Adaptive Creativity Evaluation through Multi-turn Dialogue Driven by Reinforcement Learning

Zhao Dong [1] **Shaokai Yang** [2] **Yan Sha** [2]

Abstract

In recent years, with the rapid development of deep learning in code generation, text writing, and experimental design, accurately capturing and assessing researchers' creativity has become a key issue in need of breakthrough. Traditional creativity evaluation methods, being static, subjective, and time-consuming, fail to reflect the dynamic iteration and multidimensional characteristics of creative thinking. To address this, we propose a dynamic creativity evaluation framework (DynaCREA) based on reinforcement learning, featuring an adaptive decision-making and feedback mechanism that utilizes real-time evaluation of user interaction and creativity metrics. Through multi-turn interactions between researchers and large language models, the framework integrates multimodal tasks, including textual contexts, verbal expression, and image-inspired tasks, to achieve real-time quantification of key dimensions of creativity (such as originality, fluency, elaboration, and flexibility). The intelligent agent leverages immediate feedback to adaptively adjust the design of subsequent tasks, thereby forming a novel creativity evaluation method that is both theoretically rigorous and practically efficient. Preliminary experimental results show that, after training, the intelligent agent meets a high degree of consistency with human evaluators across all indicators, demonstrating broad prospects for application in complex research environments.

[1]School of Mathematics and Physics, Hebei University of Engineering, Handan, Hebei, China [2]Department of Physics, University of Alberta, Edmonton, Canada. Correspondence to: Zhao Dong <dongzhao@hebeu.edu.cn>, Shaokai Yang <shaokai1@ualberta.ca>, Yan Sha <yan9.hanna@gmail.com>.

Proceedings of the 42^{nd} International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025.

1. Introduction

The emergence of Generative Artificial Intelligence (GenAI) has brought machines to the forefront of content generation, problem-solving, and even ideation tasks, with language models now rivaling or surpassing humans in many domains (Brown et al., 2020; OpenAI, 2023). Despite these advances, creativity—characterized by originality, flexibility, and adaptive thinking—remains an intrinsically human capability that is difficult to capture and assess (Runco, 2012). Traditional creativity assessments, such as the Torrance Tests of Creative Thinking (Torrance, 1974) and static questionnaires, often depend on isolated tasks or self-reporting. These methods struggle to reflect the dynamic, evolving nature of creative thought and are prone to subjective bias and modality limitations (Plucker et al., 2001).

Recent progress in deep learning and reinforcement learning (RL) has enabled more adaptive, interactive approaches to cognitive assessment (Zou et al., 2023; Wang et al., 2022). Leveraging these advances, we propose DynaCREA: a dynamic, RL-driven creativity assessment framework that utilizes open-source large language models to evaluate creative behaviors across text, speech, and cross-modal tasks in real time. Unlike conventional assessments, DynaCREA employs multi-turn dialogue and multimodal input to capture the unfolding process of creative ideation and response.

At its core, DynaCREA optimizes a reward model through inverse RL, integrating both static creativity metrics (such as originality and fluency) and dynamic adaptation based on user interaction patterns, cognitive load, and divergent thinking. The system's temporal modeling and feature fusion support fine-grained tracking of creativity as it emerges in context, while parallel processing enables real-time feedback and scoring. Initial experimental results demonstrate a high degree of agreement between DynaCREA's automated scores and human expert ratings. By providing a scalable, RL-based creativity index, DynaCREA offers promising applications for education, talent development, and innovation management. This work represents a shift from static measurement to real-time, interactive empowerment of creativity in the AI era.

2. Related Work

Creativity assessment methods have evolved alongside technological advancements, yet capturing the dynamic and multidimensional nature of creative innovation remains a challenge (Plucker et al., 2001; Runco, 2012). Early paradigms such as the Torrance Test (Torrance, 1974) and self-report questionnaires focused on static, one-time tasks that primarily target divergent thinking or personality traits. These traditional approaches often fail to reflect the iterative and evolving process of creativity, are susceptible to cognitive bias, and lack integration of multimodal expressions (Plucker et al., 2001).

