Designing Emotion-adaptive Human - AI Interfaces: An Empirical Study on Empathy, Trust, and Context-aware Interaction

Quan Su*
School of Materials Science and Engineering, Nanjing Institute of Technology, Nanjing 211167, China
*Corresponding email: 1726377053@qq.com

Abstract

As artificial intelligence systems become increasingly embedded in everyday human-computer interaction contexts, user expectations regarding their social and emotional capabilities are gradually shifting from a primary focus on functional efficiency toward more complex dimensions such as empathetic experience, trust formation, and contextual sensitivity. Nevertheless, despite the notable progress achieved by large language models in terms of linguistic fluency, such systems often remain confined to surface-level simulations of empathy. In response to this limitation, the present study investigates an emotion-adaptive human-AI interface that integrates real-time affect recognition, dynamic user profiling, and context-aware response modulation within a unified framework. Through comparative user study encompassing task-oriented, social, and emotionally supportive scenarios, the results suggest that emotion-adaptive mechanisms may, to some extent, enhance users' perceived empathy, trust, and engagement.

Keywords

Emotion-adaptive interaction, Human-artificial intelligence interfaces, Empathy in human-computer interaction, Context-aware systems

Introduction

As artificial intelligence systems increasingly permeate everyday human-computer interaction, their role has extended beyond instrumental problem-solving toward more socially embedded forms of engagement. Users now tend to evaluate intelligent systems not only by efficiency or accuracy but also by whether interactions feel attentive, responsive, and situationally appropriate [1]. This shift, however, does not imply a clear consensus on what constitutes social intelligence in computational systems. Empathy, trust, and contextual sensitivity are often invoked as desired qualities, yet their operational meanings remain contested, particularly when translated into design requirements for interactive systems [2].

Recent advances in large language models have enhanced conversational fluency and semantic coherence, leading to more naturalistic dialogue behaviors across a wide range of applications. Nevertheless, such improvements raise further questions. While responses may appear empathetic at the linguistic level, it remains uncertain whether these systems genuinely adapt their interaction strategies based on users' emotional and contextual states or primarily reproduce surface-level patterns learned from large-scale data [3]. This survey is notable for synthesizing reinforcement learning-based approaches and evaluation challenges, providing a structured foundation for designing controllable and goal-aligned adaptive behaviors.

In our study, this motivates framing emotion-adaptive response modulation as a tunable strategy layer, rather than if empathetic phrasing alone guarantees adaptive interaction quality. While responses may appear empathetic at the linguistic level, it remains uncertain whether these systems genuinely adapt their interaction strategies based on users' emotional and contextual states, or whether they primarily reproduce surface-level patterns learned from large-scale data. This ambiguity becomes particularly salient in interaction scenarios that require sustained emotional engagement, where perceived empathy may depend less on wording itself and more on how responses evolve over time.

Within human-computer interaction, prior studies on affective computing and social interfaces have proposed various mechanisms for emotion recognition and

adaptive response generation [4]. Many of these approaches demonstrate technical feasibility under controlled conditions, yet their translation into everyday interactive settings often reveals limitations. Emotion classifiers may struggle with ambiguous or mixed affective signals, while rule-based adaptation strategies risk oversimplifying the dynamics of human emotion. These tensions suggest that emotional adaptation should perhaps be understood not as a deterministic pipeline, but as a probabilistic and context-dependent process whose effectiveness may vary across users and situations [5]. Viewing interaction signals as stochastic and time-varying is consistent with diffusion-style perspectives that model uncertainty and evolving dynamics rather than relying on fixed, deterministic mappings. Diffusion-based generative modeling is notable for formalizing complex variability through principled stochastic processes, offering a useful conceptual analogy for handling emotional volatility and uncertainty in adaptive interaction. In our interface, affect is therefore treated as a probabilistic influence that is weighted by context and history instead of being translated into rigid "if - then" rules.

Motivated by these observations, the present study explores the design of an emotion-adaptive human-AI interface that integrates real-time affect recognition, dynamic user profiling, and context-aware response modulation within a unified framework. Our design process involved multiple iterations in which early prototypes revealed unexpected sensitivities, including overreaction to transient emotional cues and inconsistencies across interaction contexts. These challenges prompted gradual adjustments in how emotional information was weighted, interpreted, and translated into adaptive strategies [6]. This work is notable for refining damping-oriented modeling to improve stability under nontrivial dynamics, offering a transferable intuition for preventing oscillatory adaptation loops in interactive systems. Guided by this principle, we incorporate smoothing and thresholding so that brief emotional spikes bias responses gradually rather than triggering abrupt strategy reversals. These challenges prompted gradual adjustments in how emotional information was weighted, interpreted, and translated into adaptive strategies.

