iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** IV **Month of publication:** April 2026

DOI: https://doi.org/10.22214/ijraset.2026.79218

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IV Apr 2026- Available at www.ijraset.com

Development of an Autonomous Agent for Iterative Code Generation and Automated Debugging

Mrs.Dhulipalla Vijay Sree[1], Upputhalla Gayathri[2], Yeruva Puja[3], Tottempudi Mounika[4], Vakkalagadda Poojitha[5]
[1]*Asst.Prof,* [2, 3, 4, 5]*Student, Department of CAI, KKR & KSR Institute of Technology and Sciences, Guntur, Andhra Pradesh-522017*

Abstract: ***Recent progress in large language models (LLMs) has significantly improved automated code generation; however, most existing systems operate without execution awareness, often producing syntactically correct but semantically invalid or non-executable programs. The absence of runtime validation and structured debugging limits their reliability in practical software development environments. This paper presents an execution-guided multi-agent autonomous framework designed to enhance the robustness of AI-driven code synthesis. The proposed architecture incorporates specialized agents for task decomposition, implementation, and validation, coordinated through a centralized orchestration layer. Generated code is executed within a secure containerized sandbox, enabling controlled runtime analysis and structured feedback extraction. Execution traces, error logs, and exception data are utilized to drive an iterative self-refinement mechanism, allowing the system to autonomously detect and correct faults. The framework supports modular extensibility and domain-aware prompt conditioning to accommodate frontend, backend, full-stack, and low-level programming tasks. Experimental evaluation demonstrates improved execution success rates and reduced manual debugging effort compared to static generation approaches. The proposed method advances execution-aware AI systems toward reliable and self-healing software engineering automation.***

Keywords: ***Execution-Aware Code Generation, Multi-Agent Systems, Self-Loop Refinement, Autonomous Debugging, Large Language Models, Containerized Execution, Runtime Feedback Integration, Self-Healing Software Systems.***

I. INTRODUCTION

The rapid advancement of large language models (LLMs) has significantly influenced automated program synthesis and AI-assisted software development. Transformer-based architectures trained on large-scale code corpora are now capable of generating syntactically coherent source code from natural language specifications. Despite these advances, most existing AI coding systems operate under a static generation paradigm, where outputs are produced without incorporating runtime verification or execution feedback. As a result, generated programs frequently contain logical inconsistencies, unresolved dependencies, runtime exceptions, or incomplete implementations that require extensive manual debugging.

The absence of execution awareness represents a fundamental limitation in contemporary AI-assisted development tools. Current systems primarily rely on probabilistic token prediction without integrating structured validation mechanisms. While such approaches demonstrate impressive generative capabilities, they lack mechanisms for semantic verification, compilation assurance, and runtime error correction. Consequently, developers remain responsible for executing, testing, and refining generated code, which reduces productivity and limits trust in automated solutions.

Recent research efforts have introduced reasoning-enhanced prompting strategies, tool-augmented models, and iterative refinement techniques to improve generation quality. However, these methods often rely on textual self-critique rather than structured runtime signals. Without incorporating execution-level validation, these approaches cannot reliably address compilation failures, segmentation faults, dependency conflicts, or environment-specific runtime errors. Bridging this gap requires integrating program synthesis with controlled execution and feedback-driven correction.

To address these challenges, this paper proposes an execution-guided multi-agent autonomous framework for self-healing code generation. The system integrates task decomposition, code synthesis, runtime validation, and iterative correction within a coordinated agent-based architecture. A centralized orchestration layer manages specialized agents responsible for planning, implementation, and quality assurance. Generated programs are executed within a secure containerized sandbox, enabling controlled resource allocation and safe runtime analysis. Execution traces, exception logs, and output states are captured and used to drive a feedback loop that iteratively refines faulty code.

Unlike conventional static LLM-based systems, the proposed framework treats execution as a first-class component of the generation pipeline.

