, and benchmark licensing and terms of use should be followed.

The source collection should be snapshot-based. Each document should have a stable ID, version, capture date, and provenance record. When the task concerns facts that change over time, the gold annotation must define a reference date; otherwise, an updated source could make a correct historical answer appear incorrect or vice versa. Annotation guidelines should specify how to treat partial support, inferential claims, numerical rounding, and multiple acceptable citations. A subset should be independently double-annotated and adjudicated to estimate agreement and reveal ambiguous cases.

4.3. Baselines and Ablation Conditions

At minimum, the proposed evaluation should compare the following systems under the same

questions and generator settings: (1) parametric-only generation without retrieval; (2) standard top-k RAG without post-generation verification; (3) RAG with source citations requested in the prompt; (4) RAG with claim decomposition and a support verifier; (5) RAG with verification and mitigation; and (6) the full risk-sensitive pipeline with targeted re-retrieval and calibrated abstention. Where applicable, SelfCheckGPT-style sampling can be included as a complementary black-box detection baseline [5], while CoVe-style verification can be tested as a multi-step revision baseline [8].

Ablation tests should remove one component at a time: hybrid retrieval, reranking, source-quality filtering, claim decomposition, contradiction detection, citation alignment, re-retrieval, and risk-sensitive routing. Retrieval depth should be varied to test whether additional passages improve coverage or simply increase distractor exposure. A context-order ablation should move relevant evidence from the beginning to the middle and end of the prompt, motivated by evidence that long-context performance can depend on information position [9]. The study should keep the underlying language model, decoding parameters, and source corpus constant where possible.

4.4. Evaluation Metrics

No single metric captures the whole problem. Evaluation should report claim-level factual precision, evidence coverage, contradiction detection, retrieval quality, citation validity, answer utility, abstention quality, and computational cost separately. Automated metrics can reduce evaluation cost, but a manually annotated subset remains necessary because a verifier may share the same biases as the generator and because semantic entailment is difficult to measure perfectly.

Claim factual precision. The number of generated factual claims labeled supported divided by all factual claims judged; also report the unsupported and contradicted fractions separately.

Evidence coverage / recall. The proportion of gold supporting evidence retrieved within the candidate set or final context, reported at defined k values and with duplicate passages handled consistently.

Citation precision. The proportion of citations that resolve and substantively support their linked claims. Citation existence should be reported separately from semantic support.

Verification performance. Precision, recall, F1, and confusion matrix for the supported / contradicted / unresolved classifier, including results by claim type and source quality.

Answer utility. Human or benchmark assessment of relevance, completeness, coherence, and whether the response answers the question without unsupported expansion.

Abstention quality. Coverage at a fixed error rate, risk–coverage curves, and the rate of unnecessary abstention on answerable questions.

Efficiency. End-to-end latency, number of retrieval calls, model calls, context tokens, and estimated inference cost per question.

A useful summary metric may be a risk-weighted unsupported-claim rate, but the weighting scheme must be specified before evaluation and reported alongside unweighted counts. It should not obscure serious failures behind high performance on many trivial claims. Results should include confidence intervals or bootstrap uncertainty where appropriate and should be stratified by question type, source age, evidence conflict, answer length, and risk tier. If automated verification is itself assessed by another language model, a human-audited subset should estimate how often that judge agrees with expert annotators.

4.5. Reproducibility and Validity Controls

To make results reproducible, the study should publish or document the source corpus snapshot, chunking strategy, embedding model, lexical index configuration, retriever and reranker versions, generation prompt, decoding settings, verifier prompt or model, action thresholds, retry limits, and evaluation scripts. Each result should be linked to a system configuration rather than just a model name. If a closed-source model is used, its access date and exposed version identifier should be recorded because behavior can change over time.

Internal validity can be threatened by inconsistent model settings or by using different source collections for different baselines. Construct validity can be threatened when an automatic evaluator treats lexical overlap as factual support, or when a response-level score masks a few severe atomic errors. External validity may be limited if the corpus covers only English web documents or a single domain. A robust evaluation should include multiple task styles and disclose domains where the approach performs poorly. Human evaluators should be blinded to the system condition where practical, and adjudication rules should be written before score comparison.

4.6. Ethical and Data Governance Considerations

Retrieval can expose sensitive or copyrighted material, reproduce low-quality sources, or leak information across access boundaries. A deployment should enforce access controls at retrieval time, respect the license and terms governing indexed documents, minimize retention

of private queries, and avoid revealing restricted source passages in final answers. For sensitive domains, review should consider whether a generated claim could cause harm even if it is technically supported by one source. Provenance and uncertainty do not replace professional judgment or legal and organizational obligations.

The system should avoid presenting a probabilistic verifier label as an objective guarantee. Users should be able to distinguish an answer directly supported by a source from an inference or a response produced under incomplete evidence. When an error is reported, the audit trail should support correction of the source index, verifier rules, and response policy. Monitoring should measure both overconfident answers and excessive refusals, because either failure can undermine user trust and practical utility.

5. Discussion

5.1. Why Retrieval Alone Is Insufficient

RAG is a grounding architecture, not a guarantee of truth. The retriever estimates relevance, whereas the generator decides what to say, and neither operation inherently proves that a claim is factually correct. The system may retrieve a relevant but outdated source, mix facts from two entities, or fail to use a passage that is actually in the context. Long-context evidence also creates a selection problem: increasing the candidate set may improve recall while making evidence harder to attend to or increasing distraction [9]. A robust design therefore needs retrieval quality checks, claim-level verification, and an explicit policy for unresolved evidence.

A further issue is that evidence support and world truth are not identical. An answer may be perfectly faithful to a flawed source. For this reason, the framework treats source authority, date, provenance, and conflict resolution as first-class signals. For changing facts, the system should prefer sources that are authoritative for the relevant jurisdiction or domain and state the date of the information. Where high-quality sources disagree, presenting the disagreement may be more honest than compressing it into an artificially certain conclusion.

5.2. Balancing Factuality, Coverage, and Cost

Stricter verification generally adds model calls, retrieval latency, and engineering complexity. Requiring a citation for every clause can create a visually dense answer without necessarily improving factual precision; requesting several independent checks for every low-risk statement may add cost with little marginal value. The more practical approach is selective verification: devote the strongest checks to claims that are consequential, central to the answer,

unsupported by direct passages, temporally sensitive, numerically specific, or disputed by other sources. Low-risk claims with direct evidence can follow a lighter path, provided the policy is transparent and periodically evaluated.

The system must also balance precision against answer coverage. If it abstains whenever support is imperfect, it may become safe but unhelpful. If it always answers, it may appear useful while propagating unsupported claims. Risk–coverage analysis makes this trade-off visible by plotting error as the system answers a larger fraction of questions. The appropriate operating point depends on task consequences and user expectations rather than a universal threshold. The risk-sensitive controller proposed here is intended to make that choice explicit, tunable, and auditable.

5.3. Failure Modes and Limitations of the Proposed Framework

The framework has several limitations. First, claim extraction may split a statement incorrectly or omit a qualifier that is essential to its meaning. Second, the verifier may produce false support judgments, particularly for multi-hop reasoning, causal claims, tables, or documents containing complex exceptions. Third, retrieval quality depends on corpus coverage and source quality; no verification algorithm can recover evidence that is unavailable and no model can guarantee that every authoritative source is correct. Fourth, risk scores and impact tiers require domain-specific design and may introduce subjective assumptions. Fifth, repeated re-retrieval may create a feedback loop in which the system searches until it finds a passage consistent with its initial answer. A stopping rule and source-diversity policy are therefore necessary.

Sixth, evaluation benchmarks have finite scope. A model may perform well on standard fact-verification datasets yet fail on long answers, recent events, niche topics, multilingual sources, or the interaction of tables and prose. Seventh, judge models can share errors with the generator, creating correlated failure that is invisible when automated scores are treated as ground truth. The proposed protocol addresses these issues through human audits, source snapshots, error stratification, and component ablations, but the framework's effectiveness remains an empirical question until it is implemented and tested.

5.4. Directions for Future Research

Future work should test whether claim support thresholds can be calibrated across domains and model families, and whether independent verifier architectures reduce correlated error. More attention is needed for multilingual retrieval, temporal contradiction, structured data,

mathematical and causal statements, and source provenance across document transformations. Another direction is to train or adapt models to emit claim-level evidence pointers during generation, reducing the gap between answer creation and post-hoc verification. A third is to study human–AI collaboration: interfaces should expose uncertain claims and the evidence gap in ways that support efficient review rather than overwhelming reviewers with alerts.

Further studies should also explore whether source-level reliability estimates can be learned without reinforcing popularity bias, and how to evaluate conflicting sources when no single authority is definitive. Privacy-aware retrieval, permission-sensitive provenance, and defenses against malicious or prompt-injected documents are important operational extensions. These topics are complementary to factuality detection: a system that retrieves adversarial instructions or unauthorized data may generate a claim that looks evidence-supported while violating the intended security boundary.

6. Conclusion

This paper has proposed an evidence-grounded framework for detecting and mitigating factual hallucinations in large language models using retrieval-augmented generation. The central argument is that retrieval is necessary for many knowledge-intensive tasks but is not sufficient for factual reliability. An effective system must evaluate evidence coverage and quality, decompose answers into atomic claims, verify claim-to-source support, check citation alignment, and select mitigation actions according to uncertainty and risk. The framework combines these operations in an auditable pipeline and defines a proposed evaluation protocol that measures factual precision alongside retrieval quality, citation validity, answer utility, abstention, and cost.

The paper is conceptual and methodological. It provides a research blueprint based on established literature and deliberately reports no fabricated dataset, experiment, or performance improvement. Empirical work is needed to determine which components yield the greatest benefit, how their effects vary by domain, and where the cost of verification is justified. By making those questions testable and by preserving a trace from claims to evidence and decisions, the proposed framework aims to support more transparent and dependable RAG-based LLM applications.

References

[1] Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y. J., Madotto, A., & Fung, P. (2023). Survey of Hallucination in Natural Language Generation. ACM

Computing Surveys, 55(12), Article 248, 1–38. https://doi.org/10.1145/3571730

[2] Wang, Y., Wang, M., Manzoor, M. A., Liu, F., Georgiev, G. N., Das, R. J., & Nakov, P. (2024). Factuality of Large Language Models: A Survey. Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, 19519–19529. https://doi.org/10.18653/v1/2024.emnlp-main.1088

[3] Lin, S., Hilton, J., & Evans, O. (2022). TruthfulQA: Measuring How Models Mimic Human Falsehoods. Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), 3214–3252. https://doi.org/10.18653/v1/2022.acl-long.229

[4] Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Kuttler, H., Lewis, M., Yih, W.-T., Rocktaschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. Advances in Neural Information Processing Systems, 33, 9459–9474. https://proceedings.neurips.cc/paper/2020/hash/6b493230205f780e1bc26945df7481e5-Abstract.html

[5] Manakul, P., Liusie, A., & Gales, M. (2023). SelfCheckGPT: Zero-Resource Black-Box Hallucination Detection for Generative Large Language Models. Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, 9004–9017. https://doi.org/10.18653/v1/2023.emnlp-main.557

[6] Min, S., Krishna, K., Lyu, X., Lewis, M., Yih, W.-T., Koh, P., Iyyer, M., Zettlemoyer, L., & Hajishirzi, H. (2023). FActScore: Fine-grained Atomic Evaluation of Factual Precision in Long Form Text Generation. Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, 12076–12100. https://doi.org/10.18653/v1/2023.emnlp-main.741

[7] Es, S., James, J., Espinosa Anke, L., & Schockaert, S. (2024). RAGAs: Automated Evaluation of Retrieval Augmented Generation. Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations, 150–158. https://doi.org/10.18653/v1/2024.eacl-demo.16

[8] Dhuliawala, S., Komeili, M., Xu, J., Raileanu, R., Li, X., Celikyilmaz, A., & Weston, J. (2024). Chain-of-Verification Reduces Hallucination in Large Language Models. Findings of the Association for Computational Linguistics: ACL 2024, 3563–3578. https://doi.org/10.18653/v1/2024.findings-acl.212

[9] Liu, N. F., Lin, K., Hewitt, J., Paranjape, A., Bevilacqua, M., Petroni, F., & Liang, P. (2024). Lost in the Middle: How Language Models Use Long Contexts. Transactions of the Association for Computational Linguistics, 12, 157–173. https://doi.org/10.1162/tacl_a_00638

[10] Thorne, J., Vlachos, A., Christodoulopoulos, C., & Mittal, A. (2018). FEVER: A Large-scale Dataset for Fact Extraction and VERification. Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers), 809–819. https://doi.org/10.18653/v1/N18-1074

Online ISSN: 3104-9338
Print ISSN: 3104-932X

Microwave-Assisted Direct Fluorination of Polymers in Vibro-Fluidized Beds: From Reaction Kinetics to Reactor Design

Yayuan Zheng*, Yang Zhou

Department of Chemical Engineering, Sichuan University of Science & Engineering, Zigong, Sichuan 643000, P. R. China

Received: Oct10, 2026

Revised: Oct 19, 2026

Accepted: Oct 22, 2026

Published online: October 27, 2026

To appear in: *International Journal of Advanced AI Applications,* Vol. 2, No. 11 (November 2026)

*Corresponding Author: Yayuan Zheng (zhengdan0830@163.com)

Abstract. Direct fluorination of polymers is a key surface modification technology for producing high-performance fluorinated materials, but conventional batch and screw reactors face long-standing problems such as low heat transfer efficiency, severe particle attrition, and inconsistent product quality in industrial production. This paper reviews the process intensification strategy of polymer direct fluorination based on the combination of vibro-fluidized bed technology and microwave-assisted heating from the perspective of reaction engineering. The radical chain mechanism and kinetic characteristics of direct fluorination are briefly introduced, and the hydrodynamic advantages of vibro-fluidized beds in improving gas-solid contact uniformity and reducing particle attrition are analyzed. The enhancement mechanism of microwave volumetric heating on reaction kinetics and the self-regulating heating effect caused by the evolution of polymer dielectric properties are discussed. On this basis, a conceptual design of a continuous circulating vibro-fluidized bed microwave fluorination reactor with closed-loop gas purification and recycling is proposed, and its process flow, key structural modules and main operating parameters are systematically described. Finally, the main technical challenges in pilot-scale application, including particle attrition control, process safety and scale-up methods, are summarized, and corresponding future research directions are pointed out. The review can provide a reference for the upgrading and optimization of industrial polymer fluorination processes.

Keywords: *Process intensification; Vibro-fluidized bed; Microwave-assisted heating; Direct fluorination; Computational fluid dynamics; Reaction kinetics; Polymer surface modification.*

1. Introduction

Fluorinated polymers, such as fluorinated polyethylene (FPE), polyvinylidene fluoride (PVDF), and perfluoroalkoxy alkanes (PFA), possess unique properties derived from the strong electronegativity and low polarizability of the C–F bond [1,2]. These materials demonstrate remarkable chemical inertness, excellent thermal stability, superior dielectric properties, and extremely low friction coefficients, rendering them critical in harsh-environment applications including chemical containment, semiconductor fabrication, and advanced membrane technologies [1,3]. The global demand for high-performance fluorinated polymers continues to grow, driving the need for more efficient, sustainable, and controllable fluorination processes [4].

The direct fluorination of polymers using elemental fluorine (F_2) represents the most straightforward method for introducing fluorine atoms into polymer backbones or surfaces [5,6]. This gas-solid reaction proceeds via a radical chain mechanism initiated by homolytic cleavage of the C–H bond and subsequent reaction with molecular fluorine. The strong C–F bond (bond dissociation energy ~485 kJ/mol) formed during fluorination confers the exceptional properties characteristic of fluorinated materials. However, the high reactivity of elemental fluorine and the exothermic nature of the fluorination reaction pose significant challenges for reactor design and process control [6,7].

Traditional industrial fluorination methods predominantly rely on direct fluorination using elemental fluorine (F_2) or fluorinating agents in batch autoclaves or continuous screw reactors [6,8]. Using polymer monomers as raw materials and ammonium persulfate as an initiator, polymerization is initiated followed by hydrolysis. Fluorine gas is introduced into a screw reactor under high temperature and high pressure for fluorination end-capping to finally obtain fluorinated polymers (Figure 1). These approaches typically require elevated temperatures (200–400 °C) and elevated pressures to achieve adequate fluorination degrees. Such severe operating conditions lead to several inherent drawbacks: (i) non-uniform temperature distribution resulting in inconsistent product quality; (ii) significant particle attrition and agglomeration due to mechanical shear; (iii) high energy consumption and safety risks associated with handling toxic fluorine gas at high pressure; and (iv) environmental concerns related to waste gas emissions and by-product formation [9,10]. Consequently, the development of novel reactor technologies capable of simultaneously enhancing reaction efficiency, product quality, and operational safety represents an urgent research priority.