To examine how such design choices influence user experience, comparative user study was conducted across task-oriented, social, and emotionally supportive interaction scenarios. Quantitative measures and qualitative feedback jointly suggest that emotion-adaptive mechanisms may, to some extent, enhance perceived empathy, trust, and engagement when compared with a non-adaptive baseline [7]. At the same time, alternative explanations related to novelty effects, individual differences in emotional expressiveness, and contextual framing cannot be entirely ruled out [8]. These findings highlight the importance of interpreting adaptive interaction outcomes from multiple analytical perspectives rather than attributing observed effects to a single design factor.

Considering these factors, this work positions emotion-adaptive interfaces not as definitive solutions to socially intelligent interaction, but as evolving socio-technical systems that require continued empirical examination and theoretical refinement. By foregrounding both the potential and the limitations of emotional adaptation, the study aims to contribute to ongoing discussions on how empathy, trust, and context awareness can be more carefully operationalized in human-AI interaction. Further research is needed to explore how such systems perform over longer periods of use, across cultural contexts, and under conditions where emotional expectations are less explicit yet equally consequential [9].

Related work

Emotion recognition and affective computing in HCI

Research on affective computing has long sought to enable interactive systems to recognize and respond to human emotions, drawing on developments in psychology, signal processing, and machine learning [10]. Early work often relied on explicit emotion labeling schemes or rule-based mappings between detected affective cues and predefined system behaviors. More recent studies have increasingly adopted data-driven approaches, particularly deep learning-based sentiment and emotion classifiers, which report notable performance gains in benchmark evaluations [11]. Despite these advances, it remains uncertain to what extent such models can be reliably deployed in everyday human-computer interaction, where emotional expressions are frequently subtle, context-dependent, and shaped by social or cultural conventions.

A persistent challenge concerns not merely the accuracy of emotion recognition, but how recognized emotional states should meaningfully inform interaction behavior. Several systems have assumed relatively stable mappings between emotion categories and response strategies, implicitly treating affective states as actionable signals [12]. Empirical observations, however, suggest that users' emotional expressions are often ambiguous, strategically modulated, or even contradictory across interaction turns. In such cases, rigid adaptation rules may amplify misinterpretations rather than mitigate them. These findings indicate that emotion recognition, while necessary, may be insufficient on its own to support socially adaptive interaction, and that its role must be reconsidered in relation to temporal continuity and situational context [13].

Adaptive dialogue systems and context-aware interaction

Alongside affective computing, research on adaptive dialogue systems has focused on enabling conversational agents to tailor their responses based on user intent, interaction history, and contextual information. Context-aware systems typically model variables such as task state, temporal cues, or user preferences, allowing interactions to unfold beyond isolated conversational turns. While our current prototype relies primarily on interaction history and task framing, the architecture is compatible with sensor-derived context cues demonstrated in IoT (Internet of Things) +LLM (Large Language Model) activity tracking systems [14]. From a methodological standpoint, memory-augmented architectures and attention-based mechanisms have demonstrated potential for maintaining coherence across extended interactions. Nevertheless, these systems tend to prioritize informational relevance or task completion, while emotional adaptation is often treated as a secondary or peripheral concern [15].

Efforts to integrate emotional sensitivity into adaptive dialogue systems have produced mixed outcomes. Some studies report improvements in perceived naturalness or engagement, whereas others highlight inconsistencies in user perception when adaptive behaviors appear excessive, unpredictable, or misaligned with situational expectations. These divergent findings suggest that adaptation itself is not inherently beneficial, and that its effectiveness may depend on how proportionately and transparently it is applied [16]. As a result, existing research leaves open important questions regarding how emotional signals should be balanced with contextual cues, and whether adaptive systems should aim for maximal responsiveness, deliberate restraint, or a negotiated balance between the two.

Gaps and positioning of the present study

Considering these strands of research together, a gap becomes apparent between technical capability and experiential reliability in emotion-adaptive human-AI interaction. While emotion recognition and context-aware dialogue systems have each progressed substantially, their integration often presupposes idealized conditions in which emotional signals are stable, interpretable, and directly actionable. In practical interaction settings, however, emotional cues are frequently transient, ambiguous, or intertwined with situational norms, making straightforward adaptation strategies difficult to sustain over time.