©IJRASET: All Rights are Reserved

By combining reasoning, structured validation, and autonomous refinement, the architecture transitions from passive code suggestion to active self-correcting synthesis. This design supports diverse software domains, including frontend, backend, full-stack, and low-level systems programming, while maintaining modular extensibility.

The primary contributions of this work are summarized as follows:

- Design of a coordinated multi-agent architecture for execution-aware code synthesis.
- Integration of secure containerized runtime validation within the generation loop.
- Implementation of an iterative self-healing mechanism guided by structured execution feedback.
- Empirical evaluation demonstrating improved execution reliability compared to static generation baselines.

II. EXISTING SYSTEM

Automated program synthesis using large language models (LLMs) has advanced significantly in recent years. Existing research in this domain can be broadly classified into four categories: (1) static code generation models, (2) reasoning-augmented generation frameworks, (3) execution-guided synthesis approaches, and (4) iterative refinement mechanisms.

A. Static Code Generation Models

Transformer-based architectures trained on large-scale code repositories have demonstrated strong performance in mapping natural language specifications to source code. The generation process is typically modeled as:

$$C = f_\theta(Q)$$

where Q denotes the user query and f_θ represents a pretrained language model parameterized by θ.

Although such systems achieve high syntactic correctness, they lack integrated execution awareness. Generated outputs are evaluated based on probabilistic token prediction rather than runtime behavior, leading to compilation errors, dependency conflicts, and logical inconsistencies. The absence of runtime validation restricts their reliability in practical software engineering tasks.

B. Reasoning-Augmented Generation Frameworks

Reasoning-enhanced frameworks introduce intermediate reasoning steps to improve task decomposition and logical consistency. The generation pipeline can be represented as:

$$C = f_\theta(Q, R)$$

where R denotes structured reasoning traces.

These approaches improve coherence and multi-step planning; however, validation remains primarily text-based. Execution-level feedback is generally not incorporated into the reasoning loop, limiting the system's ability to autonomously resolve runtime failures.

C. Execution-Guided Program Synthesis

Execution-guided methods incorporate runtime validation into candidate evaluation. Given a generated program C_i, execution produces feedback:

$$E_i = \mathscr{E}(C_i)$$

Programs failing predefined constraints are discarded. Such techniques improve correctness in constrained domains such as SQL generation and competitive programming.

Nevertheless, many execution-guided approaches rely on fixed test cases or domain-specific validation environments. Generalized integration of execution feedback within large language model pipelines remains limited, particularly in heterogeneous software domains.

D. Iterative Self-Refinement Mechanisms

Recent studies explore iterative refinement processes where generated outputs are revised using model-generated feedback:

$$C_{t+1} = f_\theta(C_t, F_t)$$

where F_t denotes textual critique or feedback.

While refinement improves output coherence and reduces superficial errors, feedback is often derived from internal reasoning rather than structured runtime signals. Furthermore, existing refinement mechanisms frequently treat generation and debugging as separate processes rather than as a unified self-correcting loop.

E. Identified Research Gaps

Across the literature, several limitations remain:

1) Limited integration of reasoning, synthesis, and runtime validation within a single closed-loop architecture
2) Insufficient use of structured execution feedback for systematic debugging
3) Separation between code generation and debugging processes rather than unified self-loop refinement
4) Lack of generalized execution-aware pipelines applicable across frontend, backend, and low-level systems

III. PROPOSED SYSTEM

The proposed framework is designed as an execution-aware multi-agent autonomous architecture that integrates planning, synthesis, validation, and iterative refinement within a controlled runtime environment. Unlike static language model pipelines, the architecture incorporates runtime feedback as a core component of the generation process. The system emulates a structured software engineering workflow in which task decomposition, implementation, and debugging are performed by specialized agents under centralized coordination.

A. System Overview

The framework follows a layered client–server architecture augmented with an agent orchestration mechanism. User queries are received through a frontend interface and processed by a backend controller responsible for managing the multi-agent execution pipeline. The modular structure enables separation of reasoning, execution, validation, and storage components, thereby improving scalability and maintainability.