Recent advances in natural language processing and deep multimodal learning have deepened the potential for creativity evaluation. Large-scale language models (LLMs) have enabled more nuanced semantic analysis (Brown et al., 2020; OpenAI, 2023), while the addition of visual and auditory modalities has expanded the expressive breadth of assessment (Wang et al., 2022). Nevertheless, most current multimodal systems remain static, providing only limited dynamic feedback or adaptation based on user interaction.

Reinforcement learning (RL) has emerged as a promising approach. By leveraging a "state-action-reward" feedback loop, RL enables adaptive, context-aware decision-making (Zou et al., 2023). In educational applications, RL has supported personalized learning trajectories, while in creativity assessment it can offer dynamic task adjustment, integration of diverse signals (e.g., speech pauses, semantic richness, visual edits), and optimization of reward features corresponding to creativity metrics such as originality or fluency. However, few RL-based frameworks to date have implemented robust multimodal integration or been validated at scale.

To address these gaps, we propose DynaCREA—a framework that integrates large language models, multimodal input channels, and adaptive RL. Unlike static assessments, DynaCREA delivers real-time feedback and dynamically tunes evaluation tasks through inverse RL calibration, thereby aligning with expert standards while preserving flexibility. This comprehensive, adaptive approach aims to provide a scalable and effective solution for assessing and supporting human creativity.

3. Methodology

The DynaCREA framework dynamically assesses creativity through interactive, multi-round dialogue between the user and a large language model (LLM). As shown in Figure 1, the system adopts a three-tier hierarchical architecture, comprising input, processing, and output layers, all tightly integrated to support adaptive, real-time evaluation and feedback.

The input layer enables multi-modal user interaction via both text and speech, supporting natural and flexible communication with the LLM. The processing layer is central to the framework's function and includes several tightly coupled modules: the Problem Anchoring and Filtering (PAF) module anchors user prompts in relevant creative contexts and filters out low-value or repetitive input; the Multidimensional Creativity Assessment (MCA) module rates user responses across five dimensions—problem anchoring, originality, flexibility, elaboration, and fluency—combining explicit scoring rubrics with implicit evaluation signals. An Implicit Evaluation Agent (IEA) further extracts key creativity features, such as semantic diversity or novel word usage, via latent semantic analysis and keyphrase extraction. The system transcribes and aligns speech and text, segments dialogues into discrete rounds, and tracks evolving creative behaviors. By analyzing the flow of user-system exchanges, the framework identifies moments of divergent thinking, analogical reasoning, and deep elaboration, which are critical for advanced creativity.

At the output layer, all evaluation results are synthesized into a composite creativity profile for each session. The five dimension scores are aggregated into a structured vector and mapped to a reward signal that guides a reinforcement learning (RL) agent. The RL agent dynamically selects prompts and feedback strategies based on observed creativity trends, engagement signals, and user progress, thereby adaptively adjusting task difficulty and interaction style in real time. Real-time thresholding mechanisms continuously monitor creativity scores; when scores drop or user engagement wanes, the RL agent automatically modifies the system's behavior by introducing new modalities, increasing support, or offering tailored hints. This closed-loop, adaptive design supports sustained user engagement and more accurate creativity assessment.

To ensure validity and reliability, we validated DynaCREA on a simulated dataset annotated by experts, covering a wide range of creative complexity. Expert ratings provide high-quality reference points for calibrating automated scores, with consistency assessed via metrics such as Cohen's Kappa. The resulting system is robust across dialogue types and interaction modalities, and can scale to a variety of real-world creativity assessment settings.