The present study positions itself within this gap by approaching emotional adaptation as an exploratory design problem rather than a fully resolved technical solution [17]. Instead of optimizing for maximal emotional responsiveness, we examine how different degrees of emotional sensitivity, contextual weighting, and user modeling shape users' perceptions of empathy and trust. This perspective does not seek to eliminate ambiguity in emotion-adaptive interaction, but rather to render it analytically visible [18]. By foregrounding both successful adaptations and observed breakdowns during system design and evaluation, the study aims to contribute a more cautious and empirically grounded understanding of how emotion-adaptive interfaces may function in real-world human-AI interaction, while acknowledging that further research is needed to refine these approaches across broader contexts.

Methodology

System design

The design of the proposed emotion-adaptive human-AI interface was guided by the need to balance responsiveness with restraint. Rather than treating emotional adaptation as a deterministic process, the system was conceived as a layered architecture in which affective signals, contextual cues, and user-related information jointly informed response generation. At the core of the system lies a dialogue engine responsible for producing linguistically coherent responses, while

surrounding modules modulate these outputs based on emotional and contextual inputs.

Specifically, the system integrates three functional components: real-time affect recognition, dynamic user profiling, and context-aware response modulation. Emotional cues extracted from user input are represented along continuous dimensions rather than discrete categories, reflecting the observation that emotional expressions in interaction are often ambiguous or mixed. User profiles are incrementally updated based on interaction history, capturing tendencies such as preferred interaction style and typical emotional expressiveness. Contextual factors, including task type and interaction stage, are used to adjust the relative influence of emotional signals during response generation. The profiling module is designed to accommodate non-stationarity by periodically recalibrating key parameters, consistent with drift-aware perspectives in continuous deployment. Drift-aware deployment also benefits from anomaly-sensitive monitoring, because sudden shifts in behavior may reflect breakdowns, miscalibration, or context change rather than true preference updates. Physics-regularized, self-supervised anomaly detection is notable for providing an engineering-ready approach to identifying non-obvious deviations under noisy real-world signals - an idea that maps well to detect interaction misalignment in adaptive dialogue. Following this logic, we treat unexpected spikes in negative feedback or rapid tone mismatch as "interaction anomalies" that trigger recalibration rather than stronger emotional escalation.

During early system iterations, direct coupling between detected affect and response tone led to unstable interaction patterns, particularly when users expressed fleeting or contradictory emotions. As a result, emotional inputs were reweighed to function as probabilistic biases rather than explicit triggers. This design choice aligns with the broader idea that selective masking or gating can stabilize sequence-level behavior by preventing transient signals from dominating model dynamics. Dynamic focused masking is notable for showing how structured suppression of noisy cues can improve robustness in autoregressive settings - an insight that transfers naturally to affect-driven modulation in dialogue. Accordingly, we treat affect signals as soft constraints that bias tone and framing, while preserving continuity and preventing abrupt emotional "jumps". This design choice reflects an underlying assumption that emotional adaptation in interactive systems should remain adjustable and revisable, rather than fixed or exhaustive.

Experimental setup

To evaluate the effects of emotion-adaptive interaction, a within-subjects experimental design was adopted [19]. This approach allowed each participant to experience both the emotion-adaptive system and a non-adaptive baseline, thereby reducing inter-individual variability in emotional expressiveness and interaction style. Sixty participants were recruited, with attention given to diversity in age, gender, and professional background.

Participants engaged in three categories of interaction scenarios: task-oriented assistance, casual social conversation, and emotionally supportive dialogue [20]. The task-oriented scenarios were designed to reflect multi-step, decision-heavy assistance settings that have been highlighted as challenging for LLMs in realistic benchmark suites. The order of system exposure and scenario presentation was counterbalanced to mitigate learning and fatigue effects. Instruction-following evaluations that explicitly test hierarchical and multi-turn constraints indicate that seemingly fluent dialogue can still fail under layered, real-world goal structures. IHEval is notable for operationalizing "instruction hierarchy" as a measurable capability, offering a rigorous lens for stress-testing LLM behavior beyond single-turn correctness. Consistent with this perspective, our task-oriented scenarios emphasize multi-step continuity to examine whether emotion-adaptive modulation remains stable under realistic interaction complexity. Participants were instructed to interact naturally and were not explicitly informed about the adaptive mechanisms embedded in the experimental system, to minimize expectancy bias.