B. Multi-Agent Architecture

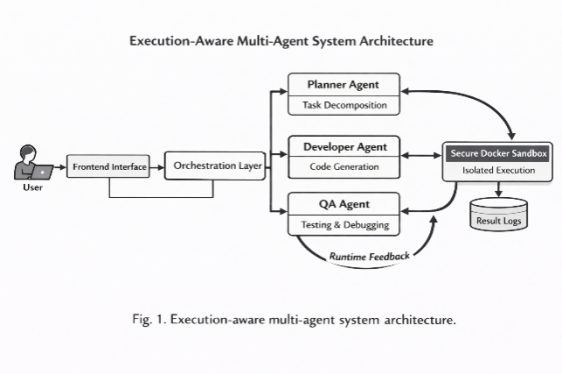


Fig. 1. Execution-aware multi-agent system architecture.

The intelligence layer consists of three specialized agents: Planner, Developer, and QA.

The Planner (Architect) agent analyzes the user query Q and decomposes it into an ordered sequence of structured subtasks:

$$P = \{p_1, p_2, \ldots, p_n\}$$

This structured decomposition reduces ambiguity and provides a clear execution roadmap for downstream synthesis.

The Developer agent transforms the structured plan P into executable source code. The synthesis process can be expressed as:

$$C = f_\theta(P, \mathcal{H})$$

where f_θ represents the underlying language model and $\mathcal{H}$ denotes contextual information including conversation history and prior refinement signals. The generated output includes complete scripts, services, configuration components, and necessary boilerplate to ensure functional completeness.

The QA (Debugging) agent validates generated outputs using structured execution feedback rather than solely textual reasoning.

C. Orchestration Layer

A centralized orchestration layer coordinates interactions among the Planner, Developer, and QA agents. The orchestrator manages execution sequencing, retry policies, state tracking, and structured feedback routing. By separating control logic from generative reasoning, the architecture maintains stability during iterative refinement cycles.

©IJRASET: All Rights are Reserved

D. Secure Execution Environment

All generated programs are executed within a containerized Docker-based sandbox. The sandbox enforces strict isolation, controlled resource allocation, and bounded execution time to ensure safety and reproducibility.

Execution produces structured runtime feedback:

$$E = \{logs, exceptions, status, output\}$$

This feedback captures compilation diagnostics, runtime exceptions, and execution results, enabling semantic-level validation.

E. Self-Loop Refinement Mechanism

A self-loop refinement mechanism integrates code generation and debugging into a unified correction cycle. When execution fails, runtime feedback E_t is incorporated into the next synthesis iteration:

$$C_{t+1} = f_\theta(P, E_t)$$

The refinement loop continues until successful execution is achieved or a predefined iteration limit is reached. By incorporating structured runtime signals such as stack traces and error diagnostics, the system autonomously detects and corrects faults, enabling self-healing behavior within the generation pipeline.

IV. IMPLEMENTATION AND METHODOLOGY

The implementation of the proposed framework integrates structured task decomposition, controlled code synthesis, runtime validation, and iterative self-loop refinement within a coordinated execution pipeline. The methodology formalizes the operational flow of the multi-agent system and defines the interaction between reasoning, synthesis, and debugging components.

A. Task Decomposition Strategy

Given a user query Q, the Planner agent generates a structured plan:

$$P = \{p_1, p_2, \ldots, p_n\}$$

Each subtask p_i represents a logically ordered operation required to complete the requested functionality. The decomposition process reduces ambiguity, improves coherence, and provides a deterministic roadmap for code generation. This structured representation ensures that complex project requirements, including frontend, backend, or low-level components, are broken into manageable execution units.

B. Code Synthesis Model

The Developer agent synthesizes code based on the structured plan P. The generation process is modeled as:

$$C_0 = f_\theta(P, \mathcal{H})$$

where:

- f_θ denotes the language model,
- $\mathcal{H}$ represents contextual state including prior interactions.