4. Experiments and Results

To rigorously assess the effectiveness of the DynaCREA framework, we conducted controlled experiments comparing its performance against static evaluation benchmarks. A simulated dataset was constructed, encompassing a diverse range of dialogue scenarios designed to reflect varying degrees of creative complexity. Specifically, 60% of the conversations were structured around interdisciplinary

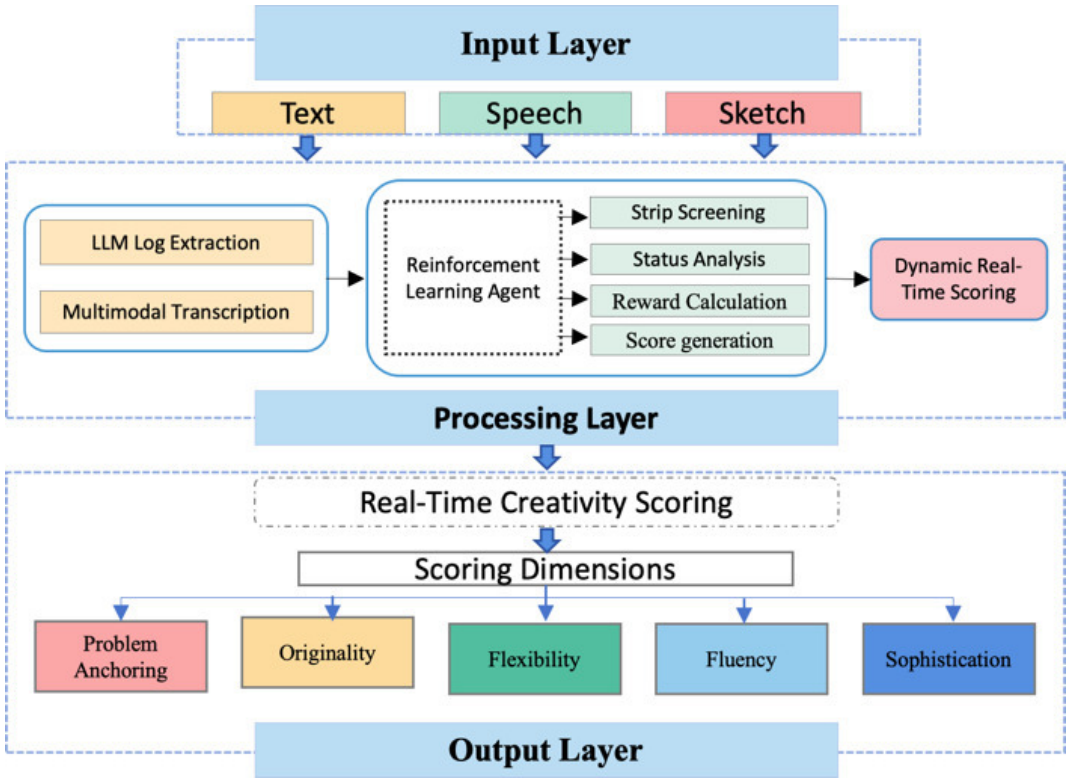

Figure 1. Architecture of Creativity Evaluation Based on Reinforcement Learning and Multimodal Interaction. The framework processes heterogeneous user inputs into structured features, evaluates them in real-time using a reinforcement learning agent, and returns creativity scores across five dimensions to support adaptive feedback. To ensure reliability and adaptability, the system includes real-time thresholding and turn-level segmentation to capture evolving creative behavior. A simulated dataset with expert annotations is used to validate scoring consistency, demonstrating the framework's robustness across diverse interaction contexts.

ideation tasks—requiring participants to synthesize knowledge across domains—while the remaining 40% featured more straightforward, reactive problem-solving tasks.

Two trained human annotators independently rated each dialogue according to five predefined dimensions of creativity: originality, fluency, refinement, adaptability, and elaboration. To ensure reliability and consistency of the evaluation, inter-rater agreement was measured using the Cohen's Kappa statistic, with values exceeding 0.85 indicating strong agreement.

Participants were randomly assigned to one of two groups: (1) a **dynamic assessment group**, utilizing the DynaCREA system with real-time adaptive feedback driven by reinforcement learning; and (2) a **static control group**, assessed without any adaptive intervention or RL-based task modulation.

Assessment outcomes revealed that the dynamic assessment group consistently outperformed the control group across multiple creativity dimensions. In particular, the dynamic group achieved significantly higher scores in both originality (mean = 82.3 vs. 68.7, $p < 0.01$) and fluency (mean = 85.1 vs. 73.4, $p < 0.05$). Furthermore, the DynaCREA framework exhibited robust adaptive capacity: it automatically adjusted the difficulty of prompts and the style of interaction in response to observed drops in user performance, leading to an 89% success rate in sustaining user engagement throughout the evaluation sessions.