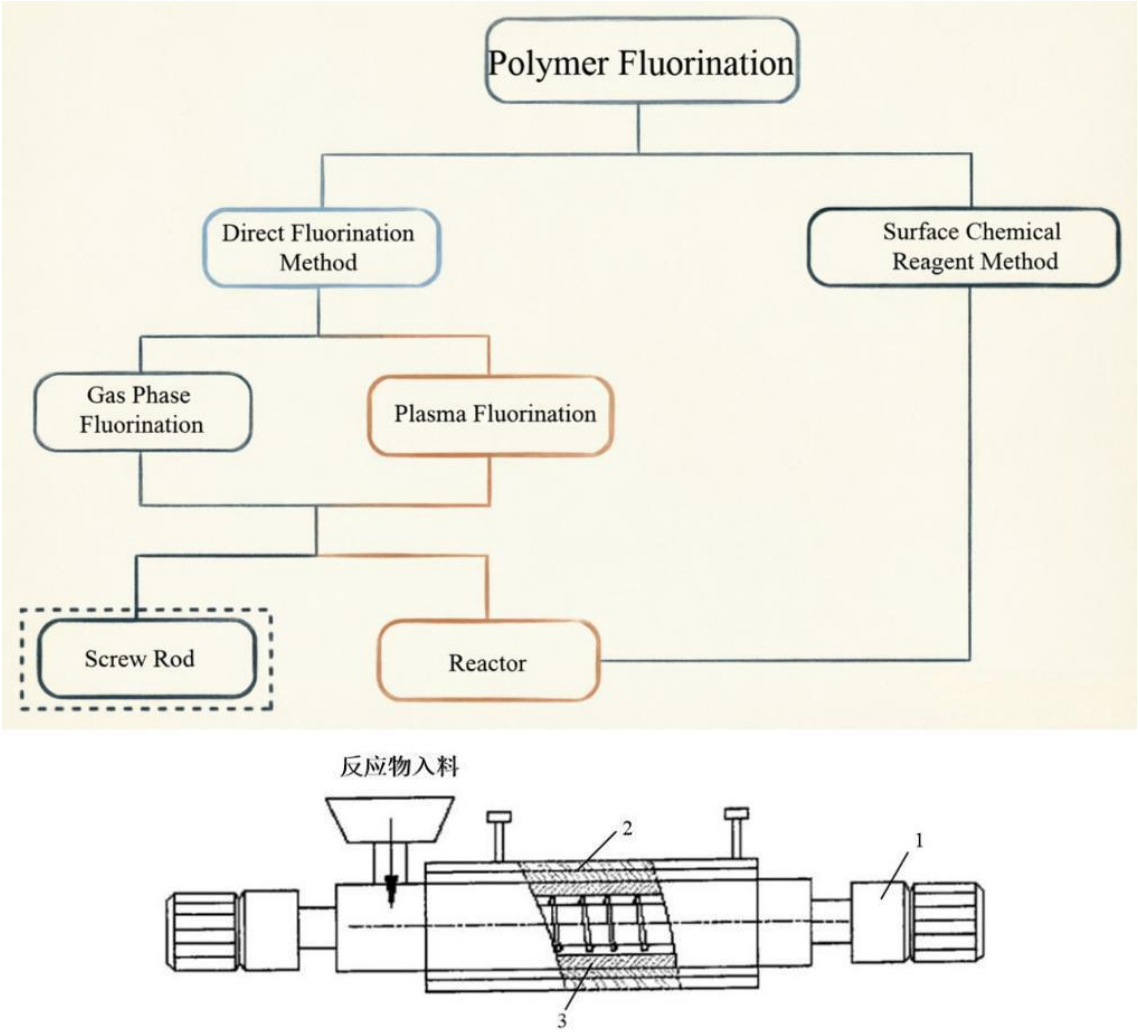


Figure 1. (a) Reactor for the production of fluorinated polymers; (b) Screw reactor (1—Motor; 2—High-temperature heating jacket; 3—Screw shaft).

Fluidized bed reactors have long been recognized as effective gas-solid contacting devices owing to their excellent heat and mass transfer characteristics, uniform temperature distribution, and continuous operation capability [11,12]. Among various fluidization regimes, vibrated fluidized beds (VFBs) offer distinct advantages for processing fine or cohesive particles that are difficult to fluidize conventionally. The application of mechanical vibration reduces the minimum fluidization velocity, suppresses bubble formation, and enhances particle mobility, thereby creating favorable conditions for surface modification reactions [13,14]. More recently, the concept of circulating vibro-fluidized beds, which integrate particle circulation with vibrational energy input, has attracted attention for processes requiring controlled residence times and continuous product removal [15,16].

Concurrently, microwave-assisted chemical processing has revolutionized heterogeneous reaction engineering by enabling rapid, volumetric, and selective heating [17,18]. Unlike

conventional conductive heating, microwave irradiation penetrates particulate materials and generates heat directly within the particles through dielectric loss mechanisms. This characteristic is particularly advantageous for fluorination reactions, where rapid heating can activate the polymer surface, reduce induction periods, and potentially lower the activation energy required for C–H bond substitution by fluorine [18,19]. The combination of microwave heating with fluidized bed technology thus presents a compelling pathway for process intensification in polymer fluorination.

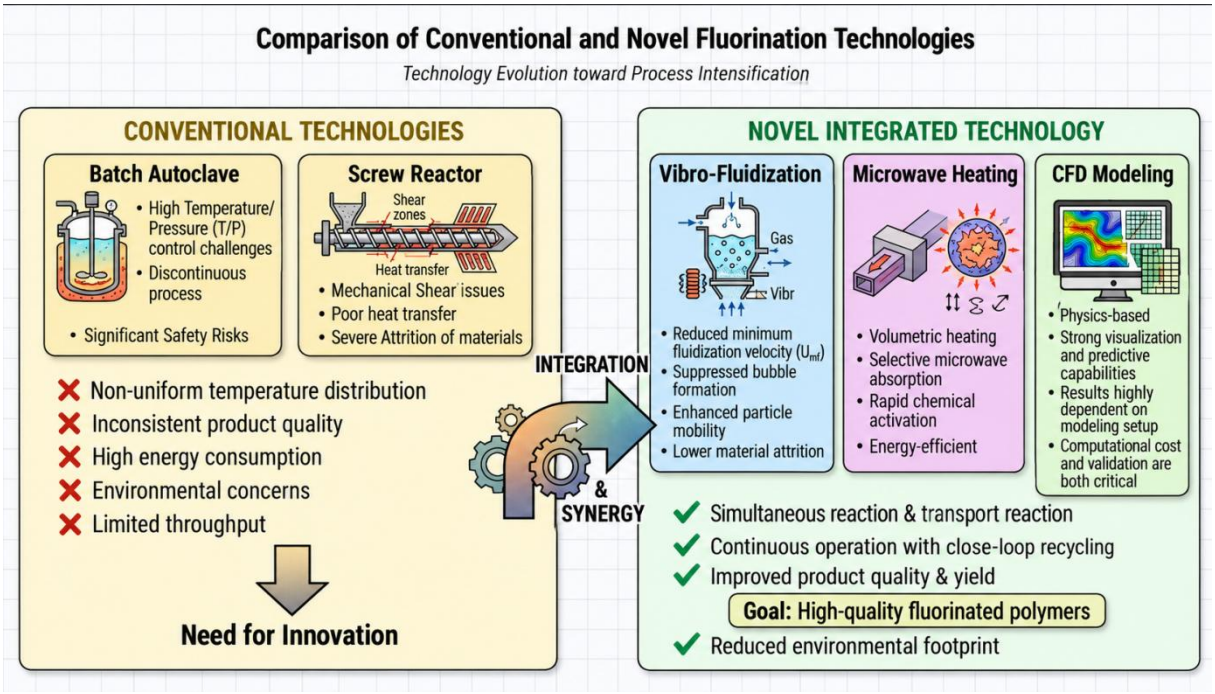


Figure 2. Comparison of conventional fluorination technologies (batch autoclave and screw reactor) and the proposed novel integrated technology.

Despite the promising features of both vibro-fluidization and microwave fluorination, their integration remains largely unexplored, particularly regarding the complex multiphase flow dynamics and reaction-transport coupling phenomena occurring within such hybrid reactors. The mesoscale flow field—characterized by particle clusters, interphase interactions, and localized reaction zones—plays a decisive role in determining reactor performance, product uniformity, and scale-up feasibility [20,21]. Computational fluid dynamics (CFD) modeling, especially coupled Euler-Euler multiphase approaches with surface reaction kinetics, has become an indispensable tool for unraveling these mesoscale phenomena and guiding reactor optimization [21,22]. Figure 2 presents a comparative overview of conventional fluorination technologies versus the proposed novel integrated approach, highlighting the synergistic advantages of combining vibro-fluidization, microwave heating, and CFD modeling. Conventional approaches suffer from non-uniform temperature distribution, inconsistent

product quality, high energy consumption, and environmental concerns. The novel integrated system combines vibro-fluidization (reduced minimum fluidization velocity, suppressed bubbles, enhanced particle mobility), microwave heating (volumetric heating, selective absorption, rapid activation), and CFD modeling (mesoscale prediction, stability analysis, scale-up guidance) to achieve simultaneous reaction and transport intensification, continuous closed-loop operation, improved product quality and yield, and reduced environmental footprint [20,23].

This review aims to provide a comprehensive analysis of the technological convergence between vibro-fluidized beds and microwave-assisted fluorination for polymer surface modification. We examine the fundamental hydrodynamics of vibrated gas-solid systems, the kinetic and thermal characteristics of microwave fluorination, and the current status of multiscale CFD modeling methodologies. Furthermore, we propose a conceptual framework for circulating vibro-fluidized bed microwave fluorination reactors and discuss the critical research needs for establishing operation stability models based on particle attrition theory and for developing coupled CFD-reaction models capable of predicting mesoscale flow fields. This synthesis is intended to guide future experimental and theoretical investigations toward the rational design of next-generation fluorination reactors.

2. Direct Fluorination of Polymers: Reaction Mechanisms and Kinetic Characteristics

2.1. Fundamental Reaction Mechanism

The direct fluorination of hydrocarbon polymers proceeds via a complex radical chain mechanism. The initiation step involves the homolytic cleavage of the C–H bond by elemental fluorine, generating hydrogen fluoride (HF) and a carbon-centered radical [5,6]:

$$R\text{–}CH_2\text{–}R' + F_2 \rightarrow R\text{–}\dot{C}H\text{–}R' + HF + F\cdot (Initiation) \quad (1)$$

The fluorine radical (F·) then propagates the chain reaction by abstracting hydrogen from adjacent carbon atoms:

$$R\text{–}\dot{C}H\text{–}R' + F_2 \rightarrow R\text{–}CHF\text{–}R' + F\cdot\ (Propagation) \quad (2)$$

The fluorine radical can also abstract hydrogen from another polymer chain, propagating the radical chain:

$$F\cdot + R\text{–}CH_2\text{–}R' \rightarrow HF + R\text{–}\dot{C}H\text{–}R' \ (Propagation) \quad (3)$$

The termination step involves radical recombination or disproportionation. The presence of oxygen in the fluorinating mixture can lead to competing oxidation reactions, resulting in the formation of oxygen-containing functional groups such as carbonyl (C=O), carboxyl (–COOH), and hydroxyl (–OH) moieties [6,10].

The fluorination of aromatic polymers, such as polystyrene, involves electrophilic substitution of hydrogen atoms on the benzene ring, as well as addition reactions at the aliphatic backbone. Recent studies have demonstrated that non-destructive hydrogen-substituting fluorination prevails under mild conditions, with fluorination occurring into both the benzene ring and the main carbon chain [7,24].

2.2. Kinetic Characteristics and Activation Energy

The kinetics of polymer direct fluorination are governed by several factors: (i) the concentration of elemental fluorine in the gas mixture; (ii) the reaction temperature; (iii) the polymer crystallinity and morphology; and (iv) the diffusion rate of fluorine through the fluorinated layer. The apparent activation energy for C–H bond fluorination typically ranges from 40–80 kJ/mol, depending on the polymer structure and fluorination conditions [6,9].

The fluorination rate often exhibits an induction period, attributed to the initial formation of fluorine radicals and the establishment of a stable radical chain. Once initiated, the reaction can proceed rapidly, with the fluorination front advancing into the polymer matrix. The thickness of the fluorinated layer (δ) typically follows a parabolic growth law with respect to time (t):

$$\delta = k \cdot t^{(1/2)} \quad (4)$$

where k is the effective rate constant incorporating both reaction and diffusion contributions.

2.3. Effect of Fluorinating Mixture Composition

The composition of the fluorinating gas mixture significantly influences the reaction kinetics and product characteristics. Typical mixtures contain 1–20 vol.% F_2 diluted with inert gases (N_2, He) and may include O_2 for oxyfluorination [25]. Higher fluorine concentrations accelerate the reaction but increase the risk of polymer degradation and thermal runaway. The addition of oxygen promotes the formation of polar functional groups, enhancing surface hydrophilicity and adhesion properties [3,10].

Table 1 summarizes the operating characteristics, product quality, and inherent limitations of conventional fluorination technologies alongside emerging reactor concepts.

Table 1. Comparison of conventional and emerging fluorination technologies.

Technology	Temperature (°C)	Pressure (MPa)	Residence time	Fluorination uniformity	Particle attrition	Energy efficiency	Key limitations
Technology	Temperature (°C)	Pressure (MPa)	Residence time	Fluorination uniformity	Particle attrition	Energy efficiency	Key limitations
Batch autoclave	150–400	0.1–5.0	Batch (hours)	Poor (hot spots)	Low (static bed)	Low (long heat-up)	Discontinuous, safety risks, thermal runaway
Screw reactor	200–350	0.1–0.5	Continuous (min)	Moderate (RTD effects)	High (mechanical shear)	Moderate (wall conduction)	Severe attrition, electrostatic caking
Conventional fluidized bed	150–300	0.1–0.3	Continuous (min)	Moderate (bubble bypassing)	Moderate (bubble eruption)	Moderate (convective)	Elutriation, poor fluidization for fines
Vibro-fluidized bed	100–250	0.1–0.2	Continuous (min)	Good (suppressed bubbles)	Low–Moderate	Improved (lower U)	Limited reactive-system data
Microwave-assisted VFB (projected)	80–200	0.1	Continuous	Excellent (volumetric heating)	Low–Moderate	High (direct heating)	Unvalidated; arcing risks; dielectric data lacking

As evident from Table 1, conventional technologies struggle to simultaneously achieve uniform fluorination, low particle attrition, and high energy efficiency. Fluidized bed reactors, by virtue of their excellent heat and mass transfer characteristics and continuous operation capability, offer a fundamentally different contacting strategy [11,12]. Among fluidization regimes, vibrated fluidized beds (VFBs) are particularly promising for processing fine or cohesive polymer particles that resist conventional fluidization [13,14]. The following section examines the hydrodynamic principles of VFBs and their implications for surface modification reactions.

3. Vibro-Fluidized Bed Technology for Enhanced Gas-Phase Fluorination

3.1. Fundamental Hydrodynamics

Vibro-fluidized beds (VFBs) represent a hybrid gas-solid contacting regime where mechanical vibration is superimposed on conventional gas-fluidization. The vibration is typically imparted to the entire bed or the gas distributor via electromagnetic or eccentric motor drives, with operating frequencies ranging from 5 to 100 Hz and amplitudes from 0.5 to 10 mm [13,26]. The vibrational energy modifies the fluidization behavior of particulate materials by reducing interparticle cohesive forces and lowering the minimum fluidization velocity (U_mf). This enables the fluidization of fine particles (Geldart Groups A and C) that otherwise exhibit channeling, slugging, or defluidization tendencies [14,27].

The hydrodynamic regime in VFBs is governed by the dimensionless vibration acceleration,

$\Gamma=A\omega^2/g$, where A is the amplitude, ω is the angular frequency, and g is the gravitational acceleration. Depending on Γ and the superficial gas velocity (U), VFBs operate in various regimes: fixed bed with vibration, vibrated fluidization, and fully fluidized states. The vibrated fluidization regime, characterized by homogeneous expansion and suppressed bubble formation, is particularly desirable for surface modification reactions requiring uniform gas-solid contact [13,14]. Figure 3 illustrates the hydrodynamic regime map of vibro-fluidized beds, identifying the optimal operating window for polymer fluorination processes. ($\Gamma=A\omega^2/g$) and superficial gas velocity normalized by minimum fluidization velocity (U/U_mf). Four distinct regimes are identified: (i) Fixed Bed with Vibration (low Γ, low U/U_mf), where particles remain essentially static despite vibrational input; (ii) Vibrated Fluidization (intermediate Γ, moderate U/U_mf), characterized by homogeneous expansion and suppressed bubble formation—this is the optimal operating window for surface modification reactions; (iii) Fully Fluidized (high U/U_mf), exhibiting conventional bubbling or turbulent behavior; and (iv) Vibration-Dominated Regime (high Γ), where excessive vibrational energy leads to particle ejection and elevated attrition risk. The optimal operating window is highlighted in green, balancing adequate fluidization quality with minimal particle attrition.