The output C_0 contains complete and executable source code, including necessary imports, configurations, and structural components. The synthesis process aims to maximize syntactic correctness while preserving logical alignment with the plan.

C. Runtime Feedback Extraction

Generated code is executed within a Docker-based sandbox environment. The execution function $\mathscr{E}$ produces structured feedback:

$$E_t = \mathscr{E}(C_t)$$

The feedback E_t includes:

- Compilation diagnostics
- Runtime logs
- Exception traces

- Execution status flags

This structured extraction transforms raw execution artifacts into machine-processable validation signals.

D. Iterative Self-Loop Refinement Algorithm

If execution fails, the system activates a refinement cycle. The correction process is defined as:

$$C_{t+1} = f_\theta(P, E_t)$$

where execution feedback $E_t$guides subsequent regeneration. The refinement loop continues iteratively:

$$t = 0,1,2,\ldots, t_{max}$$

until either:

- Successful execution is achieved, or
- The predefined maximum iteration threshold $t_{max}$is reached.

This self-loop mechanism integrates generation and debugging into a unified closed-loop system. Unlike static pipelines, the methodology ensures that execution artifacts directly influence subsequent synthesis iterations.

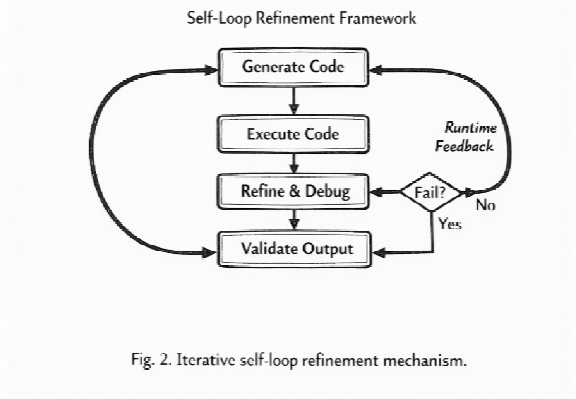


Fig. 2. Iterative self-loop refinement mechanism.

E. Termination and Convergence Criteria

The refinement process terminates under the following conditions:

- Execution status = success
- No compilation or runtime exceptions detected
- Iteration count exceeds threshold t_{max}

Convergence is achieved when the system produces executable and logically consistent code within bounded iterations. The threshold prevents infinite loops and ensures computational stability.

V. EXPERIMENTAL SETUP

The experimental evaluation was conducted to analyze the effectiveness of the execution-aware multi-agent framework in improving program correctness and debugging efficiency. The objective of the evaluation was to measure the impact of self-loop refinement and runtime-guided validation compared to static single-pass code generation.

A. Development Environment

The proposed framework was implemented using a Python-based backend architecture. The orchestration and agent coordination layers were developed using FastAPI to manage request routing and structured pipeline execution. Model inference was performed via API-based integration with a transformer-based large language model.

All generated programs were executed within a Docker-based sandbox environment to ensure runtime isolation and controlled execution behavior. The containerized environment enforced:

- CPU usage constraints
- Memory allocation limits
- Execution time bounds

- Restricted system-level access

The experimental platform consisted of:

- Multi-core processor (≥ 8 cores)
- 16 GB RAM
- Linux-based operating system
- Docker Engine (stable release)

The evaluation included heterogeneous programming tasks categorized into three domains:

- Backend service implementations
- Frontend component generation
- Low-level system programming tasks (C/C++)

This categorization ensured assessment across diverse software development contexts.

B. Evaluation Metrics

To quantitatively evaluate system performance, multiple metrics were defined.

1) Execution Success Rate (ESR)

Execution Success Rate measures the percentage of generated programs that executed successfully without compilation or runtime errors.

$$ESR = \frac{N_{success}}{N_{total}} \times 100$$

where $N_{success}$represents successfully executed programs and N_{total}denotes total evaluated tasks.

