

CAM: A Causality-Based Analysis Framework for Multi-agent Code Generation Systems

ZONGYI LYU, The Hong Kong University of Science and Technology, China and Guangzhou HKUST Fok Ying Tung Research Institute, China

ZHENLAN JI*, Nara Institute of Science and Technology, Japan

SONGQIANG CHEN, The Hong Kong University of Science and Technology, China and Guangzhou HKUST Fok Ying Tung Research Institute, China

LIWEN WANG, The Hong Kong University of Science and Technology, China

YUHENG HUANG, The University of Tokyo, Japan

SHUAI WANG, The Hong Kong University of Science and Technology, China

SHING-CHI CHEUNG*, The Hong Kong University of Science and Technology, China and Guangzhou HKUST Fok Ying Tung Research Institute, China

Despite the remarkable success that Multi-Agent Code Generation Systems (MACGS) have achieved, the inherent complexity of multi-agent architectures produces substantial volumes of intermediate outputs. To date, the individual importance of these intermediate outputs to the system correctness remains opaque, which impedes targeted optimization of MACGS designs. To address this challenge, we propose CAM, the first Causality-based Analysis framework for MACGS that systematically quantifies the contribution of different intermediate features to system correctness. By comprehensively categorizing intermediate outputs and systematically simulating realistic errors on intermediate features, we identify the important features for system correctness and aggregate their importance rankings, facilitating comprehensive analysis of MACGS.

We instantiate CAM on representative MACGS across multiple backend LLMs and datasets and conduct extensive empirical analysis on the identified importance rankings. Our analysis reveals intriguing findings: first, we uncover context-dependent features—features whose importance emerges mainly through interactions with other features, revealing that quality assurance for MACGS should move beyond module-level validation to incorporate cross-feature consistency checks; second, we reveal that hybrid backend MACGS with different backend LLMs assigned according to their relative strength achieves up to 7.3% Pass@1 improvement, underscoring hybrid architectures as a promising direction for future MACGS design. We further demonstrate CAM’s practical utility through two applications: (1) failure repair, which achieves a 73.6% success rate by optimizing top-3 importance-ranked features and (2) feature pruning, that reduces up to 33.6% intermediate token consumption with negligible or sometimes positive performance impact by pruning low-importance

*Corresponding authors.

Authors’ Contact Information: **Zongyi Lyu**, The Hong Kong University of Science and Technology, Hong Kong, China and Guangzhou HKUST Fok Ying Tung Research Institute, Guangzhou, China, zlyuaj@connect.ust.hk; **Zhenlan Ji**, Nara Institute of Science and Technology, Nara, Japan, ji.zhenlan@naist.ac.jp; **Songqiang Chen**, The Hong Kong University of Science and Technology, Hong Kong, China and Guangzhou HKUST Fok Ying Tung Research Institute, Guangzhou, China, i9s.chen@connect.ust.hk; **Liwen Wang**, The Hong Kong University of Science and Technology, Hong Kong, China, lwanged@cse.ust.hk; **Yuheng Huang**, The University of Tokyo, Tokyo, Japan, yuhenghuang42@g.ecc.u-tokyo.ac.jp; **Shuai Wang**, The Hong Kong University of Science and Technology, Hong Kong, China, shuaiw@cse.ust.hk; **Shing-Chi Cheung**, The Hong Kong University of Science and Technology, Hong Kong, China and Guangzhou HKUST Fok Ying Tung Research Institute, Guangzhou, China, scc@cse.ust.hk.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2026 Copyright held by the owner/author(s).

ACM 2994-970X/2026/10-ARTISSTA147

<https://doi.org/10.1145/3832238>

features. Our work provides actionable insights for MACGS design and deployment, establishing causality analysis as a powerful approach for understanding and improving MACGS.