During early sessions, some participants expressed uncertainty regarding the appropriate level of emotional engagement, particularly in supportive scenarios. Rather than standardizing these interactions through additional scripting, the study retained this ambiguity, as it reflects real-world conditions in which emotional expectations are often implicit. Interaction logs were recorded automatically, and post-task questionnaires were administered following each scenario [21].

Evaluation metrics

Evaluation of user experience combined quantitative performance indicators with subjective perceptual

measures. Quantitative metrics included Task Completion Rate and Average Engagement Time, capturing functional effectiveness and interaction depth. Subjective measures assessed perceived empathy, trust, and engagement using five-point Likert scales selected for their interpretability and common use in HCI (Human-Computer Interaction) research. Including completion-oriented metrics follows the motivation of emerging real-world evaluations that emphasize multi-step goal satisfaction rather than single-turn.

Data was analyzed using paired comparisons between the adaptive and non-adaptive conditions. Comparative evaluation is most informative when it triangulates multiple metrics and examines error profiles rather than relying on a single headline score. Comparative model studies are notable for emphasizing robustness checks and metric consistency across settings, offering a transferable evaluation discipline for system-to-system comparisons in HCI. Accordingly, we interpret paired differences alongside variance across participants and scenarios, treating heterogeneity as evidence about boundary conditions rather than as noise to ignore. In addition to examining mean differences, the analysis considered variance across participants and scenarios, acknowledging that emotional adaptation may not exert uniform effects. Qualitative comments were analyzed thematically to contextualize quantitative findings and to surface potential mismatches between system behavior and user expectations.

Results

Quantitative results

The quantitative results indicate consistent differences between the emotion-adaptive system and the non-adaptive baseline across multiple evaluation dimensions. Participants interacting with the emotion-adaptive system exhibited higher task completion rates and longer engagement times across all interaction scenarios. These differences were most pronounced in emotionally supportive contexts, where engagement time increased notably, suggesting that adaptive mechanisms may encourage prolonged interaction under certain conditions (as shown in Table 1).

At the same time, variability across participants and scenarios remained evident. While task-oriented interactions showed clear efficiency gains, improvements in engagement time and perceived empathy were less uniform, with some participants demonstrating only marginal changes. This heterogeneity suggests that observed quantitative improvements cannot be attributed solely to emotional adaptation, but may also reflect individual interaction styles, scenario framing, or participants' prior expectations of intelligent systems.

Table 1. Summary of quantitative measures.

Metric	Non-adaptive system	Emotion-adaptive system
Task completion rate (%)	78.4	94.6
Avg. engagement time (min)	3.2	5.7
Perceived empathy (1-5)	3.1	4.4
User trust (1-5)	3.2	4.3
Engagement (1-5)	3.4	4.5

Qualitative results

Qualitative feedback offered further insight into how participants interpreted and experienced the system's adaptive behaviors. Many participants described the emotion-adaptive interface as more attentive and socially present, frequently noting its ability to reference earlier interaction content or adjust its tone in ways that felt situationally appropriate. Such responses were especially common in emotionally supportive scenarios, where participants reported feeling acknowledged rather than merely responded to.

However, qualitative accounts also revealed moments of discomfort. Some participants perceived emotional responses as overly confident or insufficiently aligned

with their actual intentions, interpreting adaptive behavior as intrusive rather than empathetic. These reactions suggest that perceived empathy may depend less on emotional sensitivity alone and more on whether adaptive responses remain proportionate and negotiable within the interaction. In deployed settings, users' affective expressiveness and interaction preferences can drift over time, which motivates adaptive strategies that are continuously updated rather than treated as static mappings.

Taken together, the qualitative findings complicate a straightforward interpretation of adaptive success, indicating that emotional adaptation can both enhance and undermine user experience depending on alignment with users' implicit norms of appropriate emotional engagement.

Statistical significance

Paired t-tests were conducted to examine whether the observed quantitative differences reflected systematic effects rather than random variation. The results indicated statistically significant differences between the emotion-adaptive and non-adaptive systems for perceived empathy, trust, and engagement, with p-values below 0.01. Task completion rate and average engagement time also demonstrated significant improvement under the adaptive condition.

Despite these statistically significant findings, their interpretation warrants caution. Statistical significance does not preclude alternative explanations such as novelty effects, heightened participant attention due to system complexity, or scenario-specific framing effects. Moreover, statistically significant mean differences do not imply uniform benefit across individuals or interaction contexts. These considerations suggest that while the results support the potential value of emotion-adaptive interaction, further research is needed to examine the robustness and generalizability of these effects over longer-term and cross-context use.