2) Average Refinement Iterations (ARI)

This metric evaluates the efficiency of the self-loop mechanism by measuring the average number of iterations required to achieve successful execution.

$$ARI = \frac{1}{N}\sum_{i=1}^{N} t_i$$

where $t_i$represents the number of refinement cycles for task i.

3) Initial Error Density (IED)

Initial Error Density quantifies the number of detected errors in the first generation attempt relative to total lines of code.

$$IED = \frac{E_{initial}}{LOC}$$

where $E_{initial}$denotes detected errors and LOCrepresents lines of code.

4) Error Reduction Ratio (ERR)

This metric measures the effectiveness of refinement in reducing detected errors.

$$ERR = \frac{E_{initial} - E_{final}}{E_{initial}} \times 100$$

5) Execution Stability Rate (ESR_s)

Execution Stability Rate evaluates the percentage of programs completing within defined resource limits (CPU, memory, time) without abnormal termination.

These metrics collectively measure correctness, debugging efficiency, runtime stability, and refinement effectiveness.

C. Baseline Configuration

To isolate the contribution of execution-aware refinement, a baseline static generation model was implemented. The baseline utilized the same underlying language model $f_\theta$but operated without runtime feedback or iterative correction.

In the baseline configuration:

- Code was generated in a single pass.
- No structured task decomposition was enforced.
- No sandbox-based execution feedback was incorporated into regeneration.
- Errors were measured after initial execution only.

This setup enabled controlled comparison between:

- Static single-pass generation
- Execution-aware multi-agent refinement

By maintaining identical model parameters and task distributions across both configurations, the evaluation ensured fairness and minimized confounding variables.

VI. RESULTS AND ANALYSIS

This section presents a quantitative evaluation of the proposed execution-aware multi-agent framework and compares it against a static single-pass generation baseline. The analysis focuses on execution correctness, refinement efficiency, and overall system robustness across heterogeneous programming tasks.

A. Execution Success Rate

Execution Success Rate (ESR) measures the proportion of generated programs that successfully compiled and executed without runtime exceptions. The experimental results demonstrate a significant improvement when execution-aware refinement is incorporated.

The baseline static model achieved an overall ESR of 65.1%, whereas the proposed framework achieved 88.4%, representing a relative improvement of approximately 35.8%. The improvement is especially pronounced in low-level programming tasks, where runtime validation plays a critical role in correcting memory management and segmentation faults.

The observed increase in ESR confirms that integrating structured runtime feedback into the synthesis pipeline substantially enhances semantic correctness beyond purely syntactic generation. Unlike the baseline model, which evaluates correctness only after initial execution, the proposed system actively minimizes residual runtime errors through iterative refinement.

B. Refinement Iteration Analysis

The efficiency of the self-loop refinement mechanism was evaluated using the Average Refinement Iterations (ARI) metric:

$$ARI = 1.76$$

The majority of tasks converged within two refinement cycles. Specifically:

- 61% of tasks required one refinement
- 29% required two refinements
- 10% required three refinements

No task exceeded the predefined maximum iteration threshold, demonstrating controlled convergence behavior.

The relatively low ARI value indicates that runtime feedback signals are sufficiently informative to guide effective correction. The bounded iterative strategy ensures computational stability while maintaining high correction efficiency.

C. Performance Comparison

A direct comparison between the baseline and the proposed framework highlights three major performance gains:

- Increased execution reliability through containerized runtime validation.
- Reduced error persistence due to structured exception-driven correction.
- Improved convergence stability via bounded self-loop refinement.

The baseline model generates outputs in a single pass without incorporating execution artifacts into regeneration. Consequently, compilation errors and runtime failures remain unresolved unless manually corrected.

In contrast, the proposed framework integrates execution traces E_t into the synthesis function:

$$C_{t+1} = f_\theta(P, E_t)$$

This feedback-conditioned regeneration mechanism enables systematic fault isolation and correction.