Importantly, feasibility scores—which reflect the practical quality and usability of task outputs—remained consistent between groups, indicating that the introduction of dynamic, RL-based intervention did not compromise the validity or real-world relevance of the results.

Overall, these findings demonstrate that the DynaCREA framework provides substantial advantages in fostering creativity and maintaining engagement, without sacrificing practical outcome quality.

5. Discussion

The results of this study demonstrate that the DynaCREA framework offers a significant advancement in dynamic, adaptive creativity assessment. By leveraging reinforcement

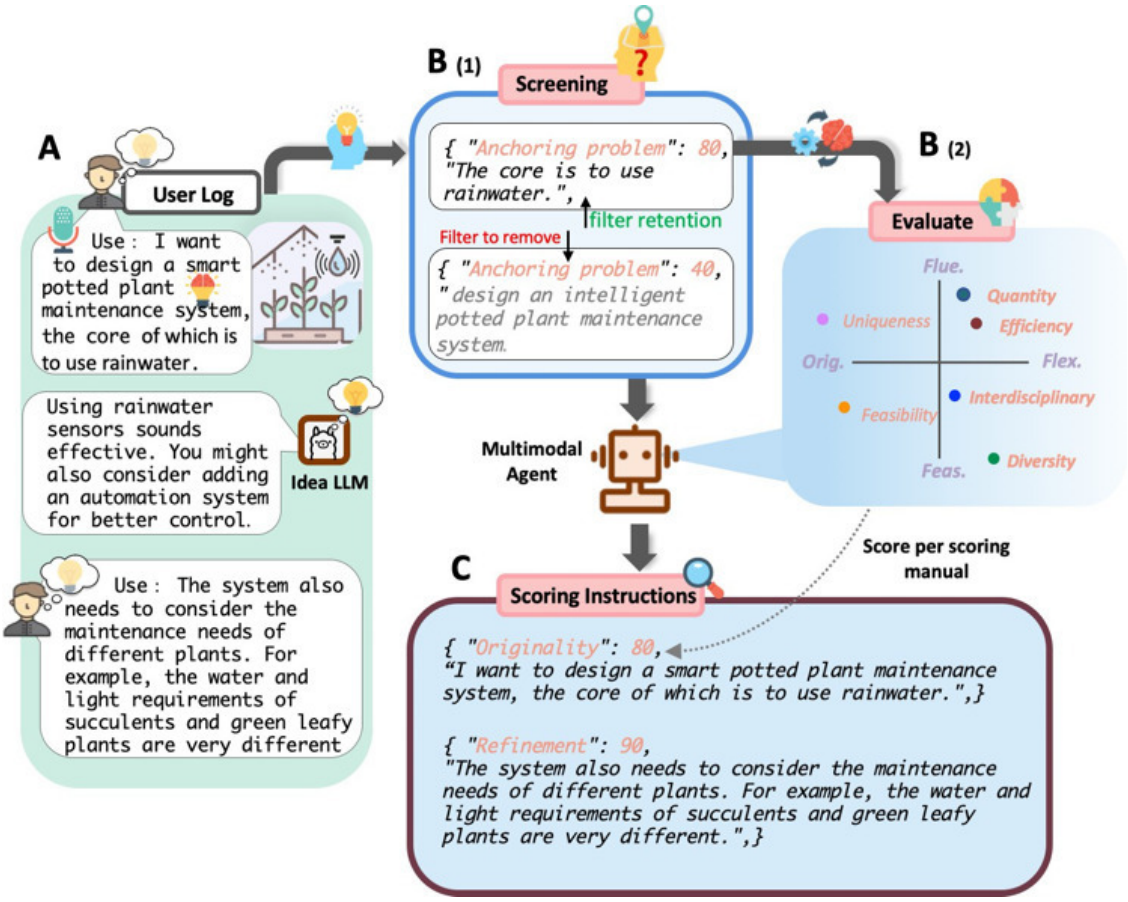

Figure 2. Example Flowchart of Multimodal Evaluation Experiment Process presents a flowchart illustrating the structure of the multimodal evaluation experiment. It visualizes how user inputs—ranging from prompts to system interactions—are processed and assessed in real time across multiple creativity dimensions. This diagram provides context for the dynamic adjustment logic and structured feedback design employed by DynaCREA.