Discussion

Empathy and social presence

The findings of this study suggest that emotion-adaptive interaction may, to some extent, enhance users' perceptions of empathy and social presence, yet these effects appear neither uniform nor unconditional. Rather than emerging as a direct consequence of emotional sensitivity alone, perceived empathy seemed to arise from a combination of factors, including continuity across interaction turns, moderation in response tone, and alignment with situational expectations. This observation complicates simplified assumptions that increasing emotional awareness will necessarily lead to more empathetic interaction.

From one perspective, the observed improvements in perceived empathy can be interpreted as evidence that adaptive systems can simulate relational cues that users associate with attentive and socially responsive behavior. At the same time, qualitative feedback revealed instances in which similar adaptive behaviors were interpreted negatively, particularly when emotional responses appeared overly confident or insufficiently grounded in the user's expressed intent. These contrasting interpretations suggest that social presence in human-AI interaction may be less dependent on the accuracy of emotion recognition and more contingent on the system's capacity to remain tentative, adjustable, and open to correction.

Considering these factors, empathy in emotion-adaptive interfaces may be better understood as a negotiated experience rather than an inherent system property. The system's role may not be to demonstrate empathy in a definitive sense, but to create conditions under which users can plausibly attribute empathic intent. This perspective points toward a design orientation that prioritizes responsiveness without overcommitment, acknowledging that further research is needed to clarify how such balance can be maintained across diverse interaction contexts and user populations.

Privacy and profiling

The incorporation of dynamic user profiling in emotion-adaptive systems introduces both practical advantages and substantive concerns regarding privacy and user agency. On one hand, maintaining interaction history and inferred preferences appeared to support continuity and personalization, which some participants associated with a sense of being understood over time. On the other hand, several participants expressed uncertainty about how much the system remembered, and whether such memory functioned in their interest or beyond their awareness.

These responses indicate that privacy in emotion-adaptive interaction cannot be reduced to data protection mechanisms alone. Even in cases where explicit identifiers are absent and data are processed locally, the

perception of being profiled may shape how users engage with the system. In this sense, privacy operates not only as a technical constraint, but also as an experiential factor that influences trust and willingness to disclose emotional information.

During the design process, attempts to increase transparency around user modeling revealed an additional tension. While explanatory cues were intended to reassure users, they occasionally disrupted interaction flow or highlighted the constructed nature of the system's empathy. This suggests that transparency and comfort do not always align straightforwardly. Balancing informative disclosure with interactional smoothness remains an open design challenge, pointing to the need for more graduated and user-controllable approaches to profiling visibility.

Ethical considerations

Beyond usability and performance, the results raise broader ethical questions concerning emotionally responsive AI systems. When systems adapt their behavior based on inferred emotional states, they inevitably participate in shaping users' emotional experiences, whether intentionally or not. This raises concerns about emotional manipulation, even in contexts where system designers primarily aim to provide support or enhance engagement.

The findings suggest that ethical risk arises not solely from system capability, but from asymmetries in understanding and control. Users may attribute intentionality or care to adaptive responses without fully comprehending the underlying mechanisms, while designers retain disproportionate influence over how emotional cues are interpreted and acted upon. Addressing this imbalance requires more than retrospective disclosure; it calls for ethical considerations to be embedded directly within interaction design decisions.

Rather than proposing fixed ethical guidelines, this study points toward an ethics of interaction that emphasizes reflexivity, user agency, and revisability. Emotion-adaptive systems should allow space for misalignment, correction, and disengagement, recognizing that emotional appropriateness is contextually situated and subject to change. Further research is needed to explore how such principles can be operationalized in practice, particularly as emotion-aware interfaces become more prevalent in sensitive domains.

Conclusion

This study does not position emotion-adaptive human-AI interaction as a definitive pathway toward "artificial empathy" but rather conceptualizes it as an evolving design and research process. By integrating affect recognition, contextual modulation, and user profiling mechanisms, the findings suggest that such approaches may, to some extent, enhance users' perceived empathy, trust, and engagement, while simultaneously revealing instances of interactional misalignment and discomfort. These tensions indicate that linear narratives of progress are insufficient to account for the practical complexity observed in emotionally adaptive interaction systems.

Funding

This work was not supported by any funds.

Acknowledgements

The author would like to show sincere thanks to those techniques who have contributed to this research.

Conflicts of Interest

The author declares no conflict of interest.