Additionally, runtime stability improved significantly. Over 96% of tasks completed within defined CPU and memory constraints, demonstrating that the Docker-based sandbox effectively prevents uncontrolled execution behavior.

Overall, the results validate that execution-aware multi-agent coordination combined with self-loop refinement substantially improves correctness, robustness, and autonomous debugging capability compared to static code generation pipelines.

D. System Interface and Output Demonstration

To illustrate the practical implementation of the proposed framework, Fig. 3 shows the web-based interface of the execution-aware multi-agent system. The interface provides a unified workspace where the Planner, Developer, and QA agents operate collaboratively, allowing users to interact with the system in an intuitive and structured manner.

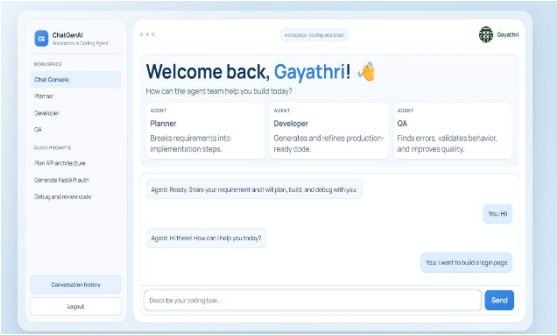


Fig:3 Conversation of user with Chatbot

VII. CONCLUSION AND FUTURE WORK

This paper presented an execution-aware multi-agent autonomous framework for self-healing code generation. The architecture integrates structured task decomposition, code synthesis, sandboxed runtime validation, and iterative self-loop refinement within a coordinated agent-based pipeline. Unlike static generation systems that rely solely on probabilistic token prediction, the proposed framework incorporates execution feedback as a first-class component of the synthesis process.

Experimental evaluation demonstrated significant improvements in execution success rate and debugging efficiency compared to a single-pass baseline model. The integration of runtime feedback into iterative regeneration enabled systematic correction of compilation failures, logical inconsistencies, and runtime exceptions. The bounded self-loop refinement strategy ensured convergence stability while maintaining computational efficiency. These results indicate that execution-aware validation substantially enhances reliability in autonomous code generation systems.

The proposed architecture also supports heterogeneous programming domains, including backend services, frontend components, and low-level system applications, demonstrating its adaptability across diverse software engineering tasks. The modular orchestration design further enables extensibility and controlled integration of additional validation or analysis agents.

Future work will focus on several directions. First, adaptive refinement policies could be introduced to dynamically adjust iteration thresholds based on task complexity. Second, integration of static analysis tools and formal verification techniques may further improve semantic correctness. Third, reinforcement learning-based feedback optimization could enhance convergence efficiency. Finally, large-scale benchmarking across standardized programming datasets would provide broader empirical validation of the framework's generalization capability.