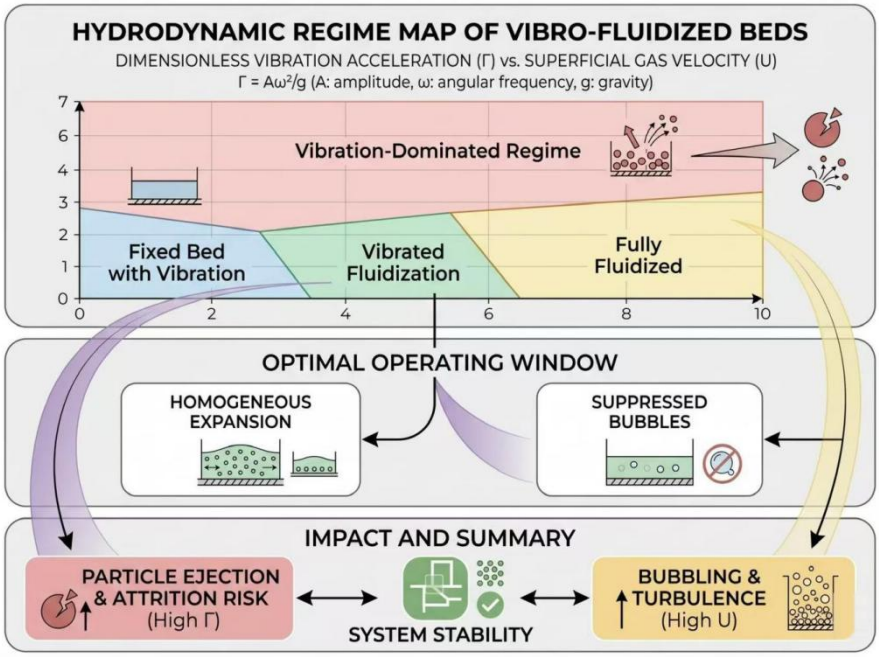


Figure 3. Hydrodynamic regime map of vibro-fluidized beds as a function of dimensionless vibration acceleration.

3.2. Advantages for Surface Modification Processes

For polymer fluorination, VFBs offer several critical advantages over conventional fluidized beds and mechanical reactors. First, the reduced bubble activity and enhanced particle mobility

promote uniform exposure of particle surfaces to the fluorinating gas, minimizing the formation of unreacted cores. Second, the gentle mechanical agitation prevents the severe attrition associated with high-velocity bubbling or turbulent fluidization, thereby preserving particle morphology and reducing fines generation [13,28]. Third, the ability to operate at lower gas velocities decreases the elutriation of fine particles and reduces the gas load on downstream separation and purification equipment.

3.3. Circulating Vibro-Fluidized Bed Concept

The circulating vibro-fluidized bed (CVFB) extends the VFB concept by incorporating a particle circulation loop, wherein particles are continuously fed, processed, and withdrawn while the gas phase undergoes separation, purification, and recycle [15,29]. This configuration decouples the residence time from the gas contact time, allowing independent optimization of reaction duration and gas utilization. For fluorination processes, the circulation loop facilitates temperature control through staged microwave heating, enables continuous product removal, and supports the integration of in-line gas purification systems for fluorine recycling [15,16].

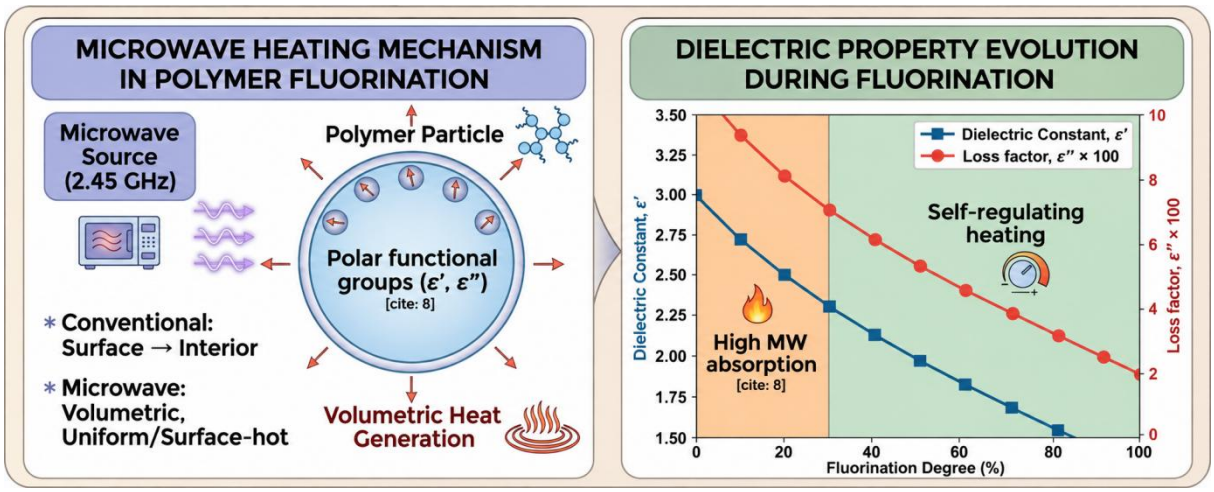
4. Microwave-Assisted Fluorination: Kinetics and Thermal Effects

4.1. Microwave Heating Mechanisms in Polymer Fluorination

Microwave heating relies on the interaction of electromagnetic radiation (typically 2.45 GHz in industrial applications) with polar molecules or conductive materials. In polymer fluorination, microwave energy is absorbed primarily by the polar functional groups on the polymer surface and by any moisture or residual solvents, generating rapid volumetric heating [17,19]. Unlike conventional heating, where thermal gradients develop from the surface inward, microwave heating can establish more uniform temperature profiles or deliberately create surface-hot conditions favorable for initiating fluorination reactions.

The dielectric properties of polymers—specifically the dielectric constant (ε') and loss factor (ε")—determine their microwave absorption efficiency. Most fluorinated polymers exhibit low dielectric loss, meaning that once fluorination begins, the microwave coupling efficiency may decrease, providing a self-regulating heating mechanism that prevents overheating of fully fluorinated regions [17,30]. This selective heating behavior is advantageous for achieving controlled surface fluorination without bulk degradation. Figure 4(a) is the schematic illustration of microwave heating mechanism in polymer fluorination. Microwave irradiation (2.45 GHz) penetrates the polymer particle and generates volumetric heat through interaction

with polar functional groups on the surface (dielectric constant ε', loss factor ε''). Unlike conventional surface-to-interior conductive heating, microwave heating establishes uniform or surface-hot temperature profiles, with heat radiating outward from the particle interior. Figure 4(b) shows the evolution of dielectric properties (ε' and ε'') as a function of fluorination degree. In the initial stage (0–30% fluorination, orange shaded region), high microwave absorption occurs due to abundant polar C–H bonds. As fluorination progresses (30–100%, green shaded region), the decreasing dielectric loss factor provides a self-regulating heating mechanism that prevents overheating of fully fluorinated regions, enabling controlled surface modification without bulk degradation.



(a) (b)

Figure 4. (a) Schematic illustration of microwave heating mechanism in polymer fluorination; (b) Evolution of dielectric properties (ε' and ε'') as a function of fluorination degree.

4.2. Reaction Kinetics Enhancement

Microwave irradiation can influence chemical reaction kinetics through both thermal and non-thermal (specific microwave) effects. In fluorination reactions, the rapid heating rates achievable under microwave conditions (up to 100 °C/min or higher) can significantly reduce the induction period associated with fluorine radical formation and surface activation [18,31]. Studies on microwave-assisted direct fluorination have reported reduced apparent activation energies and enhanced reaction rates compared to conventional isothermal heating, although the extent of non-thermal contributions remains debated in the literature [31,32].

For polymer particles in a fluidized state, microwave heating offers the additional benefit of simultaneous heating of the particle ensemble without relying on gas-phase heat transfer. This is particularly important in vibro-fluidized beds, where the gas flow rates may be insufficient to provide the heat required for endothermic or high-temperature reactions through convective

transfer alone. The coupling of microwave power input with gas flow rate and vibration parameters creates a multi-parameter operating space that can be optimized for simultaneous reaction and transport intensification [18,30].

4.3. Dielectric Property Evolution During Fluorination

A critical aspect of microwave-assisted fluorination is the evolution of polymer dielectric properties as fluorination progresses. In the initial stage (0–30% fluorination), high microwave absorption occurs due to abundant polar C–H bonds. As fluorination progresses (30–100%), the decreasing dielectric loss factor provides a self-regulating heating mechanism that prevents overheating of fully fluorinated regions, enabling controlled surface modification without bulk degradation (Figure 4b). This dielectric evolution has important implications for reactor design and process control, as the microwave power input must be adjusted dynamically to maintain optimal heating rates throughout the fluorination process [17,30].

4.4. Safety and Process Control

Microwave heating in fluorination environments requires careful engineering to prevent arcing, plasma formation, and localized overheating. The presence of fluorine gas, which is highly reactive and can form conductive by-products, necessitates the use of corrosion-resistant microwave window materials (e.g., alumina, quartz) and rigorous gas sealing [30,33]. Additionally, the design of microwave applicators—such as multimode cavities or traveling-wave systems—must account for the load characteristics of fluidized particle beds, which present a dynamically changing dielectric medium. Advanced process control strategies incorporating real-time temperature monitoring (e.g., infrared thermography) and microwave power modulation are essential for safe and reproducible operation [30,33].

5. Multiscale Flow Field Modeling and CFD Approaches

5.1. Overview of Multiphase Modeling Approaches

Two primary CFD approaches are used for modeling gas-solid flows in fluidized beds: Euler-Euler (two-fluid model, TFM) and Euler-Lagrange (discrete element method, DEM). Table 2 compares their key characteristics and applicability to vibro-fluidized bed fluorination reactors.

5.2. Euler-Euler Multiphase Modeling

The computational modeling of gas-solid flows in vibro-fluidized beds presents significant challenges due to the complex interplay between particle-particle, particle-wall, and particle-

gas interactions. The Euler-Euler (two-fluid model, TFM) approach treats both the gas and solid phases as interpenetrating continua, with closure relations derived from kinetic theory of granular flow (KTGF) [21,34]. This framework is computationally efficient for industrial-scale simulations and has been widely applied to conventional fluidized beds.

Table 2. Comparison of Euler-Euler and DEM models for gas-solid flow simulation.

Model Type	Advantages	Disadvantages	Applicability
Euler-Euler (TFM)	Computationally efficient; suitable for industrial-scale simulations; well-established closure models	Cannot directly resolve individual particle collisions; requires empirical drag correlations	Reactor-scale flow field prediction; parametric studies; scale-up analysis
Euler-Lagrange (DEM)	Accurately simulates particle-particle and particle-wall interactions; can model particle attrition and breakage	Computationally expensive; limited to small-scale systems ($\leq 10^6$ particles)	Particle-level mechanism studies; attrition model development; mesoscale structure validation

For VFBs, the TFM must be augmented to account for vibrational effects, typically through modified drag laws, additional body forces representing vibrational acceleration, and enhanced granular temperature dissipation terms [13,35]. The choice of drag model—such as the Gidaspow, Syamlal-O'Brien, or Wen-Yu correlations—significantly influences the predicted bed expansion, pressure drop, and bubble characteristics. Recent studies have emphasized the importance of using drag models that incorporate the effect of heterogeneous structures (clusters) at mesoscales, such as the Energy-Minimization Multi-Scale (EMMS) drag model, as standard homogeneous drag models often overpredict the fluidization quality [36,37].

5.3. Particle Attrition and Stability Modeling

Particle attrition in fluidized beds results from impact, abrasion, and compression forces experienced by particles during bubble eruption, bed expansion, and wall collisions. In vibrated systems, attrition is further influenced by the vibrational intensity and the resulting particle kinematics. Understanding attrition is critical for fluorination processes because excessive fines generation increases the specific surface area uncontrollably, alters the reaction rate, and complicates downstream gas-solid separation [28,38].

Theoretical models for particle attrition in fluidized beds typically combine fracture mechanics principles—such as the indentation fracture theory—with population balance frameworks to predict the evolution of particle size distributions [28,39]. For VFBs, the attrition rate can be correlated with the granular temperature, which quantifies the fluctuating kinetic

energy of particles. By establishing relationships between operating parameters (vibration frequency, amplitude, gas velocity), granular temperature, and attrition rates, it is possible to define "stable operating windows" where attrition is minimized while maintaining adequate fluidization quality [38,40].

5.4. Mesoscale Flow Field Characteristics

The mesoscale in gas-solid fluidization refers to the intermediate scale between individual particle interactions (microscale) and reactor-scale hydrodynamics (macroscale), encompassing structures such as particle clusters, streamers, and bubbles. In VFBs, mechanical vibration suppresses mesoscale heterogeneity by breaking up particle clusters and reducing bubble size, leading to a more homogeneous flow structure [13,41]. This homogenization is beneficial for fluorination reactions, as it reduces gas bypassing and improves interphase contact efficiency.

However, when chemical reactions are introduced, the mesoscale flow field becomes coupled with reaction heat release and species transport. Localized reaction zones may form in regions of high solid concentration or preferential gas flow, creating mesoscale temperature and concentration gradients that are not captured by macroscopic reactor models [20,21]. The prediction of these coupled phenomena requires CFD models that resolve or accurately model mesoscale structures, such as the Energy-Minimization Multi-Scale (EMMS) drag model or filtered TFM approaches [37,42].

5.5. CFD-Chemical Reaction Coupling

The integration of surface reaction models into CFD simulations of fluidized beds enables the prediction of reaction conversion, selectivity, and product distribution alongside hydrodynamic behavior. For fluorination reactions, the surface reaction model must account for the heterogeneous nature of the gas-solid reaction, where fluorine diffuses to the polymer surface and reacts with C–H bonds to form C–F bonds, potentially releasing HF as a by-product [6,7]. In Fluent or similar CFD platforms, surface reactions are typically implemented through user-defined functions (UDFs) that specify the reaction stoichiometry, rate expressions, and Arrhenius parameters [43].

The resulting coupled CFD-reaction model provides detailed spatial distributions of velocity, solid volume fraction, temperature, and species concentrations within the reactor. These distributions constitute the "mesoscale flow field" and are essential for identifying dead zones, short-circuiting paths, and regions of excessive particle heating or attrition. Such insights are invaluable for guiding the geometric design of the reactor, the placement of microwave

applicators, and the selection of operating parameters [21,44].

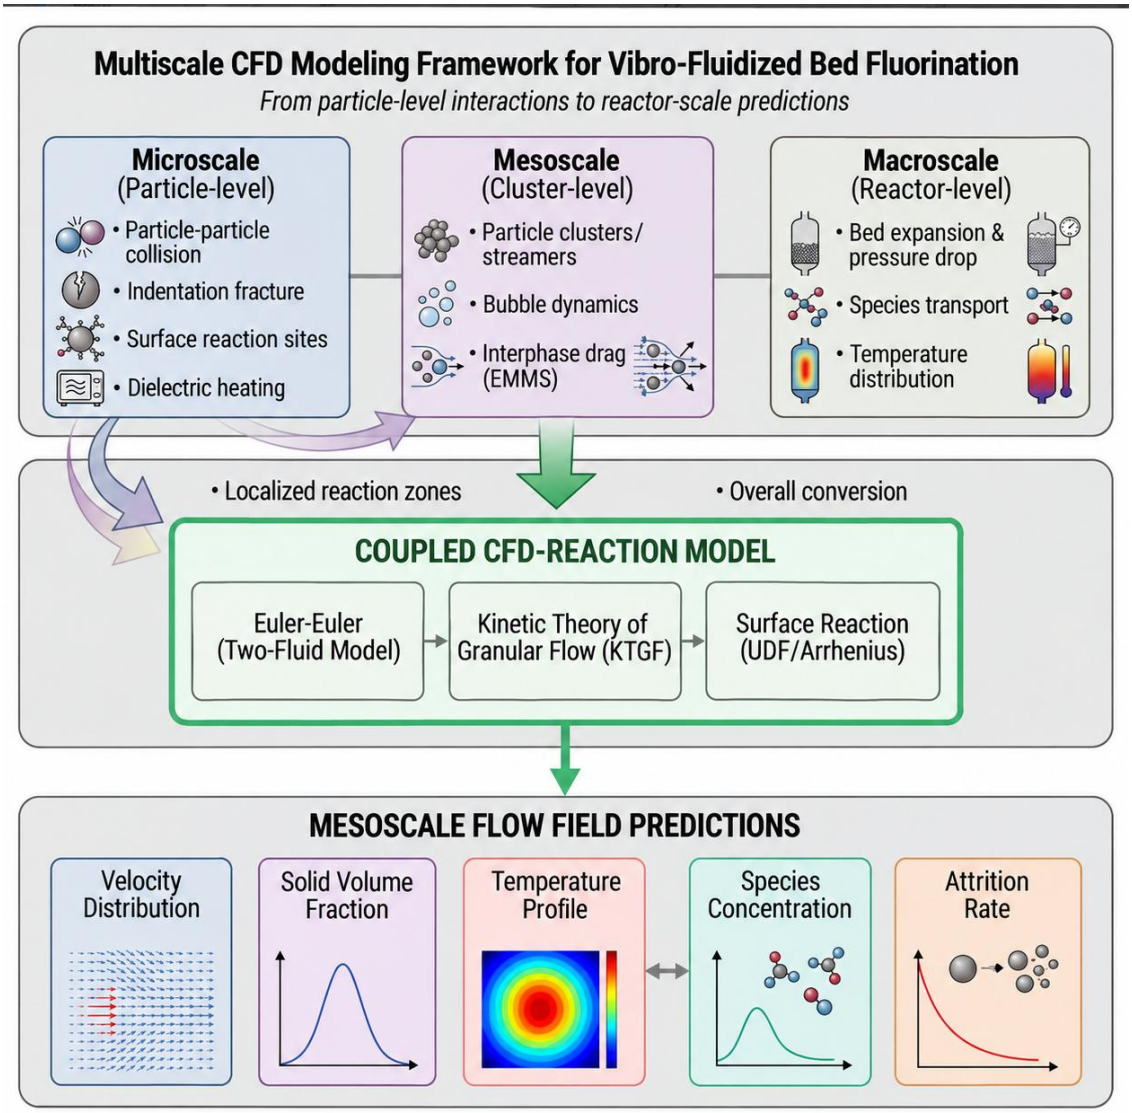


Figure 5. Multiscale CFD modeling framework for vibro-fluidized bed fluorination reactors.