CCS Concepts: • **Software and its engineering** → *Software defect analysis; Automatic programming.*

Additional Key Words and Phrases: multi-agent systems for code generation, causal analysis

ACM Reference Format:

Zongyi Lyu, Zhenlan Ji, Songqiang Chen, Liwen Wang, Yuheng Huang, Shuai Wang, and Shing-Chi Cheung. 2026. CAM: A Causality-Based Analysis Framework for Multi-agent Code Generation Systems. *Proc. ACM Softw. Eng.* 3, ISSTA, Article ISSTA147 (October 2026), 24 pages. <https://doi.org/10.1145/3832238>

1 Introduction

Multi-Agent Code Generation Systems (MACGS) have emerged as a transformative paradigm in automated software development, demonstrating remarkable capabilities in generating high-quality code through sophisticated agent collaboration [13, 25, 27, 51, 77]. By decomposing complex programming tasks into specialized subtasks and coordinating multiple agents with distinct roles, MACGS have achieved substantial improvements over single-LLM approaches [19, 21, 47, 71] across diverse benchmarks [2, 7, 38, 75]. Despite their demonstrated effectiveness, MACGS present fundamental challenges in understanding and analyzing their intermediate outputs. The multi-agent architecture inherently produces diverse and complex intermediate outputs across agents [68], making it exceptionally difficult to comprehend their impact on the overall system. This opacity substantially compromises the robustness of these new-generation software—LLM-based agent systems—and impedes further optimization of MACGS designs.

Several approaches can be applied to analyze intermediate outputs in MACGS, yet each suffers from notable limitations. Manual analysis, while intuitive, is prohibitively expensive and inherently subjective, limiting its scalability and reproducibility. LLM-based evaluation methods [48, 79, 82] offer automation but exhibit low accuracy and introduce reliability concerns, as shown by our preliminary study in Sec. 3.2. Currently, MACGS developers predominantly rely on all-but-one ablation designs that disable entire agents or modules to assess their impact on final code quality [13, 40, 77]. However, such coarse-grained approaches fail to capture the nuanced influence of specific intermediate components and lack generalizability across different MACGS architectures. Collectively, these limitations underscore the need for a fine-grained analysis method that balances analytical rigor with practical cost.

To address these limitations, we leverage *actual causality* [23] to analyze the causal relationships between intermediate outputs and final code correctness in MACGS. However, conducting causal analysis on MACGS requires overcoming substantial technical challenges. First, it is challenging to decompose complex intermediate outputs into structured features and construct the causal graph for causal analysis. Second, causality-based methods typically rely on a substantial amount of data to guarantee the reliability of analysis [14, 45, 72]. In the context of MACGS, this requirement manifests as the systematic simulation of a variety of realistic execution states, particularly errors, which is a non-trivial task [68]. Third, the computational cost of MACGS execution is another noticeable obstacle. Unlike traditional software systems [60, 67] that operate with minor cost, merely one single execution of MACGS may involve multiple LLM calls, with up to 10^4 LLM token consumption [68]. This prohibitive cost makes it infeasible to exhaustively explore all execution states of the target software like prior work [14]. Instead, it is imperative to design an efficient execution state exploration algorithm that can strategically make trade-offs between analysis comprehensiveness and computational cost.

To tackle these challenges, we present CAM, the first **Causality-based Analysis** framework for **MACGS** to identify the relative importance of different intermediate components to final code

correctness. Specifically, we first develop a systematic categorization method to model complex intermediate outputs into a variety of variables, or *features*, which capture key aspects of intermediate outputs from various perspectives and are generalizable across various MACGS. We then construct a causal graph based on these features to facilitate subsequent causal analysis. To systematically simulate realistic errors, we design an LLM-based approach that comprehensively conducts counterfactual intervention on intermediate outputs. Additionally, regarding the execution state exploration challenge, we introduce a novel notion, *influence set*, that reflects the unique error propagation mechanism of multi-agent systems considering LLMs' self-correction ability [11]. Based on this insight, our proposed algorithm can efficiently prune unnecessary states and achieve an ideal balance between reliability and computational cost.

To gauge the reliability of CAM, we empirically validate CAM's reliability by comparing its identified feature importance against manual analysis, demonstrating strong agreement with Kendall's correlation coefficients [34] ranging from 0.76 to 0.91. On this basis, we conduct a comprehensive analysis of feature importance patterns that lead to intriguing findings for MACGS design and deployment. First, we uncover a previously overlooked phenomenon: *context-dependent features*—features whose importance manifests primarily through interactions with other features. For instance, in 78.8% of cases, `Program_Lang` (programming language for code implementation) affects the system correctness only when simultaneously intervened with other features. This context-dependency indicates that failures arise not only from individual feature errors, but may also stem from subtle incompatibilities between seemingly correct intermediate features. This finding reveals a critical insight for improving MACGS: quality assurance for MACGS should move beyond module-level validation to incorporate cross-feature consistency checks. Furthermore, we reveal that different backend LLMs exhibit distinct capabilities across subtasks, motivating hybrid multi-backend architectures for MACGS, where different backend LLMs are assigned to specific stages. Empirical validation of a hybrid backend MACGS achieves up to 7.3% Pass@1 improvement compared to uniform backend configurations, underscoring the potential of hybrid architectures as a promising direction for future MACGS design. Moreover, we demonstrate CAM's practical utility through two causality-guided downstream applications: (1) failure repair, which achieves a 73.6% success rate by optimizing only the top-3 ranked features, and (2) feature pruning, which reduces intermediate output token consumption by up to 33.6% with negligible or sometimes positive performance impact. We further demonstrate the generalizability of CAM to prominent MACGS and a reasoning model backend, establishing CAM as a broadly applicable analysis tool.