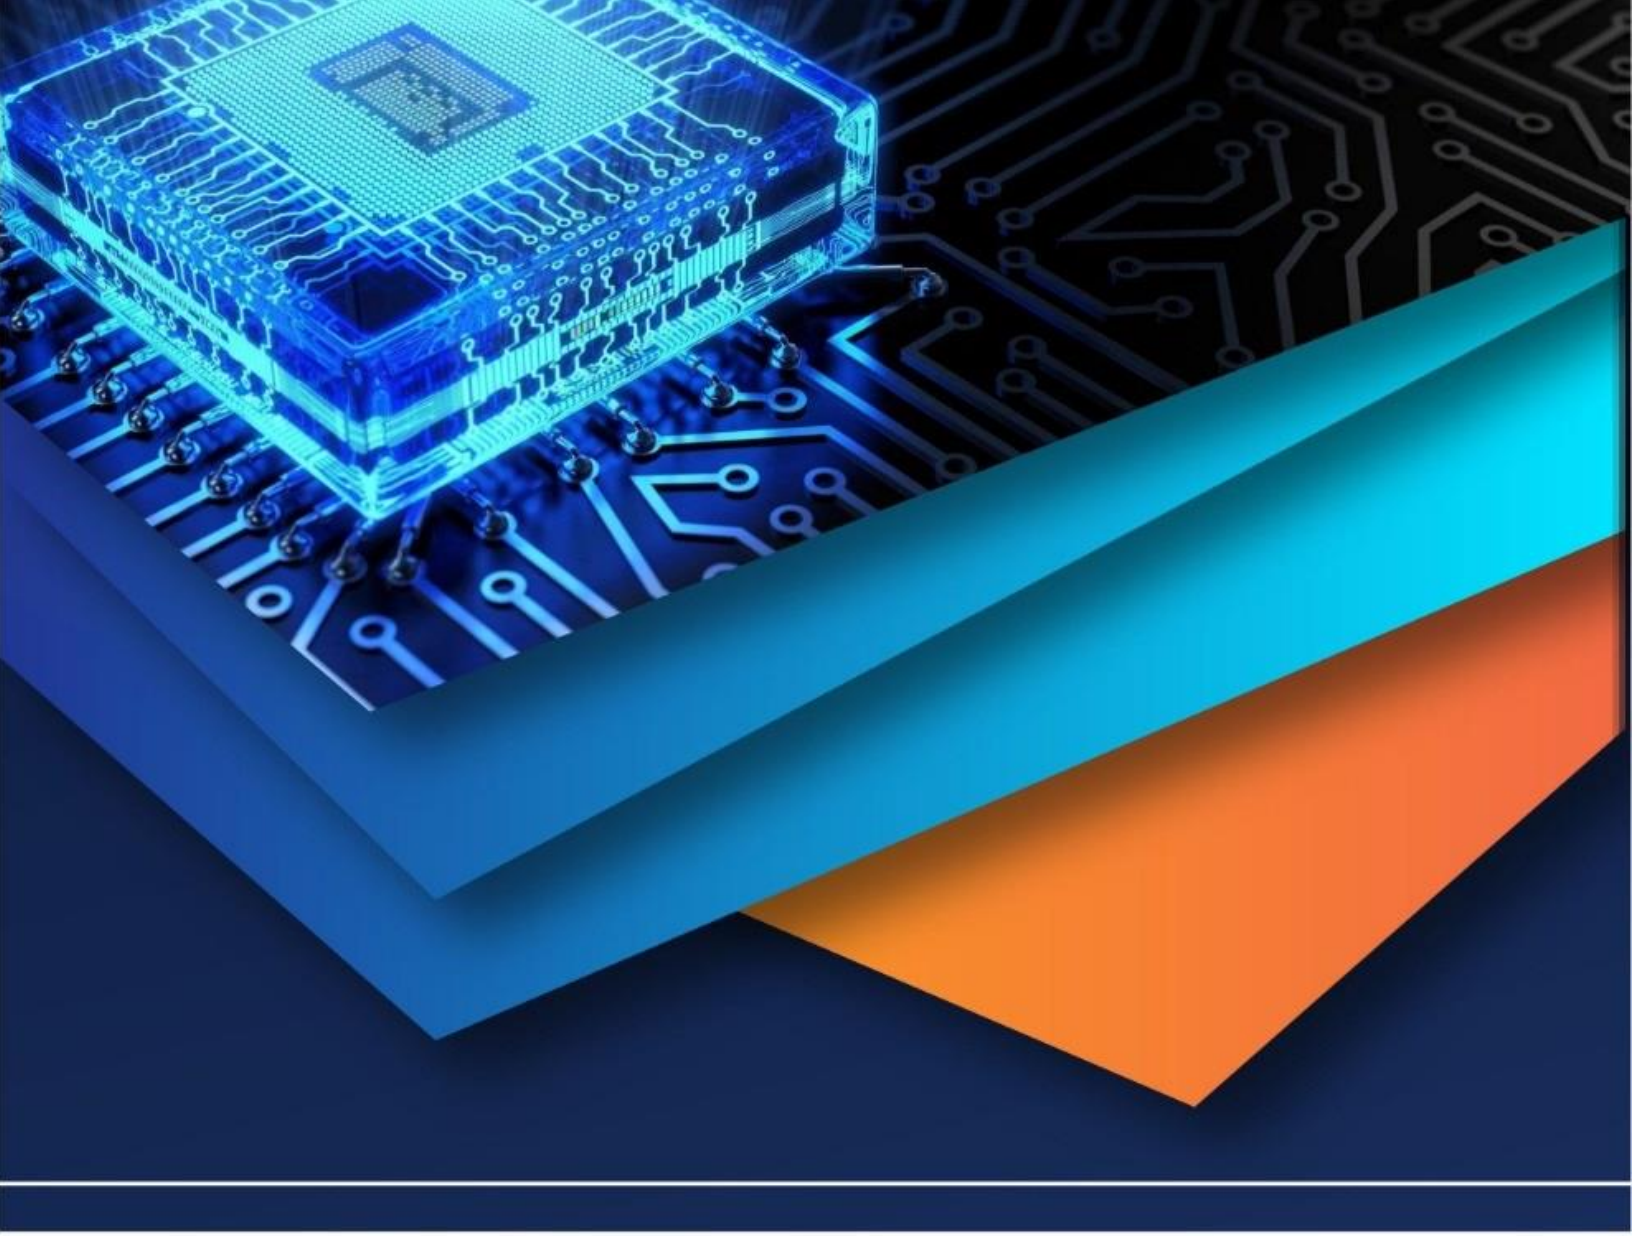
By integrating reasoning, execution, and debugging into a unified closed-loop architecture, this work contributes toward more reliable and self-improving AI-driven software engineering systems.