Figure 5 illustrates the multiscale CFD modeling framework that integrates microscale, mesoscale, and macroscale phenomena through coupled Euler-Euler and surface reaction models. The framework spans three hierarchical scales: (i) Microscale (particle-level), encompassing particle-particle collisions, indentation fracture mechanics, surface reaction site distribution, and dielectric heating effects; (ii) Mesoscale (cluster-level), capturing particle clusters and streamers, bubble dynamics, interphase drag via EMMS corrections, and localized reaction zones; and (iii) Macroscale (reactor-level), predicting bed expansion, pressure drop, species transport, temperature distribution, and overall conversion. These scales are coupled through the Euler-Euler two-fluid model (TFM), kinetic theory of granular flow (KTGF) closures, and surface reaction models implemented via user-defined functions (UDFs) with

Arrhenius kinetics. The integrated model outputs critical mesoscale flow field predictions including velocity distribution, solid volume fraction, temperature profile, species concentration, and attrition rate.

6. Integrated Reactor Design: Circulating Vibro-Fluidized Bed with Microwave Fluorination

6.1. Process Flow and System Integration

The proposed integrated reactor system combines the continuous processing capability of circulating fluidized beds with the enhanced fluidization quality of mechanical vibration and the rapid heating benefits of microwave irradiation. The system comprises five integrated modules: (1) Feed System—polymer particles are sized and introduced with inert gas (N_2 or He); (2) Vibro-Fluidization (Pre-treatment Zone)—particles are preheated and homogenized via a vibrated distributor with controlled vibration parameters; (3) Microwave Fluorination (Core Reaction Zone)—fluorine gas is introduced and microwave power (2.45 GHz magnetrons) is applied for surface fluorination with volumetric heating and rapid activation; (4) Separation System (Recovery Zone)—cyclone separator performs gas-solid separation and product cooling; (5) Purification & Recycle (Closed-loop System)—solid scrubbers ($KF/NaF/AlF_3$) remove HF and recover F_2 for recycling [15,45]. Product collection yields high-purity, uniformly fluorinated polymers with preserved morphology. Key operating parameters for optimization include vibration (f: 5–100 Hz, A: 0.5–10 mm, Γ: 1–5), gas flow (U/U_mf: 1.2–3.0, F_2/N_2 ratio), microwave (power: 500–2000 W), and particle properties (d_p: 100–500 μm, residence time).The process flow follows a continuous closed-loop sequence, as illustrated in Figure 6. Polymer particles are first fed into the vibro-fluidization zone where mechanical vibration ensures homogeneous particle distribution and eliminates channeling or slugging. Microwave heating is then activated, and fluorine gas is introduced into the fluidized bed to initiate the surface fluorination reaction. The gas-solid mixture subsequently enters the cyclone separator, where denser fluorinated particles are recovered as product while the lighter gas phase (containing unreacted F_2 and by-product HF) is directed to the purification tower. The solid scrubber mixture ($KF/NaF/AlF_3$) selectively adsorbs HF through chemical reaction, yielding purified fluorine that is recycled back to the reactor via the circulation pump and heat exchanger. This closed-loop design minimizes fresh fluorine consumption and eliminates hazardous waste gas emissions.

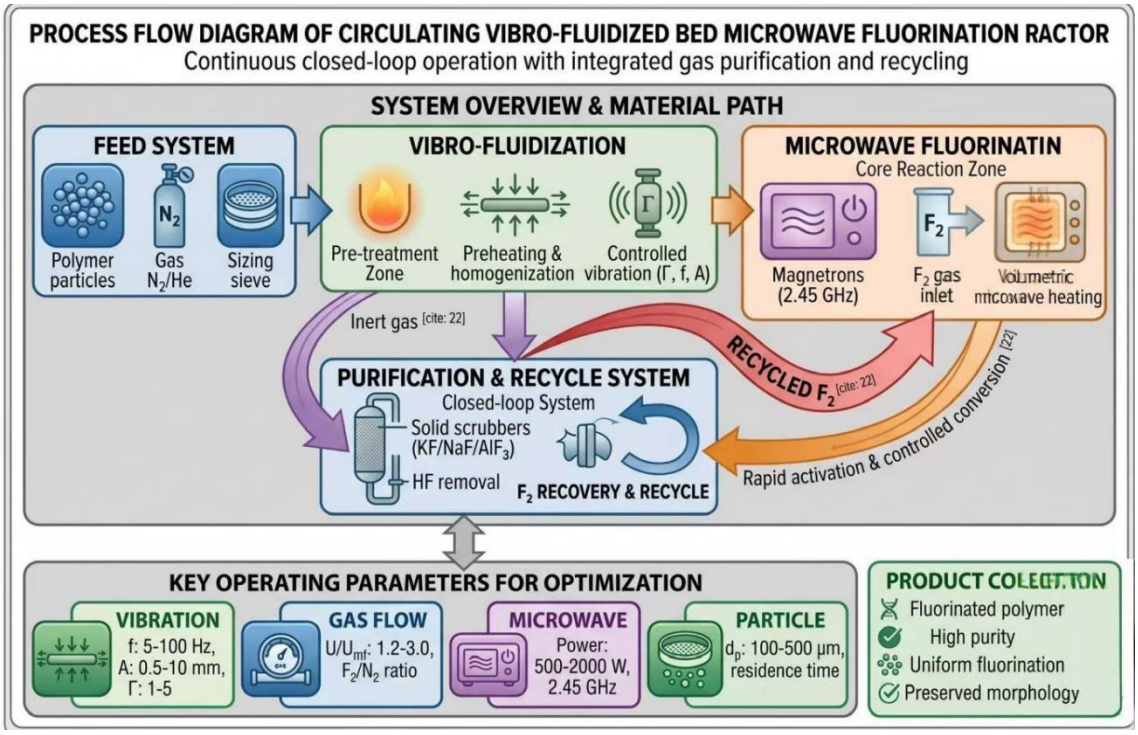


Figure 6. Process flow diagram of the circulating vibro-fluidized bed microwave fluorination reactor.

6.2. Reactor Configuration and Structural Design

Figure 7 presents the detailed structural schematic of the circulating vibro-fluidized bed microwave fluorination reactor, illustrating the spatial arrangement and interconnection of all functional components. As depicted in Figure 7, the reactor system comprises five integrated functional modules that operate in a closed-loop continuous mode. The vibro-fluidization module forms the core reaction environment. Polymer particles are introduced through the feed inlet (1) via the feed pipe (2) into the main reactor body, where they rest on the gas distributor plate (15). The vibration motor (12), mounted beneath the reactor, imparts mechanical oscillation to the entire bed structure, which is supported by spring isolators (13) to minimize vibration transmission to the surrounding infrastructure. This vibration reduces the minimum fluidization velocity and suppresses bubble formation, creating a homogeneous gas-solid contacting regime essential for uniform surface fluorination [13,14].

The microwave heating module is positioned at the top of the reactor. Three microwave generators (3) operating at 2.45 GHz are arranged in series above the quartz window (4), which serves as a critical barrier that allows microwave penetration while preventing the highly corrosive fluorine atmosphere from entering the microwave cavity. The quartz material (typically high-purity fused silica) is selected for its low dielectric loss and excellent chemical resistance to fluorine and hydrogen fluoride at elevated temperatures. Microwave energy

penetrates the particle bed and generates volumetric heat through dielectric loss mechanisms, rapidly activating the polymer surface for fluorination reaction.

The gas distribution and reaction module is integrated within the reactor body. Fluidizing gas (initially inert N_2 or He for preheating, subsequently mixed with F_2 for reaction) enters through the distributor plate (15) and fluidizes the particle bed. The gas outlet (5) at the top collects the reaction gas mixture containing unreacted fluorine and by-product hydrogen fluoride. The observation port (16) allows for visual monitoring of bed behavior and accommodates thermocouples or optical fibers for temperature measurement.

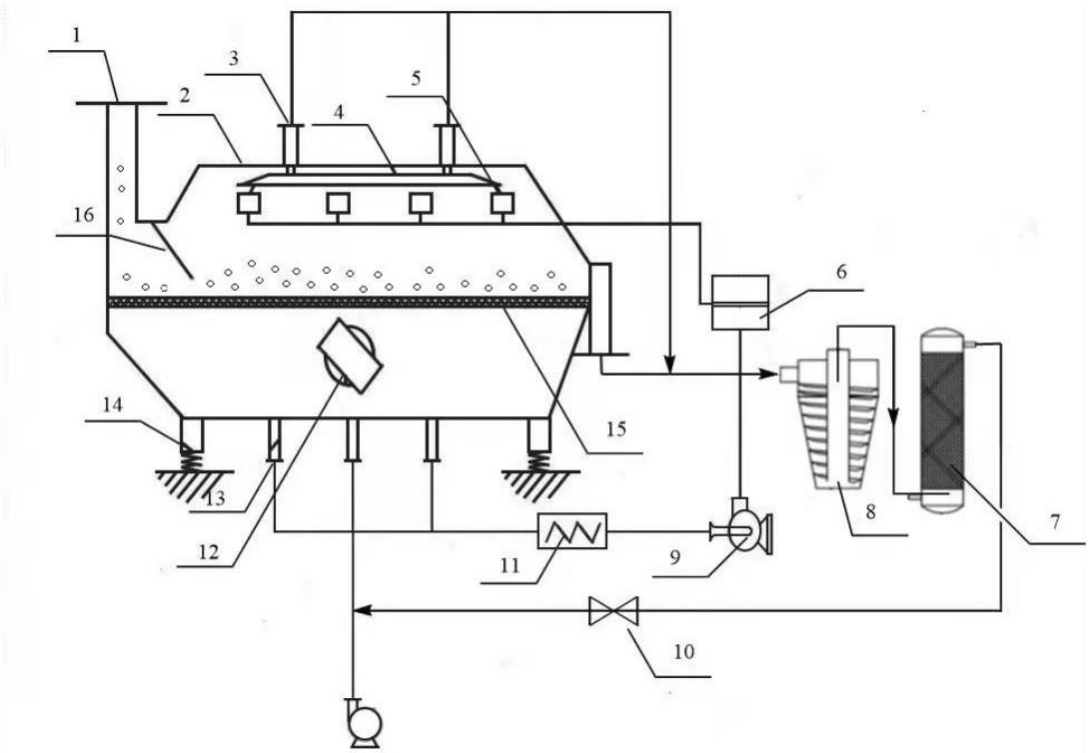


Figure 7. Schematic diagram of the circulating vibro-fluidized bed microwave fluorination reactor. (1) Feed inlet; (2) Gas distributor plate; (3) Microwave generators; (4) Quartz microwave window; (5) Gas outlet; (6) Cyclone separator; (7) Purification tower (packed with $KF/NaF/AlF_3$); (8) Gas recycle pipeline; (9) Circulation pump; (10) Control valve; (11) Heat exchanger; (12) Vibration motor; (13) Spring support; (14) Product discharge; (15) Distributor support; (16) Bed observation port.

The separation and recovery module is located downstream of the reactor. The gas-solid mixture exits the reactor and enters the cyclone separator (6) through the connecting pipeline (8). The cyclone achieves centrifugal separation: denser fluorinated polymer particles are thrown to the wall and collected at the bottom, while the lighter gas phase (containing F_2 and HF) exits through the central vortex finder and proceeds to the purification tower (7). The product discharge (14) enables continuous removal of the fluorinated polymer product, which may be cooled and packaged directly or subjected to further post-treatment.

The purification and recycle module ensures environmental compliance and economic viability. The purification tower (7) is packed with a solid scrubber mixture of potassium fluoride (KF), sodium fluoride (NaF), and aluminum fluoride (AlF_3), which possesses high porosity and specific surface area for efficient HF adsorption. The scrubbing reaction can be represented as:

$$KF + HF \rightarrow KHF_2 \quad (5)$$

The purified fluorine gas, now depleted of HF, is driven by the circulation pump (9) through the heat exchanger (11) and the control valve (10), then returned to the reactor bottom for reuse. The heat exchanger regulates the temperature of the recycled gas, either preheating cold feed gas or cooling hot recycle streams to maintain optimal reaction conditions. The control valve modulates the gas flow rate to maintain the desired superficial velocity relative to the minimum fluidization velocity (typically $U/U_mf = 1.2–3.0$ in the stable operating window) [14,15].

6.3. Multi-Physics Coupling at the Particle Scale

Within the reactor core, three fundamental physical processes occur simultaneously and interactively at the particle scale: microwave heating, gas-solid interphase forces, and particle collision with attrition. Figure 8 presents a local magnification of these coupled mechanisms, which collectively determine the reactor performance, product quality, and operational stability.

Microwave heating (Figure 8a) occurs through the interaction of 2.45 GHz electromagnetic radiation with polar functional groups on the polymer surface. The dielectric loss mechanism converts electromagnetic energy into thermal energy directly within the particle volume, establishing rapid temperature rise rates (up to 100 °C/min) that significantly reduce the induction period for fluorine radical formation. As fluorination progresses, the decreasing dielectric loss factor of the polymer provides a self-regulating heating effect, preventing thermal degradation of the core material while maintaining sufficient surface activation for continued reaction.

Gas-solid interphase forces (Figure 8a) govern the hydrodynamic behavior of the particle bed. The primary forces include: (i) drag force, arising from the relative velocity between gas and particle phases and representing the dominant mechanism for particle suspension and transport; (ii) lift force, generated by velocity gradients and rotation of particles in the shear flow; and (iii) virtual mass force, accounting for the acceleration of surrounding gas as particles accelerate. In vibro-fluidized beds, these forces are superimposed on the vibrational body force, creating complex particle trajectories that enhance mixing and surface renewal while suppressing bubble

coalescence.