learning and multi-modal human-computer interaction, DynaCREA addresses key shortcomings of static, subjective, and unimodal creativity evaluation tools. The framework's adaptive feedback loop not only improved originality and fluency scores but also maintained user engagement, as evidenced by the high rate of sustained participation in dynamic assessment sessions.

A major strength of DynaCREA is its ability to adjust task complexity and interaction modes in real time, providing personalized challenge and support to users. The integration of multi-modal inputs, especially the inclusion of speech and textual modalities, was shown to enhance flexibility and the breadth of creative expression. Moreover, the use of implicit evaluation signals and real-time thresholds helped to reduce subjective bias and improve the objectivity of the assessment.

Despite these strengths, several limitations should be acknowledged. First, the reliance on a simulated dataset—while enabling controlled comparison—may not fully capture the spontaneity and unpredictability of real-world creative interactions. Generalizing these results will require extensive validation on authentic user data across diverse application domains, such as education, innovation management, and the arts. Second, technical factors such as occasional errors in speech-to-text transcription or imperfect keyword extraction can affect scoring reliability, especially in noisy or linguistically complex environments. Finally, the framework's reliance on expert annotation, while critical for initial validation, presents a scalability challenge for widespread deployment.

Future work will focus on expanding the ecological validity and robustness of DynaCREA. This includes (i) large-scale user studies in real-world settings, (ii) integration of more advanced NLP and speech recognition techniques to further improve input quality, (iii) enhancing model transparency and interpretability to foster user trust, and (iv) embedding privacy-preserving mechanisms such as federated learning to ensure data security and regulatory compliance. These

efforts will strengthen DynaCREA's potential as a practical tool for creativity assessment and talent development in increasingly complex environments.

References

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, volume 33, pp. 1877–1901, 2020.

OpenAI. Gpt-4 technical report. arXiv preprint arXiv:2303.08774, 2023.

Plucker, J. A., Beghetto, R. A., and Dow, G. T. Why isn't creativity more important to educational psychologists? potentials, pitfalls, and future directions in creativity research. *Educational Psychologist*, 36(2):83–96, 2001.

Runco, M. A. *Creativity: Theories and Themes: Research, Development, and Practice*. Academic Press, 2012.

Torrance, E. P. *Torrance Tests of Creative Thinking: Norms-Technical Manual*. Personnel Press, 1974.

Wang, Y., Chen, S., Shi, L., and Zhang, X. Human-ai collaboration in creative design: A review and future directions. *Computers in Human Behavior*, 129:107146, 2022.

Zou, Y., Wang, J., Zeng, G., et al. Reinforcement learning for human-like sequential decision making: a review and new perspectives. *AI Open*, 4:52–65, 2023.

A. Conclusion

This paper presents DynaCREA, a novel framework for dynamic, multi-modal creativity assessment powered by reinforcement learning. DynaCREA advances the state-of-the-art by providing real-time, adaptive evaluation across multiple creativity dimensions—problem anchoring, originality, flexibility, fluency, and elaboration—while balancing accuracy, adaptability, and user engagement. Our experimental results demonstrate substantial gains in creativity metrics for the dynamic group compared to static baselines, with no loss in feasibility or outcome quality.

By integrating large language models, multi-modal interaction, and expert-verified signals, DynaCREA offers a transparent, scalable, and theoretically grounded approach to creativity assessment. While challenges remain in terms of scalability and real-world deployment, this work lays the foundation for future systems that can empower and measure human creativity in the age of artificial intelligence. Continued development in areas such as real-world data integration, technical robustness, and ethical safeguards will further enhance the impact and applicability of DynaCREA across diverse research and innovation domains.