Overall, we summarize our contributions as follows:

- We identify and address a critical challenge in MACGS: understanding how intermediate outputs causally influence final code correctness. For the first time, we introduce actual causality to systematically analyze MACGS, establishing causality analysis as an effective approach for understanding and improving MACGS.
- We propose CAM, a comprehensive causality-based analysis framework that automatically quantifies each intermediate feature's importance on system correctness. CAM first models complex intermediate outputs into structured features, and then systematically identifies critical features by efficiently conducting counterfactual interventions through influence set analysis.
- We conduct extensive empirical analysis on representative MACGS across multiple backend LLMs and datasets, yielding actionable insights for MACGS design and deployment. We further demonstrate the practical utility of CAM through two successful applications: causality-guided failure repair and feature pruning.

2 Background

2.1 Multi-agent Code Generation Systems

With the rapid advancement of LLM-based agents [24, 74, 78, 80], multi-agent code generation systems (MACGS) have emerged as a predominant paradigm for automated code generation [13, 25, 27, 51, 77]. The fundamental methodology underlying MACGS involves decomposing the code generation process into a sequence of specialized subtasks, each delegated to distinct agents with specific responsibilities. The division of labor enables MACGS to develop a more comprehensive and nuanced understanding of programming tasks compared to single-LLM approaches [17, 22, 30].

A typical workflow of MACGS includes three main phases: planning, coding and refinement [24, 41, 77]. Upon receiving a user's problem specification, MACGS conduct multi-step planning where the original problem undergoes systematic decomposition, requirement analysis, and implementation design. This planning process serves as a critical guidance mechanism for subsequent code generation [42], producing substantial intermediate outputs that capture different facets of problem understanding. Following the planning phase, MACGS proceeds to generate code based on the planning outputs and conducts iterative refinement to ensure code quality and correctness. Among the various MACGS architectures, MetaGPT [25] stands as one of the most influential and widely-adopted frameworks, inspiring numerous subsequent MACGS implementations [27, 43, 77].

While MACGS facilitate sophisticated problem solving, they generate substantial volumes of intermediate outputs in addition to the final code [68], including problem understanding, algorithm design, and implementation approaches [24, 41]. The individual contributions of these intermediate outputs to the final code remain opaque, which hinders the targeted optimization of MACGS [39, 57].

2.2 Actual Causality

As a canonical technique that is extensively applied to analyze complex software systems [3, 14, 45, 72], actual causality [23] is effective for identifying the underlying causes of observed behaviors and outcomes within intricate systems. Typically, actual causality proceeds by examining counterfactual scenarios on causal graph: given an observed outcome, actual causality determines which variables are the cause for the outcome to occur. In this paper, we leverage the following concepts in actual causality to systematically analyze MACGS.

Definition 1. (Causal graph). In actual causality, a causal graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a directed acyclic graph (DAG) that encodes the causal dependencies among variables in the system. Each vertex in $\mathcal{V}$ represents a variable, and each directed edge $(V_i, V_j) \in \mathcal{E}$ indicates that V_i directly influences V_j . The acyclic structure ensures a well-defined causal ordering and prevents circular dependencies that would render causal analysis ill-posed. In this paper, we construct causal graph for MACGS through systematic decomposition in Sec. 4.1.

Definition 2. (Actual cause). An actual cause [23] is defined through three conditions (AC1, AC2, AC3) that respectively capture actuality, counterfactual dependence, and minimality. Formally, $\vec{X} = \vec{x}$ is an *actual cause* of ϕ in the causal setting $(M, \vec{u})$ if the following three conditions hold:

AC1 (Actuality): $M, \vec{u} \models \vec{X} = \vec{x} \wedge \phi$. This ensures that both the cause and the effect actually occur.