Particle collision and attrition (Figure 8b) represent the primary degradation mechanism for polymer particles in the vibrated bed. Three modes of attrition are operative: (i) impact attrition from particle-particle collisions driven by vibrational and fluid dynamic forces; (ii) abrasion

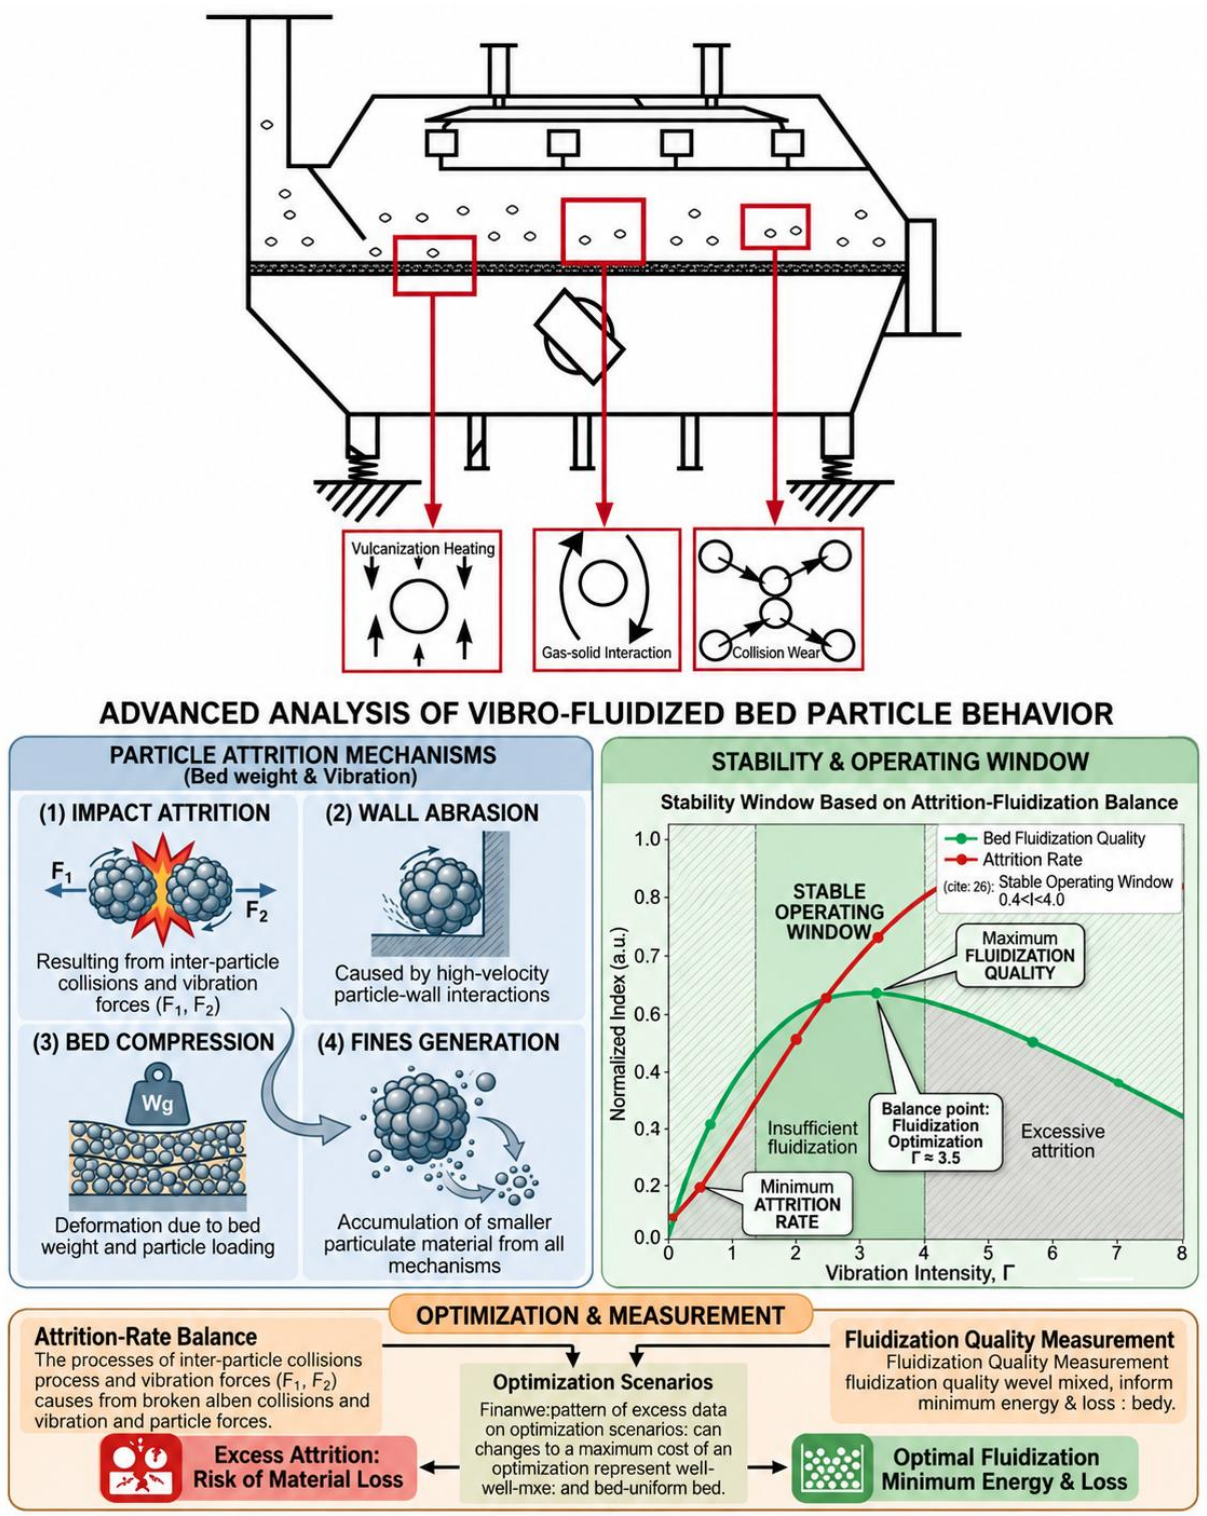


Figure 8. (a) Multi-physics coupling mechanisms at particle scale in microwave fluorination reactors; (b) Attrition modes and stable operating window of vibro-fluidized beds.

from repeated particle-wall contacts at the reactor periphery; and (iii) compression attrition under the static load of the overlying bed. The attrition rate correlates with granular temperature—a measure of particle fluctuating kinetic energy—and must be balanced against fluidization quality to define the stable operating window (see Figure 8).

The coupling among these three mechanisms creates a complex multi-physics environment. Microwave heating rates depend on local particle concentration (affected by fluidization quality), while fluidization quality is influenced by particle size distribution (evolving due to attrition). Meanwhile, attrition rates are temperature-dependent because polymer mechanical properties soften with heating. This three-way coupling necessitates the multiscale CFD modeling framework described in Section 5 for predictive reactor design and optimization [21,40].

6.4. Simultaneous Intensification of Reaction and Transport

The central innovation of this hybrid reactor lies in the simultaneous intensification of reaction and transport (mass, heat, momentum transfer) processes. Vibro-fluidization enhances the external mass transfer coefficient by reducing gas bypassing and improving particle surface renewal. Microwave heating provides rapid, volumetric thermal activation of the polymer surface, reducing the reaction induction time and potentially lowering the overall activation energy. The circulating configuration ensures continuous product removal and fresh surface exposure, preventing over-fluorination and enabling steady-state operation. The synergy between these mechanisms can be quantified through process metrics such as the fluorination rate per unit energy input, the product uniformity index, and the particle attrition rate [7,20].

6.5. Scale-Up Considerations

The scale-up of vibro-fluidized bed reactors from laboratory to industrial scale requires careful consideration of vibration transmission, microwave field uniformity, and gas distribution. Mechanical vibration becomes increasingly difficult to apply uniformly as reactor diameter increases, potentially leading to dead zones or excessive attrition at the walls. Multi-point vibration excitation or segmented vibrator designs may be necessary for large-scale units [12,13]. Microwave scale-up presents analogous challenges, as electromagnetic field uniformity degrades in large multimode cavities. Traveling-wave or slotted-waveguide applicators may offer better scalability than simple cavity designs [30,33].

7. Challenges and Future Research Directions

Despite the promising conceptual and preliminary theoretical framework for circulating vibro-fluidized bed microwave fluorination reactors, several significant challenges remain.

Experimental Validation: The construction and operation of pilot-scale integrated reactors involve substantial engineering complexity, particularly regarding the safe handling of fluorine gas under microwave irradiation and vibrational excitation. Comprehensive experimental campaigns are needed to validate hydrodynamic models, measure reaction kinetics under microwave conditions, and characterize product quality [35,46].

Particle Attrition and Stability: While theoretical attrition models based on granular temperature and indentation fracture theory provide a foundation, the specific attrition behavior of polymer particles under combined vibrational and chemical stress is poorly understood. In-situ particle size and morphology monitoring techniques are needed to develop robust stability models [28,40].

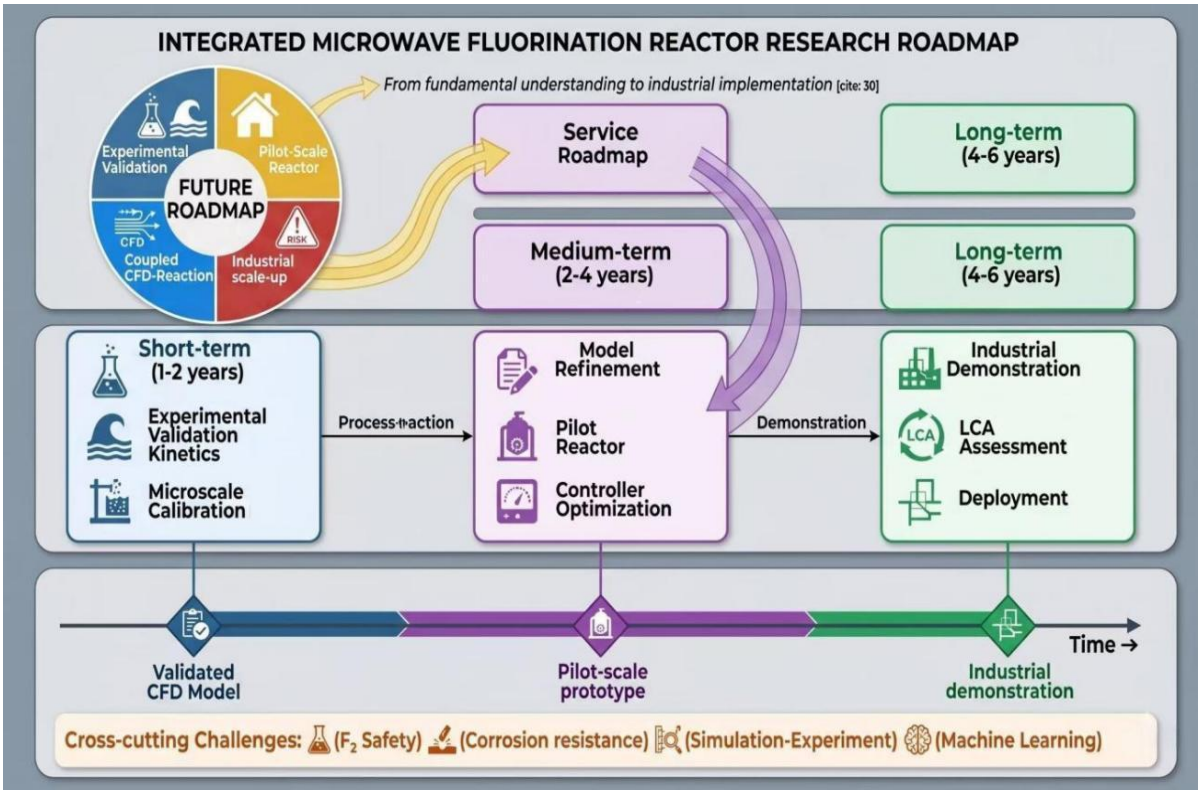


Figure 9. Future research roadmap for integrated microwave fluorination reactors, progressing from fundamental understanding to industrial implementation across three time horizons.

Coupled Model Refinement: Current CFD-reaction coupling approaches rely on simplified global kinetics and homogeneous surface assumptions. Future models should incorporate distributed surface reactivity, multistep reaction mechanisms, and the evolution of polymer

dielectric properties during fluorination. The incorporation of machine learning algorithms for real-time model adaptation based on sensor data represents an emerging opportunity [37,47,48].

Dielectric Property Characterization: Systematic measurement of dielectric constant (ε') and loss factor (ε'') as functions of fluorination degree, temperature, and frequency is essential for optimizing microwave power input and predicting heating behavior. Such data are currently scarce for partially fluorinated polymers [17,30].

Environmental and Economic Optimization: The closed-loop gas purification and recycle system must be optimized for energy efficiency and scrubber regeneration cycles. Life-cycle assessment (LCA) studies comparing the integrated reactor with conventional technologies are needed to quantify environmental benefits and guide industrial adoption [45,47]. Figure 9 outlines the proposed future research roadmap.

8. Conclusion

The development of next-generation fluorination reactors capable of producing high-quality fluorinated polymers with improved efficiency and reduced environmental impact requires the convergence of advanced reactor engineering, novel heating technologies, and multiscale computational modeling. The circulating vibro-fluidized bed microwave fluorination reactor represents a promising paradigm that simultaneously addresses the limitations of conventional batch and screw reactors by enhancing gas-solid contact, enabling rapid volumetric heating, and supporting continuous closed-loop operation.

This review has highlighted the fundamental principles and recent advances in the constituent technologies: direct fluorination reaction kinetics and mechanisms, vibro-fluidized bed hydrodynamics, microwave-assisted reaction kinetics with dielectric property evolution, and Euler-Euler CFD modeling with surface reaction coupling. The integration of these elements into a unified reactor system, supported by particle attrition-based stability models and mesoscale flow field analysis, provides a rational pathway for process intensification in polymer fluorination.

Future research should focus on experimental validation of integrated reactor concepts, refinement of multiscale CFD-chemical reaction models incorporating dielectric property evolution, and the development of scale-up strategies that preserve the mesoscale homogeneity and reaction uniformity achieved at laboratory scale. This review provides a systematic framework for the design of circulating vibro-fluidized bed microwave fluorination reactors from the perspective of fluorine chemistry and reaction engineering.

Acknowledgments

This work was supported by the Sichuan Science and Technology Program (No. 2025NSFSC2070).

References

[1] Belov, N. A. (2021). Effect of direct fluorination on the transport properties and swelling of polymeric materials: A review. Polymers, 13(18), 3094.
[2] Ameduri, B. (2009). From vinylidene fluoride (VDF) to the applications of VDF-containing polymers and copolymers: recent developments and future trends. Chem. Rev., 109(12), 6632-6686.
[3] Doronin, F. A., Rytikov, G. O., Evdokimov, A. G., et al. (2023). Biostable fluorine-containing coatings on the surface of polymers. Coatings, 13(2), 424.
[4] Grand View Research. (2025). Fluorinated polymers market size, share & trends analysis report.
[5] Suchet, N., Batisse, N., Dubois, M., et al. (2025). Surface chemistry stabilization by esterification of polypropylene treated by direct fluorination. ChemPlusChem, 88(8), e202500128.
[6] Belov, N. A., Alentiev, A. Y., Bogdanova, Y. G., et al. (2020). Direct fluorination as method of improvement of operational properties of polymeric materials. Polymers, 12(12), 2836.
[7] Batisse, N., Bonnet, P., Claves, D., et al. (2025). Versatility of gas–solid fluorination for multiple functionalities. C. R. Chimie, 28(G1), 11-25.
[8] Kharitonov, A. P. (2000). Practical applications of the direct fluorination of polymers. J. Fluor. Chem., 103(2), 123-127.
[9] Kharitonov, A. P., Taege, R., Ferrier, G., et al. (2005). Direct fluorination—Useful tool to enhance commercial properties of polymer articles. J. Fluor. Chem., 126(2), 251-263.
[10] Tressaud, A., Durand, E., Labrugere, C., et al. (2007). Modification of surface properties of carbon-based and polymeric materials through fluorination routes. J. Fluor. Chem., 128, 378-391.
[11] Kunii, D., & Levenspiel, O. (1991). Fluidization Engineering (2nd ed.). Butterworth-Heinemann.
[12] Wen, Z. Q., Zhang, X. B., & Luo, Z. H. (2023). A unified CFD methodology for gas-solid fluidized beds based on the quasi-three-fluid approach. Chem. Eng. Sci., 273, 118645.
[13] Lehmann, S., Buchholz, M., Jongsma, A., et al. (2021). Modeling and flowsheet simulation of vibrated fluidized bed dryers. Processes, 9(1), 52.
[14] Mawatari, Y., Tatemoto, Y., & Noda, K. (2003). Prediction of minimum fluidization velocity for vibrated fluidized bed. Powder Technol., 131(1), 66-73.
[15] Ge, X., Zhang, J., Wang, Y., et al. (2023). Modeling and simulation of a multizone circulating reactor for polyethylene production with internal cooling. Polymers, 15(19), 3954.
[16] Li, S., & Shen, Y. (2020). Numerical study of gas-solid flow behaviors in the air reactor of coal-direct chemical looping combustion with Geldart D particles. Powder Technol., 361, 74-86.
[17] Adam, M., et al. (2024). Understanding microwave interactions with polymers to enable advanced plastic chemical recycling. Polymer Testing, 137, 108483.
[18] Kappe, C. O. (2004). Controlled microwave heating in modern organic synthesis. Angew. Chem. Int. Ed., 43(46), 6250-6284.