References

[1] Hao, Y. C., Chen, L. W., Li, J., Guo, Y., Su, X., Shu, M., Yan, C. H. (2021) Metal-organic framework membranes with single-atomic centers for photocatalytic CO_2 and O_2 reduction. *Nature Communications*, 12(1), 2682.

[2] Yin, M. (2025) Data quality control in semiconductor manufacturing through automated ETL processes and class imbalance handling techniques. *Journal of Industrial Engineering and Applied Science*, 3(6), 13-22.

[3] Liu, J., Wang, H., Sun, W., Liu, Y. (2022) Prioritizing autism risk genes using personalized graphical models estimated from single-cell rna-seq data. *Journal of the American Statistical Association*, 117(537), 38-51.

[4] Lin, A. (2025) Toward regulatory compliance in DAO governance: from regulatory rule engines to on-chain audit report generation. *Journal of World Economy*, 4(6), 12-20.

[5] Chen, Y. (2025) Interpretable automated machine learning for asset pricing in US capital markets. *Journal of Economic Theory and Business Management*, 2(5), 15-21.

[6] Yin, M. (2025) A data-driven approach for real-time bottleneck detection and optimization in semiconductor manufacturing using active period method and visualization. *Academic Journal of Natural Science*, 2(4), 19-26.

[7] Yin, M. (2025) Defect prediction and optimization in semiconductor manufacturing using explainable AutoML. *Academic Journal of Natural Science*, 2(4), 1-10.

[8] Chen, Y. (2025) A comparative study of machine learning models for credit card fraud detection. *Academic Journal of Natural Science*, 2(4), 11-18.

[9] Zhu, H., Luo, Y., Liu, Q., Fan, H., Song, T., Yu, C. W., Du, B. (2019) Multistep flow prediction on car-sharing systems: a multi-graph convolutional neural network with attention mechanism. *International Journal of Software Engineering and Knowledge Engineering*, 29(11n12), 1727-1740.

[10] Liu, Z. (2025) Reinforcement learning for prompt optimization in language models: a comprehensive survey of methods, representations, and evaluation challenges. *ICCK Transactions on Emerging Topics in Artificial Intelligence*, 2(4), 173-181.

[11] Li, Y., Cheng, H., Qin, Q. (2025) Evaluations and improvement methods of deep learning ability in blended learning. *International Journal of e-Collaboration (IJeC)*, 21(1), 1-17.

[12] Wang, H., Li, Q., Liu, Y. (2022) Regularized Buckley-James method for right-censored outcomes with block-missing multimodal covariates. *Stat*, 11(1), e515.

[13] Wang, H., Li, Q., Liu, Y. (2023) Adaptive supervised learning on data streams in reproducing kernel Hilbert spaces with data sparsity constraint. *Stat*, 12(1), e514.

[14] Gupta, K., Zhang, Y., Gunasekaran, T. S., Krishna, N., Pai, Y. S., Billinghurst, M. (2024) Caevr: Biosignals-driven context-aware empathy in virtual reality. *IEEE Transactions on Visualization and Computer Graphics*, 30(5), 2671-2681.

[15] Chen, Y. (2025) Artificial intelligence in economic applications: stock trading, market analysis, and risk management. *Journal of Economic Theory and Business Management*, 2(5), 7-14.

[16] Malchanau, A., Petukhova, V., Bunt, H. (2018) Towards integration of cognitive models in dialogue management: Designing the virtual negotiation coach application. *Dialogue & Discourse*, 9(2), 35-79.

[17] Chen, Y. (2025) Generative diffusion models for option pricing: a novel framework for modeling volatility dynamics in us financial markets. *Journal of Industrial Engineering and Applied Science*, 3(6), 23-29.

[18] Lin, A. (2025) Low-barrier pathways for traditional financial institutions to access Web3: compliant wallet custody and asset valuation models. *Frontiers in Management Science*, 4(6), 80-86.

[19] Wang, P., Wang, H., Li, Q., Shen, D., Liu, Y. (2024) Joint and individual component regression. *Journal of Computational and Graphical Statistics*, 33(3), 763-773.

[20] Jin, Y., Li, Z., Zhang, C., Cao, T., Gao, Y., Jayarao, P., Yin, B. (2024) Shopping mmlu: a massive multi-task online shopping benchmark for large language models. *Advances in Neural Information Processing Systems*, 37, 18062-18089.

[21] Chen, Y. (2025) Leveraging LSTM networks for vehicle stability prediction: a comparative analysis with traditional models under dynamic load conditions. *Computing and Interdisciplinary Science*, 1(2), 15-22.