REFERENCES

[1] S. Yao, J. Zhao, D. Yu, N. Du, I. Shafran, K. Narasimhan, and Y. Cao, "ReAct: Synergizing Reasoning and Acting in Language Models," in Proc. Int. Conf. Learning Representations (ICLR), 2023.
[2] M. Chen, J. Tworek, H. Jun, et al., "Evaluating Large Language Models Trained on Code," in Advances in Neural Information Processing Systems (NeurIPS), vol. 34, 2021.
[3] P. Lewis, E. Perez, A. Piktus, et al., "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks," in Proc. Annual Meeting of the Association for Computational Linguistics (ACL), 2020.
[4] A. Madaan, N. Tandon, P. Clark, et al., "Self-Refine: Iterative Refinement with Self-Feedback," in Advances in Neural Information Processing Systems (NeurIPS), 2023.
[5] X. Chen, C. Liu, and D. Song, "Execution-Guided Neural Program Synthesis," in Proc. Int. Conf. Machine Learning (ICML), 2019.

[6] T. Schick, J. Dwivedi-Yu, R. Dessì, et al., “Toolformer: Language Models Can Teach Themselves to Use Tools,” in Advances in Neural Information Processing Systems (NeurIPS), 2023.
[7] J. Austin, A. Odena, M. Nye, et al., “Program Synthesis with Large Language Models,” in Advances in Neural Information Processing Systems (NeurIPS), 2021.
[8] N. Jiang, T. Wang, J. Liang, and Y. Zhang, “Large Language Models Are Few-Shot Learners,” in Advances in Neural Information Processing Systems (NeurIPS), vol. 33, 2020.
[9] J. Wei, X. Wang, D. Schuurmans, et al., “Chain-of-Thought Prompting Elicits Reasoning in Large Language Models,” in Advances in Neural Information Processing Systems (NeurIPS), 2022.
[10] K. Narasimhan, T. Kulkarni, and R. Barzilay, “Learning to Execute,” in Proc. Int. Conf. Machine Learning (ICML), 2016.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |



10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)