[19] Mingos, D. M. P., & Baghurst, D. R. (1991). Applications of microwave dielectric heating effects to synthetic problems in chemistry. Chem. Soc. Rev., 20(1), 1-47.
[20] Li, J., Huang, W., & Chen, J. (2021). Possible roadmap to advancing the knowledge system and tackling challenges from complexity. Chem. Eng. Sci., 237, 116548.
[21] Li, J., Ge, W., Zhang, J., et al. (2013). Multi-scale compromise and multi-level coarse-graining in CFD modeling of gas-solid reactors. Chem. Eng. Sci., 81, 65-74.
[22] Sundaresan, S., Ozel, A., & Kolehmainen, J. (2018). Toward constitutive models for momentum, species, and energy transport in gas-particle flows. Annu. Rev. Chem. Biomol. Eng., 9, 61-81.
[23] Song, Y., Zhou, T., Bai, R., et al. (2023). Review of CFD-DEM modeling of wet fluidized bed granulation and coating processes. Processes, 11, 382.
[24] Lagow, R. J., & Margrave, J. L. (1979). Direct fluorination: a "new" approach to fluorine chemistry. Prog. Inorg. Chem., 26, 161-210.
[25] Blinov, I. A., Mukhortov, D. A., Yampolskii, Y. P., et al. (2020). Direct fluorination of poly-2,6-dimethyl-1,4-phenylene oxide in perfluorinated liquid medium. J. Fluor. Chem., 234, 109526.
[26] Bratu, E. A., & Jinescu, G. (1971). Vibrofluidization. Rev. Roum. Chim., 16(8), 1219-1226.
[27] Geldart, D. (1973). Types of gas fluidization. Powder Technol., 7(5), 285-292.
[28] Ghadiri, M., & Zhang, Z. (2002). Impact attrition of particulate solids. Part 1: A theoretical model of chipping. Chem. Eng. Sci., 57(17), 3659-3669.
[29] Zhuang, Y. (2025). Markov-chain-based statistic model for predicting particle movement in circulating fluidized bed risers. Processes, 13(5), 5907.
[30] Metaxas, A. C., & Meredith, R. J. (1983). Industrial Microwave Heating. IET.
[31] Perreux, L., & Loupy, A. (2001). A tentative rationalization of microwave effects in organic synthesis according to the reaction medium, and mechanistic considerations. Tetrahedron, 57(45), 9199-9223.
[32] Stuerga, D., & Gaillard, P. (1996). Microwave athermal effects in chemistry: a myth's autopsy. J. Microw. Power Electromagn. Energy, 31(2), 87-100.
[33] Bradbury, M. J. (2001). Industrial applications of microwave processing. Mater. World, 9(6), 21-23.
[34] Gidaspow, D. (1994). Multiphase Flow and Fluidization: Continuum and Kinetic Theory Descriptions. Academic Press.
[35] Corcoles, J. I., Acosta-Iborra, A., Almendros-Ibanez, J. A., et al. (2021). Numerical simulation of a 3-D gas-solid fluidized bed: Comparison of TFM and CPFD numerical approaches and experimental validation. Adv. Powder Technol., 32(10), 3689-3705.
[36] Tian, Y., Lu, B., Li, F., et al. (2020). A steady-state EMMS drag model for fluidized beds. Chem. Eng. Sci., 219, 115616.
[37] Nikolopoulos, A., Stroggylidis, A., Alobaid, F., et al. (2023). Development of an artificial neural network EMMS drag model for the simulation of fluidized beds in chemical looping combustion. Chem. Eng. Sci., 281, 119424.
[38] Salman, A. D., & Hounslow, M. J. (2004). Particle breakage. Powder Technol., 143-144, 1.
[39] Fulchini, F., Nan, W., Ghadiri, M., et al. (2017). CFD-DEM analysis of particle attrition in a jet in a fluidised bed. EPJ Web Conf., 140, 07017.
[40] Goh, W. P., & Ghadiri, M. (2025). DEM simulation of particle attrition in mechanofusion device. Powder Technol., 456, 120822.
[41] Zhou, T., & Li, H. (2015). Mesoscale structures in fluidized beds. Adv. Powder Technol., 26(3), 673-684.
[42] Wang, W., & Li, J. (2022). Exploring a unified EMMS drag model for gas-solid

fluidization. Chem. Eng. Sci., 258, 117766.
[43] ANSYS Inc. (2021). ANSYS Fluent User's Guide (Release 2021 R2).
[44] Xie, N., Battaglia, F., & Pannala, S. (2008). Effects of using two-versus three-dimensional computational modeling of fluidized beds: Part I, hydrodynamics. Powder Technol., 182(1), 1-13.
[45] Sim, Y., Park, S., Kim, S., et al. (2024). Process intensification for enhanced fluoride removal and recovery as calcium fluoride using a fluidized bed reactor. Int. J. Mol. Sci., 25(9), 4646.
[46] Xi, K., Kovar, T. M., Fullmer, W. D., et al. (2021). CFD-DEM study of bubble properties in a cylindrical fluidized bed of Geldart Group D particles and comparison with prior MRI data. Powder Technol., 389, 75-84.
[47] Paavani, K., Agarwal, K., Alam, S. S., et al. (2025). Advances in plastic to fuel conversion: reactor design, operational optimization, and machine learning integration. Sustainable Energy & Fuels, 9, 54-71.
[48] Farizhandi, A. A. K., Zhao, H., & Lau, R. (2016). Modeling the change in particle size distribution in a gas-solid fluidized bed due to particle attrition using a hybrid artificial neural network-genetic algorithm approach. Chem. Eng. Sci., 155, 210-220.

Online ISSN: 3104-9338
Print ISSN: 3104-932X

Recursive Adversarial Engine: A Stability-Oriented Architecture for Reliable Large Language Model Multi-Agent Systems

Jianhua Fang

Shihao Nine Laboratory

Received: Oct 15, 2026

Revised: Oct 21, 2026

Accepted: Oct 26, 2026

Published online:October 30, 2026

To appear in: International Journal of Advanced AI Applications, Vol. 2, No. 11 (November 2026)

∗ Corresponding Author:

JianhuaFang

(shardylab@sina.com)

Abstract. Large language model (LLM) based multi-agent systems are rapidly migrating from laboratory prototypes to production deployments in finance, scientific research, industrial digital twins, and urban scheduling. As deployment scales, however, a family of systemic pathologies persists: decision oscillation driven by zero-sum agent interactions, cascade propagation of hallucinated content, long-horizon behavioral drift, value misalignment, and unresolved inter-agent conflict. Existing orchestration paradigms, which rely on instruction dispatch, greedy optimization, or static rules, address only shallow coordination and provide no principled mechanism for convergence. This paper proposes the Recursive Adversarial Engine (RAE), a runtime architecture for multi-agent systems organized around three propositions: bounded adversarial interaction supplies the error-detection signal, recursive correction supplies the error-elimination mechanism, and dynamic checks and balances supply the stability guarantee. We specify a four-layer reference architecture comprising adversarial verification, recursive correction, checks-and-balances scheduling, and steady-state output, and we characterize its convergence behavior through a Lyapunov-style stability argument. A qualitative comparison with self-play, multi-agent debate, alignment, and consensus frameworks delineates the paradigm difference. Engineering implications for high-stakes domains, long-horizon robustness, auditability, and middleware integration are analyzed, and open challenges are identified.

Keywords: *Convergence; Hallucination; Autonomy; Checks and balances; Robustness; Interpretability*

1. Introduction

The past three years have witnessed a decisive architectural shift in applied artificial intelligence: from monolithic large language models toward distributed multi-agent systems in which specialized LLM-powered agents collaborate, negotiate, and supervise one another [1]. Financial decision clusters, collaborative research platforms, industrial digital twins, urban traffic orchestration, and autonomous fleet coordination increasingly adopt multi-agent topologies because they decompose complex tasks, isolate failure domains, and exploit heterogeneous capabilities that no single model instance possesses [1,2]. Industry surveys document hundreds of deployed frameworks and a rapidly consolidating engineering practice around role specialization, tool use, and inter-agent communication protocols [1].

This migration, however, has exposed a layer of systemic pathologies that single-model engineering never confronted. An empirical failure taxonomy constructed from more than 1,600 annotated multi-agent traces identifies fourteen recurring failure modes spanning specification errors, inter-agent misalignment, task verification failures, and malformed coordination [3]. Four of these pathologies dominate production incident reports. First, decision oscillation: agents locked in zero-sum bargaining revise proposals indefinitely without converging, consuming compute while degrading output quality. Second, hallucination propagation: a fabricated fact accepted by one agent becomes premise material for its peers, converting an isolated generation error into a system-wide belief [4,5]. Third, long-horizon drift: systems scheduled for uninterrupted 24/7 operation accumulate small biases that individual checkpoints fail to remove, eventually departing from their operating envelope. Fourth, value boundary erosion: agents with growing autonomy progressively overstep policy constraints that static rule engines were expected to enforce.

The economic weight of these pathologies is worth making explicit, because it explains why they resist the incremental fixes that individually look sufficient. Oscillation is paid in latency and compute: an unconverged negotiation loop burns inference budget on every iteration while the user waits. Propagation is paid in trust: a single fabricated statistic that survives into a client deliverable can cost more reputation than the entire system earns in a quarter, and the propagation dynamic makes the cost super-linear in system size, since each additional agent multiplies the surfaces on which an error can land. Drift is paid in operations: deployments that would otherwise run unattended accumulate silent deviations that eventually demand human firefighting at inconvenient hours. Boundary erosion is paid in liability: an agent action that oversteps a compliance envelope converts a technical defect into a legal event. A runtime

architecture that addresses these four costs jointly, rather than one at a time, is therefore not a luxury but the precondition for the multi-agent economics to close.

Mainstream mitigations address these pathologies only partially. Instruction-dispatch orchestrators and greedy task routers resolve shallow coordination but carry no representation of inter-agent conflict. Reinforcement learning from human feedback (RLHF) and its variants align a model to human preferences during training but operate as an external patch that cannot inspect the runtime interaction between autonomous agents [6,7]. Self-play and adversarial training inject conflict productively during training, yet their zero-sum objective incentivizes a winner rather than a stable collective, and their convergence behavior remains fragile in practice [8,9]. Distributed consensus protocols guarantee agreement on shared state but presuppose correct components; they silently replicate a hallucinated value across replicas. What is missing is a runtime architecture in which conflict, correction, and control are designed as a single convergent loop rather than bolted together post hoc.

This paper addresses that gap. We propose the Recursive Adversarial Engine (RAE), a stability-oriented runtime architecture for LLM-based multi-agent systems. RAE rests on three design propositions: (i) bounded, non-destructive adversarial interaction is the principal source of error-detection signal in a closed multi-agent system; (ii) single-pass verification cannot eliminate systematic bias, and recursive, layered correction is necessary for convergence; (iii) long-run stability arises from dynamic checks and balances among agents rather than from centralized static control. These propositions are realized in a four-layer reference architecture—adversarial verification, recursive correction, checks-and-balances scheduling, and steady-state output—each specified at the level of modules, data flow, and interface contracts so that it can be directly implemented on contemporary agent frameworks.

The contributions of this paper are fourfold. First, we formulate a coherent conceptual framework that converts the three propositions into an implementable architecture, with explicit definitions of the deviation signal, the recursive correction operator, and the scheduling policy (Section 3). Second, we characterize the convergence behavior of the correction loop through a Lyapunov-style stability argument, identifying the conditions under which recursive correction decreases system deviation and the mechanisms that guard against its failure modes (Section 3.7). Third, we delineate the paradigm difference between RAE and four incumbent families—self-play, multi-agent debate, alignment pipelines, and consensus protocols—through a structured qualitative comparison (Section 4.5). Fourth, we analyze the engineering implications of the architecture for high-stakes deployment, long-horizon robustness,

auditability, and middleware integration, and we state its limitations candidly: RAE is a conceptual framework with a formal skeleton, and its claims require empirical validation on benchmark workloads before adoption in safety-critical settings (Section 4.6).

The remainder of the paper is organized as follows. Section 2 reviews related work in five strands. Section 3 develops the RAE architecture and its stability analysis. Section 4 discusses engineering significance, comparisons, and limitations. Section 5 concludes.

2. Related Work

2.1. Adversarial Training and Self-Play

Adversarial formulation has repeatedly proven to be a powerful engine of capability. Generative adversarial networks (GANs) established that two networks locked in a minimax contest can implicitly learn rich distributions, and successive refinements—Wasserstein objectives, gradient penalty, and two-time-scale updates—steadily improved the stability of the underlying game [8,10,11,12]. Nevertheless, the GAN literature is equally a catalogue of pathologies: non-convergence, mode collapse, and oscillation between generator and discriminator persist even in well-tuned systems, and formal guarantees are limited to local Nash equilibria under restrictive assumptions [9,12]. Self-play curricula in game playing share the same structure: an agent improves by contesting copies of itself until the opponent pool saturates. Both lines teach an essential lesson that RAE takes seriously—adversarial pressure is a potent source of learning signal, but unbounded or mis-objectived conflict destabilizes the very system it is meant to improve. The zero-sum objective that serves a two-player training game is counterproductive when the objective is a cooperative, long-lived multi-agent deployment: it incentivizes a victor rather than a stable equilibrium of the collective.

A second, less discussed distinction concerns the lifetime of the adversarial relationship. In training games, conflict is a phase: once the opponent pool saturates, adversarial pressure is withdrawn and the artifact ships. In a deployed multi-agent system there is no post-training moment—the collective runs for months, its membership changes, its task distribution drifts, and the adversarial relationship, if it is to keep detecting errors, must be a permanent, institutionalized feature of the runtime rather than a curriculum that ends. None of the incumbent training-time paradigms were designed for permanent conflict under deployment constraints: budgets, auditability, and graceful degradation are not concerns a GAN ever had. The architecture proposed in this paper takes permanent, bounded conflict as its starting design condition.

2.2. Multi-Agent Debate and LLM Collectives

Within the LLM era, the closest neighbors of RAE are debate-style collectives. Irving et al. proposed AI safety via debate, in which two agents argue opposing positions and a judge model evaluates the exchange, with the aim of surfacing flaws that obfuscated single-model outputs would hide [13]. Du et al. demonstrated that having multiple LLM instances propose and critique successive answers improves factuality and reasoning on arithmetic and strategic tasks [2]. Framework surveys catalog a growing zoo of role-played collectives—programmer, critic, planner, executor—that operationalize division of labor with conversational protocols [1]. Two structural limitations recur across these systems. First, the debate terminates by design after a fixed number of rounds, with a judge or majority vote as the arbiter; there is no mechanism that distinguishes residual disagreement caused by genuine task difficulty from disagreement caused by persistent hallucination, and no guarantee that the accepted answer is internally consistent across the collective. Second, participants in the debate are peers with equal standing and static permissions; a participant that repeatedly fabricates evidence exerts the same influence in round ten as in round one. RAE inherits the adversarial-interview insight but embeds it in a recursive loop with per-agent accountability and graded intervention.

2.3. Alignment Pipelines and Safety Layers

A parallel literature hardens individual models against unsafe or unfaithful behavior. RLHF established preference learning from human comparisons as the dominant alignment signal [6]; Constitutional AI replaced part of the human loop with principle-guided self-critique [7]; automated red-teaming scales adversarial probing of model behavior before deployment [14,15]. Safety agendas across the field enumerate concrete hazards—negative side effects, reward hacking, scalable oversight—and unsolved problems that reservation must confront [16-18]. These pipelines are indispensable, but they are properties of a model at training time, not of a fleet at run time. Hallucination surveys document that no current mitigation eliminates fabrication, and that residual errors are precisely the substrate on which multi-agent propagation operates [4,5]. RAE is complementary to this line: it assumes aligned component models and targets the emergent failure modes of their interaction—oscillation, propagation, drift—which training-time alignment cannot reach.

The complementarity has a practical corollary for system builders. Because alignment quality and runtime stability are different properties, neither can substitute for the other in a deployment argument: a perfectly aligned model can still participate in an oscillating collective, and an impeccably scheduled collective still inherits the fabrication tendencies of its members. The

engineering consequence is that acceptance criteria for multi-agent systems need two columns—one for component-level alignment evidence and one for system-level stability evidence—and much of the current difficulty in certifying multi-agent deployments reflects the absence of a standard form for the second column. Providing a candidate for that form is among the intentions of the architecture developed below.

2.4. Stability Theory and Self-Regulating Systems

The third strand supplies RAE's analytical vocabulary. Lyapunov stability theory provides the canonical condition under which a dynamical system returns to an equilibrium: the existence of a scalar energy-like function that decreases along all trajectories [19-21]. Autopoietic and self-organizing systems research describes how living and computational systems maintain identity through continuous internal reorganization, and how apparently chaotic collectives can exhibit phase transitions between order and disorder [22,23]. Recursion itself has been studied since the foundational work of computability theory, with fixed-point theorems guaranteeing that self-referential definitions admit unique stable interpretations under contraction conditions [24,25]. Hofstadter's treatment of strange loops further illuminates how tangled hierarchies resolve into stable levels [26]. RAE imports two ideas from this tradition: the requirement that its correction loop admit a decreasing deviation functional, and the insistence that stability be an emergent property of mechanism design rather than an externally imposed constraint.