AC2 (Counterfactual Dependence): There exists a set $\vec{W}$ of variables in $\mathcal{V}$ and a setting $\vec{x}'$ of the variables in $\vec{X}$ such that if $(M, \vec{u}) \models \vec{W} = \vec{w}^*$, then: $(M, \vec{u}) \models [\vec{X} \leftarrow \vec{x}', \vec{W} \leftarrow \vec{w}^*] \neg \phi$. This condition captures the counterfactual dependence of ϕ on $\vec{X} = \vec{x}$ under the contingency $\vec{W} = \vec{w}^*$.

AC3 (Minimality): $\vec{X}$ is minimal; no proper subset of $\vec{X}$ satisfies AC1 and AC2: if $\{\vec{X}, \vec{Y}\}$ causes outcome ϕ but $\{\vec{X}\}$ alone suffices, then only $\{\vec{X}\}$ constitutes the actual cause, not $\{\vec{X}, \vec{Y}\}$.

In this work, we analyze the intermediate outputs of MACGS through systematically identifying the important outputs to system correctness, leveraging the definition of actual cause.

Definition 3. (Responsibility) The responsibility quantifies the relative importance of different causes. For a cause $\vec{X} = \vec{x}$ of ϕ , its responsibility is defined as:

$$dr(\vec{X} = \vec{x}, \phi) = \frac{1}{1 + |\vec{W}|} \quad (1)$$

where $|\vec{W}|$ denotes the minimum cardinality of contingency sets $\vec{W}$ required to establish the counterfactual necessity of $\vec{X} = \vec{x}$. Intuitively, responsibility measures how many additional variables must be held or fixed for the causal effect to manifest. In this work, we measure the importance of different features in MACGS through aggregating their corresponding contribution, which is inspired by the concept of responsibility in actual causality.

3 Motivation

3.1 Significance of Intermediate Outputs

While MACGS demonstrate impressive capabilities in automated code generation, they suffer from notable robustness challenges [42]. The generation process in MACGS is inherently complex, involving multiple stages of information transformation across multiple agents [24, 41]. Consequently, errors introduced in the intermediate outputs can propagate through the pipeline and ultimately compromise system correctness [11]. To empirically investigate this impact, we randomly sample 100 failed cases from MetaGPT [25] and manually analyze their root causes. We identify three main failure causes: *Intermediate flaw* (flawed intermediate outputs generated by agents), *Overthinking* (exceeding iteration limits without reaching solution), and *Underthinking* (premature termination before sufficient exploration). We then attempt to recover these failures and re-execute MetaGPT: for intermediate flaws, we manually correct the erroneous outputs; for overthinking and underthinking, we identify the problematic phase and prompt re-execution from that checkpoint, following existing works [65, 70]. As shown in Table 1, 83% of failures stem from intermediate flaws, substantially higher than overthinking (12%) and underthinking (5%), highlighting the importance of intermediate outputs in system correctness. Moreover, correcting intermediate flaws resolved 78 out of 83 cases, demonstrating that intermediate output errors are both the most prevalent and most recoverable. These findings underscore the criticality of systematically exploring the importance of intermediate outputs and analyzing their robustness and impact on the overall system.

Table 1. Distribution and recoverability of failure causes.

Failure Cause	Count	# Fixed
Intermediate flaw	83	78
Overthinking	12	4
Underthinking	5	2
Total	100	82

3.2 Analytical Challenges and Limitations of Existing Methods

Despite their significance, the large volume of intermediate outputs generated by MACGS presents formidable analytical challenges. The complex dependencies among these outputs make it exceedingly difficult to disentangle their individual contributions to final outcomes, underscoring the pressing need for systematic analysis. Before introducing our proposed framework, we first examine existing methods and discuss their inherent limitations.

A straightforward approach to tackle this challenge is through manual inspection. Domain experts can examine failed code generation instances, trace through intermediate outputs, and determine the important parts. However, manual inspection suffers from severe scalability limitations and is inherently subjective, making it unsuitable for systematic evaluation.

Table 2. Accuracy of LLM-based automatic cause identification.