2.5. Positioning

Across the four strands, no existing framework simultaneously provides bounded adversarial signal generation, recursive correction with convergence accountability, and dynamic checks-and-balances control as one integrated runtime. Self-play and GANs supply conflict without convergence guarantees suitable for deployment; debate supplies signal without enforcement; alignment supplies constraint without runtime visibility; consensus supplies consistency without error detection. RAE is positioned as the integration layer that closes this gap for LLM-based multi-agent systems.

3. The Recursive Adversarial Engine

3.1. Design Rationale: Three Foundational Propositions

RAE's architecture is derived from three propositions about the nature of error and stability in autonomous multi-agent systems.

Proposition 1: bounded adversarial interaction is the principal negative-entropy source of a

closed multi-agent system. A collective of mutually agreeable agents constitutes a structurally closed loop in which errors circulate unchallenged; politeness gradients suppress dissent exactly where dissent is most informative. Controlled, boundary-respecting adversarial behavior—cross-examination of proposals, mutual consistency interrogation, compliance scanning—exposes latent defects, identifies decision bias, and escapes local optima, supplying the system with corrective signal it cannot generate internally [23]. The qualifier "bounded" is load-bearing: unbounded conflict degrades into the oscillation and attrition documented in the adversarial-training literature [9,12], so adversarial exchanges must be licensed, scoped, and terminated by protocol rather than by agent discretion.

Proposition 2: single-pass verification cannot eradicate systematic deviation; recursive iteration is necessary for convergence. Hallucination and value drift are characterized by concealment and propagation: an error removed from one layer re-enters through another, and a correction applied without re-inspection introduces fresh bias [4,5]. Single-shot cross-checking therefore only trims surface errors. Multi-round recursion—correct locally, verify globally, re-tune locally, confirm global consistency—drives deviation down layer by layer, provided the loop is contracting rather than divergent. This is the same insight that underlies fixed-point semantics of self-referential definitions [25] and that motivates retaining the full recursion path as first-class system state.

Proposition 3: steady-state balance is achieved by dynamic checks and balances, not by static control. Fully centralized control suffocates agent autonomy and re-creates the single point of failure that multi-agent design was meant to remove; unconstrained autonomy dissolves the system into the oscillation of Proposition 1's negative case. The viable middle is mechanism design: permissions, decision weights, and intervention thresholds that adjust continuously to observed agent behavior, so that the system trades autonomy against safety along a controlled frontier [22]. RAE operationalizes this as a scheduling layer with graded authority and graduated circuit-breaking.

The three propositions are not independent axioms but a division of labor. Proposition 1 concerns signal: it explains where corrective information comes from in a closed system, namely from licensed conflict among the agents themselves. Proposition 2 concerns dynamics: it explains how that information must be consumed—recursively, with re-verification—if it is to eliminate rather than merely relocate error. Proposition 3 concerns control: it explains how the first two can operate for the long run without either freezing the system or dissolving it, by making authority itself a state variable that responds to evidence. The four-layer architecture of

Sections 3.3 through 3.6 realizes this division directly: Layer 1 instantiates the signal, Layer 2 the dynamics, Layer 3 the control, and Layer 4 the closure that turns each run's experience into the system's own prior.

3.2. Formal Problem Statement

Consider a multi-agent system whose collective state at discrete time t is a vector s_t comprising agent outputs, intermediate reasoning artifacts, and interaction logs. Let $D: S \rightarrow R_{\geq 0}$ be a deviation functional that aggregates, over the state, the four pathology dimensions introduced in Section 1: factual inconsistency with grounded evidence, cross-agent disagreement, policy-boundary violation, and unresolved conflict. The runtime goal of RAE is to transform an arbitrary operating state into a steady state—a state $s*$ satisfying:

$$|D(s^*)| \leq \varepsilon \tag{1}$$

where ε is the application's tolerance vector, while preserving agent autonomy and task progress. An idealized one-step rendering of the engine's action is the composed operator $s_{t+1} = C(B(s_t))$, where B denotes bounded adversarial verification (deviation measurement) and C denotes recursive correction with checks-and-balances scheduling. The architecture question is how to realize B and C so that the composed map is contracting in the region of interest, i.e., so that there exists a comparison functional $V(s) = \|D(s)\|^2$ whose value decreases along successive iterations:

$$V(s_{t+1}) \leq \gamma \cdot V(s_t), \quad 0 < \gamma < 1 \tag{2}$$

until the stopping condition (1) is met or a safety override fires. Equation (2) is the interface between RAE's engineering layers and classical stability theory [19,21]; Section 3.7 discusses which architectural mechanisms make the inequality plausible and what happens when it fails.

3.3. Layer 1: Adversarial Verification

The verification layer is the engine's sensory apparatus. It deploys four cooperating modules. The cross-adversarial review module pairs agents so that each decision proposal, output artifact, and reasoning path is examined by an independent peer tasked with falsification rather than agreement; the examiner's objective is to construct the strongest refutation, and its reward excludes agreement outcomes, inverting the sycophancy gradient of cooperative collectives. The proposal cross-check module enforces mutual inspection of decision proposals for internal contradiction, unsupported premises, and inferential gaps. The fact-consistency module anchors claims against available grounded sources and against the system's own verified memory,

flagging fabricated or unverifiable content [4]. The boundary-compliance scanner screens every candidate action against the applicable policy envelope—legal constraints, ethical guidelines, permission scopes—classifying borderline behavior for escalation rather than silent passage [7].

The layer's output is a structured deviation inventory: each finding carries a deviation type, affected agents, evidence pointers, a risk grade, and a propagation estimate. This structure is deliberate. First, it converts diffuse conflict into a machine-actionable signal, realizing the transformation from unregulated dispute to controlled verification data. Second, it makes the deviation functional D of Section 3.2 computable at runtime rather than merely definable. Third, it is the primitive on which auditability downstream depends: every accepted deviation becomes a traceable event in the system's evidentiary record.

3.4. Layer 2: Recursive Correction

The correction layer consumes the deviation inventory and operates as a bounded recursion. Each cycle executes four stages: local repair, in which the owning agent revises the flagged artifact under the examiner's findings; global association check, in which a coordination pass re-verifies every state element downstream of the repair, since the deviation inventory's propagation estimate defines the re-verification frontier; secondary local tuning, in which residual inconsistencies surfaced by the global check are patched; and global consistency confirmation, in which the cycle terminates only when the deviation functional over the affected region falls below the local tolerance. Crucially, each recursion round re-inspects the previous round's repairs, so corrections cannot silently introduce new deviations of the kind they removed.

Three design decisions distinguish this loop from a naive retry. First, termination is measured, not counted: the loop exits on deviation collapse (equation (2) satisfied locally), on progress stagnation, or on budget exhaustion—not after a fixed number of chat turns, which is the debated norm in debate collectives [2]. Second, the complete recursion path—the inventory that triggered the cycle, each intermediate repair, each re-verification verdict—is retained as immutable system state. This record is the raw material of interpretability (Section 4.3) and the substrate for post hoc analysis of correction quality. Third, the loop is designed with contraction in mind: by reusing the same verifier family across rounds and requiring monotone inventory shrinkage, the architecture aims at the fixed-point condition under which self-referential correction settles to a unique consistent state [24,25]. Contraction is an engineering target to be monitored, not a guarantee bestowed by design; Section 3.7 makes the failure modes explicit.

3.5. Layer 3: Checks-and-Balances Scheduling

The scheduling layer converts verification outcomes into control actions, implementing Proposition 3's dynamic checks and balances. It maintains, per agent, a continuously adjusted decision weight and permission scope. Weight adaptation is evidence-driven: an agent whose outputs repeatedly attract high-risk deviations, or whose adversarial exchanges exceed the behavioral envelope, has its temporary weight reduced, which degrades its influence on collective decisions and routes its proposals through strengthened verification. Permission adjustment is symmetrical—an agent with a sustained clean record earns wider autonomy—which preserves the incentive for reliable behavior rather than merely punishing deviation.

Conflict arbitration ranks concurrent disagreements by a fixed lexicographic order: safety impact first, responsibility tier second, decision weight third. This ordering ensures that when, for example, a planning agent and a compliance agent deadlock, the compliance constraint prevails regardless of the planners' seniority. The layer also owns the graded circuit breaker. A circuit breaker is triggered when global conflict rate, deviation density, or recursion stagnation crosses its threshold; the response is graduated—throttling non-critical agents, freezing discretionary tasks, rolling back to the last verified state—rather than a binary halt, so that core mission progress and safety invariants are protected at minimal operational cost. Circuit breaking deliberately mirrors established practice in control systems, where protective action must be fast, layered, and reversible [20,21].

fast, layered, and reversible [20,21].

The graded design also reconciles two requirements that a binary kill-switch cannot jointly satisfy: containment granularity and recoverability. Because interventions attach to agents and task classes rather than to the system as a whole, a defecting summarizer is throttled while a serving path continues; and because the response to any trigger is a sequence—throttle, freeze, roll back to the last certified state—an intervention that proves unnecessary is undone by progressing back down the same sequence. Trigger thresholds themselves are adaptive at slow timescales: sustained near-threshold behavior in a task class tightens that class's thresholds preemptively, which is the mechanism through which the scheduling layer contributes to the long-horizon drift control discussed in Section 4.2.

3.6. Layer 4: Steady-State Output

The output layer integrates the corrected state with the scheduling layer's verdicts and emits the system's deliverables. Its obligations are threefold. It produces the final decision or artifact, certified as globally consistent, factually grounded, and policy-compliant with respect to the

deviation inventory. It produces the run's evidentiary bundle: the recursion verification report, the conflict-resolution report, and a robustness score that summarizes deviation trajectory, correction depth, and circuit-breaker activity during the run. And it closes the loop at the temporal scale of a single run: the evidentiary bundle feeds weight adaptation at the scheduling layer and rule refinement at the verification layer, so that each run's correction experience becomes the next run's prior. Table 1 summarizes the four layers, their core modules, and the property each guarantees.

Table 1. The four-layer architecture of the Recursive Adversarial Engine.

Layer	Core modules	Primary input	Primary output	Guarantee contributed
Adversarial verification	Cross-adversarial review; proposal cross-check; fact-consistency; boundary-compliance scan	Agent proposals, outputs, reasoning paths	Structured deviation inventory with risk grades	Detectability of hallucination, conflict, and violation
Recursive correction	Local repair; global association check; secondary tuning; consistency confirmation	Deviation inventory	Repaired state; immutable recursion path	Monotone deviation reduction toward steady state
Checks-and-balances scheduling	Weight adaptation; permission scoping; conflict arbitration; graded circuit breaking	Deviation statistics; conflict telemetry	Adjusted weights, scopes, interventions	Dynamic equilibrium of autonomy and safety
Steady-state output	Result integration; report generation; robustness scoring	Corrected state; scheduling verdicts	Certified deliverables; audit bundle; feedback to Layers 1 and 3	Global consistency, traceability, run-level closure

3.7. Convergence Behavior and Failure Modes

The stability argument for RAE is comparative and conditional, and it is worth stating precisely what is and is not claimed. What is claimed: the architecture instantiates the structure of a stabilizing feedback system, in the sense of Lyapunov's classical criterion [19]. The deviation functional $V(s) = \|D(s)\|^2$ of equation (2) plays the role of an energy function; the verification layer provides continuous measurement of V; the correction layer supplies a repair operator whose cycle semantics require inventory shrinkage before termination; and the scheduling layer bounds the influence of any component whose behavior would drive V upward. Under two assumptions—(A1) verifier judgments are more often right than the agents they

judge, so that each repair reduces the true deviation in expectation; and (A2) correction resources are sufficient to reach the affected region before propagation outruns repair—each recursion cycle is expected to contract V, and the system drifts toward the steady state of equation (1) [20,21]. Assumption A1 is the epistemic load-bearing wall: it is supported operationally by grounding the fact-consistency module against external evidence rather than agent consensus [4], and by the lexicographic precedence of compliance arbitration, but it is an empirical property of a deployment, not a theorem.

The failure modes are correspondingly explicit. If verifiers systematically err (A1 violated), the loop can converge to a confidently wrong steady state; the audit bundle is designed precisely to make this detectable post hoc, and weight adaptation shrinks the influence of agents whose findings are later overturned. If propagation outruns correction (A2 violated), the graded circuit breaker freezes discretionary activity before the deviation spreads system-wide, trading availability for integrity. If adversarial pressure exceeds the behavioral envelope, oscillation re-emerges; the boundary scanner and conflict-rate threshold catch it and throttle the exchange. Finally, the engine itself consumes compute; recursion depth and arbitration frequency are governed by explicit budgets so that stability expenditures remain bounded. These mechanisms do not prove convergence—they bound the modes of divergence, which for a deployable system is the practically meaningful guarantee [9,12].

practically meaningful guarantee [9,12].

It is useful to relate this stability picture to the detect-respond discipline of mature reliability engineering. The verification layer implements detection with coverage obligations—every certified output must carry a deviation inventory, so undetected deviation can pass only through verifier failure, never through absence of testing. The correction and scheduling layers implement response with measured escalation, and the audit bundle implements post-incident analysis with complete evidence. The mapping is deliberate: it means RAE's operational health can be monitored with the same service-level vocabulary—detection latency, correction success rate, breaker activation rate—that production systems already use, and its stability claims are therefore falsifiable through ordinary operational telemetry rather than through bespoke evaluation alone [16,17].

The fourth mechanism addresses the cost side of the ledger. Recursive correction multiplies inference expenditure precisely when the system is least healthy, and an unpriced loop can turn an incident into an incident plus a budget overrun. RAE therefore attaches budgets to recursion at two granularities: per-cycle token and call budgets enforced by the scheduling layer, and per-

run cumulative budgets whose exhaustion converts the run from corrective to conservative mode—delivering with an explicit uncertified annotation rather than consuming unlimited resources. The annotation pathway preserves the honesty property of Section 3.6 under scarcity: when the engine cannot afford to verify, it says so on the artifact instead of pretending the risk away.

4. Engineering Implications and Discussion

4.1. Safety Envelopes for High-Stakes Domains

The first deployment dividend of the architecture appears in high-reliability domains: quantitative finance, clinical decision support, legal drafting, and public-sector approval workflows. In these settings the binding constraint is not capability but error tolerance: regulatory regimes demand that decision error, fabrication, and conflict rates remain inside auditable bounds [16,18]. RAE's contribution to this constraint is architectural rather than statistical. The deviation inventory gives risk a runtime representation with explicit grades; the recursive correction loop drives residual deviation toward the tolerance envelope before output certification; and the steady-state output refuses certification when the envelope is not met, converting a silent quality failure into an explicit, logged refusal. The comparison with incumbent pipelines is direct: RLHF-style alignment shapes component models before deployment [6,7], but the acceptance decision at run time still requires evidence that this particular output is deviation-free; RAE supplies exactly that evidence chain.

Concretely, consider investment-research report generation under a suitability regime. The verification layer's boundary scanner enforces that every performance claim carries a computation basis and that restricted-product language cannot reach the draft stage; the recursion loop requires that a corrected valuation be re-derived, not merely re-worded; and the steady-state output certifies the report together with its deviation inventory, so the compliance function receives, for each document, the list of what was found and what was corrected. From the regulator's perspective the difference is qualitative: instead of a vendor attestation that models were aligned in training, the institution can produce per-artifact process evidence—finding, repair, re-verification—of the kind audit standards already presuppose [6,18].

4.2. Long-Horizon Robustness

The second dividend concerns systems scheduled for uninterrupted operation: industrial IoT coordination, urban scheduling, communication network optimization. Long-horizon deployments fail not through any single catastrophic error but through drift—small biases that

pass each individual checkpoint yet compound across operating hours, a dynamic reminiscent of self-organized criticality in which accumulating stress releases through avalanches of varying size [23]. RAE counteracts drift through two mechanisms with complementary time constants. At the fast timescale, every run's recursion loop removes deviation down to tolerance before certification, so each delivered artifact starts from a verified base. At the slow timescale, the evidentiary bundles accumulate into the scheduling layer's statistics, and sustained patterns—rising correction depth for a given agent, increasing conflict density in a given task class—trigger preventive weight and permission adjustments before thresholds are breached. The engineering claim is that this converts drift from an unbounded random walk into a regulated process with observable early-warning indicators, thereby reducing emergency intervention and manual oversight cost. Again, the claim is architectural; validating it requires longitudinal field measurement, which we flag in Section 4.6.

which we flag in Section 4.6.

A metropolitan scheduling cluster makes the two-timescale mechanism concrete. Hourly, signal-timing and transit-consolidation proposals pass through adversarial verification against live sensor feeds; a misparsed congestion reading is caught as factual inconsistency before it contaminates a corridor plan, corrected in one recursion cycle, and recorded. Weekly, the scheduling layer aggregates the correction record: if one perception agent's findings are increasingly overturned, its weight decays before its errors compound; if a task class shows rising correction depth, its verification rules are tightened. Neither timescale requires the system to halt. The design intent is that drift is consumed continuously at the run level and structurally at the season level, so the accumulation dynamics that produce avalanche-style failures [23] never find an unmonitored interval in which to operate.

4.3. Interpretability and Auditability

The third dividend addresses the black-box problem that currently blocks adoption in audited industries. Interpretability is usually pursued by explaining a model's internal computation, an approach whose outputs are themselves unverifiable generations. RAE relocates interpretability from model introspection to process evidence: because every deviation finding, every recursive repair, every arbitration decision, and every circuit-breaker activation is written to the immutable run record, the system's behavior is reconstructible as a causal narrative without interpreting any neural internals. An auditor can answer the questions that regulation actually asks—why was this output certified, what deviations were found, who intervened, what was overridden—entirely from the record [7,14]. The robustness score attached to each run gives

oversight bodies a single comparable metric across deployments. This evidence-first stance parallels the direction of the AI-safety field toward scalable oversight and evaluation: rather than demanding that a model explain itself, demand that the process leave evidence [17,18].

demand that the process leave evidence [17,18].

The practical shape of such an audit is worth sketching. Asked why a loan-advisory output was certified, the institution does not interrogate a model; it replays the run record: the deviation inventory attached to the certified state, the two high-risk findings (an unsupported income inference and a policy-boundary approach), the recursion path that repaired the first and re-derived the second, the arbitration decision that resolved a disagreement between the compliance and advisory agents, and the robustness score contextualized against the deployment's history. Every element is a logged system event rather than a generated explanation, so the narrative is reconstructible from evidence that existed before the question was asked—an epistemically sturdier position than post hoc rationalization [14].

4.4. Middleware Integration and Ecosystem Effects

The fourth dividend is integrative. RAE is specified as a runtime layer with narrow interfaces—deviation inventories in, certified states and verdicts out—that do not presuppose a particular agent framework, model provider, or communication protocol. This positions the engine as middleware: heterogeneous LLM agents, open-source agent frameworks, and domain-specific modules can be attached beneath the verification layer without modification, and receive scheduling and certification services from above. The practical consequence is a standardization path for organizations that currently face high bespoke integration costs when composing multi-agent systems from diverse components [1,3]. Ecosystem-wise, uniform deviation semantics is the deeper prize: once verification findings, arbitration records, and robustness scores share a schema, cross-organizational benchmarking of reliability becomes possible in a way that isolated deployments have not supported, and failure knowledge can accumulate across the community rather than within single engineering teams [3].

The middleware posture also carries a governance interface. Because deviation grades, arbitration orders, and circuit-breaker thresholds are configuration rather than code, an adopting organization encodes its own risk appetite—what counts as high risk, which constraints rank above throughput, when availability yields to integrity—without modifying the engine. This separation matters operationally: risk appetite differs across deployments of the same architecture, and hard-wiring any single appetite into engine code would either over-constrain permissive domains or under-protect regulated ones [16,18].

Standardization of the interface is likewise what makes the ecosystem claim concrete. If deviation inventories, arbitration records, and robustness scores converge on shared schemas, then third parties can build cross-framework dashboards, benchmark suites can compare reliability across agent frameworks on equal terms, and incident reports become machine-ingestible by the community at large. The history of infrastructure software suggests that such shared observational schemas—not any single framework's superiority—are what eventually consolidate a fragmented ecosystem, because they convert every participant's operational experience into a public good that all participants, including the framework vendors, consume [1].

4.5. Comparison with Mainstream Paradigms

Table 2 positions RAE against the four incumbent families along five dimensions: the objective of adversarial interaction, the correction strategy, the control structure, the convergence story, and the characteristic failure each leaves unaddressed. The comparison substantiates the paradigm claim made in Section 1: existing frameworks each solve a genuine sub-problem—training signal, factuality improvement, behavioral constraint, or state agreement—and RAE's contribution is the integration of all four concerns under a single runtime contract oriented to collective stability.

Three contrasts deserve emphasis. Against self-play, RAE rejects the zero-sum premise: adversarial energy exists to serve collective convergence, not to crown a victor, which is why examiner incentives exclude agreement outcomes while the collective objective remains cooperative [8,9]. Against debate collectives, RAE replaces round-counting with measured termination and static peer standing with evidence-driven weight adaptation, so that a habitually fabricating participant cannot steer the collective indefinitely [2,13].

A further engineering pain point in deployed collectives deserves mention: protocol fragility under participant turnover. Debate-style systems typically assume a fixed roster and a well-formed transcript; in production, agent versions are redeployed mid-conversation, contexts are truncated by budget, and participants join or leave. Framework surveys record malformed interactions and unreliable tool behavior among the dominant failure categories observed in practice [1,3]. RAE's verification layer is comparatively robust to churn because its unit of trust is the graded finding rather than the participant: a newly deployed agent inherits verification treatment from its role and history class, not from conversational seniority, and a redeployed agent's unverified proposals simply re-enter the inventory like any others. Against alignment and consensus, RAE operates inside the runtime window where emergent interaction failures

live, detecting what training-time pipelines cannot see and reconciling what consensus cannot question [6,7].

Table 2. Qualitative comparison of RAE with mainstream multi-agent and alignment paradigms.

Paradigm	Representative methods	Adversarial objective	Correction strategy	Control structure	Residual gap
Self-play / adversarial training	GANs and variants [8,10,11,12]; self-play curricula	Zero-sum; a winner must emerge	Training-time optimization	Implicit in game dynamics	Deployment-time stability unaddressed; oscillation persists [9]
Multi-agent debate	Debate collectives [2,13]	Improve answer quality via critique	Fixed-round judging / voting	Static peer roles	No per-agent accountability; residual disagreement unclassified
Alignment pipelines	RLHF [6]; Constitutional AI [7]; red-teaming [14,15]	Shape model behavior to preferences	Pre-deployment training	External to runtime	Runtime interaction failures out of scope
Distributed consensus	Paxos-family protocols	None (assumes correct components)	State reconciliation	Quorum rules	Replicates corrupted content; no error detection
RAE (this work)	Four-layer architecture (Section 3)	Bounded, non-zero-sum verification signal	Recursive correction with convergence monitoring	Dynamic checks and balances	Empirical validation pending (Section 4.6)

4.6. Limitations and Open Challenges

Candor about scope is a design obligation for a framework paper. RAE as presented here is a conceptual architecture with a formal skeleton; no claims in this paper rest on experimental measurement. Six limitations bound its current status. First, empirical validation is pending: the architecture's central predictions—that measured recursive correction reduces deviation monotonically in realistic workloads, and that weight adaptation degrades influence of unreliable agents fast enough to matter—require controlled studies on multi-agent benchmarks and the failure taxonomies now available [3]. Second, the deviation functional D aggregates heterogeneous dimensions (factual, consistency, compliance, conflict) whose weights are inherently contestable; choosing them is a governance decision as much as a technical one. Third, computational overhead is real: adversarial verification and multi-round recursion multiply inference cost, and the cost–stability frontier must be mapped per application class before deployment commitments. Fourth, verifier reliability is the epistemic foundation (assumption A1, Section 3.7), and correlated verifier–agent failures—a shared base model

replicating a blind spot—could defeat the architecture's core redundancy; verifier diversity is therefore not an optimization but a safety requirement. Fifth, adversarial dynamics introduce their own risks: examiner collusion, reward gaming of deviation grades, and strategic behavior around circuit-breaker thresholds are all plausible and require mechanism-design countermeasures of the kind the red-teaming literature studies [14,15]. Sixth, the Lyapunov-style analysis is an engineering analogy with explicit assumptions, not a theorem; a formal treatment—deviation functionals with approximation guarantees, contraction conditions for LLM-mediated repair operators—remains open, as does the interface to training-time alignment so that runtime and pre-deployment mechanisms reinforce rather than duplicate each other [17,18]. These challenges define the experimental agenda that must precede any safety-critical adoption.

must precede any safety-critical adoption.

A validation agenda can nonetheless be stated precisely, and its feasibility is itself an argument for the architecture. On benchmarks with constructable ground truth, the deviation functional's measurement validity can be established against annotated failures of the kind the MAST taxonomy provides [3]; recursion efficacy is testable as inventory shrinkage per cycle under ablations that disable individual layers; weight-adaptation dynamics are observable as influence curves for seeded unreliable agents; and overhead is measurable as the cost–stability frontier across recursion budgets. Each experiment maps onto an architectural mechanism rather than onto an emergent property, which is what distinguishes evaluating an architecture from evaluating a prompt arrangement, and each produces telemetry that downstream deployments would monitor in production anyway [17].

5. Conclusion

This paper proposed the Recursive Adversarial Engine, a stability-oriented runtime architecture for large language model multi-agent systems, motivated by four production-scale pathologies that incumbent orchestration, debate, alignment, and consensus frameworks each leave partially unaddressed: decision oscillation, hallucination propagation, long-horizon drift, and value-boundary erosion. RAE's design rests on three propositions—bounded adversarial interaction as the system's error signal, recursive correction as its error-elimination mechanism, and dynamic checks and balances as its stability guarantee—realized in a four-layer architecture specified at module and interface level: adversarial verification, recursive correction, checks-and-balances scheduling, and steady-state output.

The paper's analysis yields three takeaways. First, conflict and stability are not opposites in multi-agent design: the architecture converts adversarial pressure into a controlled verification signal and therefore into convergent correction, inverting the zero-sum premise of self-play while retaining its generative energy [8,9]. Second, stability must be engineered as an emergent property of mechanism design—measured termination, evidence-driven weights, graded circuit breaking—rather than asserted by static rules, which connects runtime multi-agent engineering to the classical vocabulary of stability theory [19-21]. Third, auditability is an architectural by-product that RAE obtains nearly for free: because every deviation, repair, arbitration, and intervention is recorded, process evidence replaces model introspection as the substrate of oversight [17,18].

The framework's limitations are stated without hedging in Section 4.6, and they define the road ahead: empirical evaluation against multi-agent failure taxonomies, formalization of deviation functionals and contraction conditions, cost–stability characterization across application classes, and verifier-diversity requirements. The destination justifies the effort. As multi-agent systems assume responsibility for decisions in finance, medicine, infrastructure, and government, the engineering question that matters is no longer whether agents can collaborate, but whether their collaboration can be made to converge, to remain within its value boundaries, and to account for itself. This paper contributes an architecture aimed at exactly that question.

References

[1] Guo, T., Chen, X., Wang, Y., Chang, R., Pei, S., Chawla, N. V., Wiest, O., & Zhang, X. (2024). Large language model based multi-agents: A survey of progress and challenges. arXiv preprint arXiv:2402.01680.

[2] Du, Y., Li, S., Torralba, A., Tenenbaum, J. B., & Mordatch, I. (2024). Improving factuality and reasoning in language models through multiagent debate. In Proceedings of the 41st International Conference on Machine Learning (ICML).

[3] Cemri, M., Pan, M. Z., Yang, S., Agrawal, L. A., Chopra, B., Tiwari, R., Keutzer, K., Parameswaran, A., & Song, D. (2025). Why do multi-agent LLM systems fail? arXiv preprint arXiv:2503.13657.

[4] Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y. J., Madotto, A., & Fung, P. (2023). Survey of hallucination in natural language generation. ACM Computing Surveys, 55(12), 1-35.

[5] Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H., Chen, Q., & Liu, T. (2023). A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. arXiv preprint arXiv:2311.05232.

[6] Christiano, P. F., Leike, J., Brown, T., Martic, M., Legg, S., & Amodei, D. (2017). Deep reinforcement learning from human preferences. In Advances in Neural Information Processing Systems (Vol. 30).

[7] Bai, Y., Kadavath, S., Kundu, S., Askell, A., Kernion, J., Jones, A., et al. (2022).

Constitutional AI: Harmlessness from AI feedback. arXiv preprint arXiv:2212.08073.
[8] Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. In Advances in Neural Information Processing Systems (Vol. 27).
[9] Mescheder, L., Geiger, A., & Nowozin, S. (2018). Which training methods for GANs do actually converge? In Proceedings of the 35th International Conference on Machine Learning (ICML).
[10] Arjovsky, M., Chintala, S., & Bottou, L. (2017). Wasserstein generative adversarial networks. In Proceedings of the 34th International Conference on Machine Learning (ICML).
[11] Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., & Courville, A. (2017). Improved training of Wasserstein GANs. In Advances in Neural Information Processing Systems (Vol. 30).
[12] Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., & Hochreiter, S. (2017). GANs trained by a two time-scale update rule converge to a local Nash equilibrium. In Advances in Neural Information Processing Systems (Vol. 30).
[13] Irving, G., Christiano, P., & Amodei, D. (2018). AI safety via debate. arXiv preprint arXiv:1805.00899.
[14] Ganguli, D., Hernandez, D., Lovitt, L., Askell, A., Bai, Y., Chen, A., et al. (2022). Red teaming language models to reduce harms: Methods, scaling behaviors, and lessons learned. arXiv preprint arXiv:2209.07858.
[15] Perez, E., Huang, S., Song, F., Cai, T., Ring, R., Aslanides, J., Gladstone, A., Irving, G., & Christiano, P. (2022). Red teaming language models with language models. In Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP).
[16] Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., & Mané, D. (2016). Concrete problems in AI safety. arXiv preprint arXiv:1606.06565.
[17] Hendrycks, D., Carlini, N., Schulman, J., & Steinhardt, J. (2021). Unsolved problems in machine learning safety. arXiv preprint arXiv:2109.13916.
[18] Russell, S. (2019). Human compatible: Artificial intelligence and the problem of control. Viking.
[19] Lyapunov, A. M. (1992). The general problem of the stability of motion. International Journal of Control, 55(3), 531-534.
[20] Slotine, J. J. E., & Li, W. (1991). Applied nonlinear control. Prentice Hall.
[21] Khalil, H. K. (2002). Nonlinear systems (3rd ed.). Prentice Hall.
[22] Varela, F. J. (1979). Principles of biological autonomy. North-Holland.
[23] Bak, P., Tang, C., & Wiesenfeld, K. (1987). Self-organized criticality: An explanation of the 1/f noise. Physical Review A, 38(1), 364-374.
[24] Turing, A. M. (1936). On computable numbers, with an application to the Entscheidungsproblem. Proceedings of the London Mathematical Society, s2-42(1), 230-265.
[25] Banach, S. (1922). Sur les opérations dans les ensembles abstraits et leur application aux équations intégrales. Fundamenta Mathematicae, 3, 133-181.
[26] Hofstadter, D. R. (1979). Gödel, Escher, Bach: An eternal golden braid. Basic Books.

Impressum

Founders	Zhengjie Gao, Xinyu Song
Editor in Chief	Ao Feng, Chengdu University of Information Technology, China
Executive Editor	Zhengjie Gao, Geely University of China, China
Editorial Board	Jing Hu, Huazhong University of Science and Technology, China Xiaohu Du, Huazhong University of Science and Technology, China Xiangkui Li, Harbin University of Science and Technology, China Zuopeng Liu, Goettingen University, Germany Xinyu Song, Geely University of China, China
Young Editorial Board	Min Liao, Geely University of China, China Tao Zheng, Geely University of China, China Chong Li, Chongqing University, China Ruiqin Fan, Sehan University, Korea Ziyang Liu, Jiangsu Normal University, China Qiwei Liu, Urumqi Vocational University, China Minqiu Kuang, Hunan Agricultural University, China

Published By	**Hong Kong Dawn Clarity Press Limited** Rm 9042, 9/F, Block B Chung Mei Centre, 15-17 Hing Yip Street, Kwun Tong, Kowloon, Hong Kong e-mail: ijaaa@dawnclarity.press ***International Journal of Advanced AI Applications*** is published monthly.

Editorial Policy	***International Journal of Advanced AI Applications*** is directed to the international communities of scientific researchers in artificial intelligence, computers and electronic, from the universities, research units and industry.To differentiate from other similar journals, the editorial policy of IJAAA encourages the submission of original scientific papers that focus on the integration of the advanced AI applications. In particular, the following topics are expected to be addressed by authors: (1) Natural Language Processing (NLP): Conversational AI, machine translation, sentiment analysis, and context-aware dialogue systems. (2) Smart Cities and IoT Integration: AI for traffic optimization, energy management, waste reduction, and urban infrastructure. (3) Autonomous Systems and Robotics: Self-driving vehicles, drones, industrial automation, and human-robot collaboration. (4) Edge AI and Distributed Systems: Real-time processing, federated learning, and low-latency AI at the network edge. (5) Creative and Generative AI: Art, music, and content generation using generative adversarial networks (GANs) and transformers. (6) AI in Education and Industry: Adaptive learning platforms, intelligent tutoring systems, and AI-driven supply chain optimization. Ethical and Explainable AI (XAI): Fairness, transparency, and accountability in real-world AI deployment.