LLM	Accuracy
GPT-4o	41.0%
Qwen-2.5-Coder	37.0%
DeepSeek-Coder-V2	34.0%

An alternative approach leverages LLM reasoning capabilities [48, 79, 82]. We evaluate this through a preliminary study where we task state-of-the-art LLMs to identify important intermediate outputs responsible for MACGS failures and compare the results with human annotations. Specifically, we randomly select 100 MetaGPT failure cases and provide the LLM with all intermediate outputs and the final incorrect code. We then prompt it to identify failure-causing outputs. As shown in Table 2, even the best model (GPT-4o) achieves only 41% accuracy. This poor performance may stem from LLMs frequently identifying outputs merely correlated with failures rather than actually causing them, aligning with recent findings that LLMs struggle with analyzing complex chain processes [31, 35, 73].

Other approaches rely on all-but-one studies that remove entire agents or modules to examine their effects on MACGS performance [13, 43, 77]. However, such approaches suffer from significant limitations in analytical granularity. Agent outputs typically comprise multiple parts with distinct semantics and functions, yet all-but-one studies treat them as monolithic units [13, 27, 40]. This overlooks the intricate coupling among different parts: they are often cross-dependent within agents, and their collective effects on downstream performance are non-additive [8, 50]. Consequently, removing an entire agent conflates distinct contributions of different parts, yielding unreliable results that lack actionable insights. For example, in MetaGPT, the output of product manager can be decomposed into several semantic fields, directly disabling it fails to identify the individual contribution of each field. Therefore, we aim to answer the following critical question: how to systematically analyze the intermediate outputs' importance on the final output of MACGS?

3.3 Actual Causality for MACGS Analysis

As introduced in Sec. 2.2, actual causality [23, 49] provides an ideal foundation for analyzing MACGS intermediate outputs. Unlike correlation-based analysis [46, 52], which merely identifies statistical associations between intermediate outputs and failures, actual causality establishes genuine causal relationships through counterfactual reasoning and systematically conducts interventions to distinguish causation from correlation [14, 72].

As illustrated in Fig. 1, suppose MACGS receive the task of generating a stable sort function, with the following intermediate outputs: (1) Design *D* specifies a quicksort-based algorithm, and (2) Implementation steps *I* details the logic. The final code based on *D* and *I* implements quicksort, which is an unstable sort. Using correlation-based analysis, we might observe that both *D* and *I* are present in most failed executions of this task, but this does not tell us which was the actual cause. In contrast, actual causality systematically tests counterfactuals: Would replacing *D* with a bubble-sort algorithm fix the failure? Would replacing *I* alone suffice? Through these interventions, actual causality determines that *D*'s choice of quicksort is the minimal sufficient cause, as changing it resolves the failure regardless of *I*, while changing *I* alone cannot guarantee correctness.

Leveraging actual causality, we aim to develop a comprehensive framework specifically tailored for analyzing MACGS, grounded in the following causal assumptions. First, the causal graph must be acyclic (see Sec. 2.2), so that the causal relationships among intermediate outputs and final outcomes can be represented as a DAG. Second, edges in the causal graph must respect a temporal ordering that reflects the sequential generation of intermediate outputs in the actual MACGS workflow. Third, interventions on intermediate outputs must eliminate confounding factors and yield realistic erroneous values, so that observed performance changes can be attributed to the interventions




Fig. 1. Illustration of causality.

rather than random perturbations. However, conducting comprehensive causal analysis on MACGS subject to these assumptions presents several challenges:

(I) Challenge of Causal Modeling. It is difficult to systematically decompose complex intermediate outputs of MACGS into meaningful features for causal analysis, since different MACGS may include distinct workflows and organizations.

(II) Challenge of Causal Graph Construction. It is challenging to construct causal graph for MACGS satisfying the requirements of acyclicity and temporal ordering. Iterative refinement processes of MACGS may lead to circular dependencies between features, which violates the acyclicity requirement of causal graphs. For example, MetaGPT's implementation phase might identify missing details, prompting updates to earlier design documents.

(III) Challenge of Realistic Error Simulation. Systematically simulating realistic errors for MACGS poses two key difficulties: 1) the crafted errors should be semantically coherent within MACGS workflow, yet sufficiently different from the original output, and 2) they reflect realistic issues that could occur during agent collaborations [55].

(IV) Challenge of Limited Computation Budget. It is infeasible to exhaustively test all execution states of MACGS, since each MACGS execution incurs significant overhead [68]. Therefore, it is challenging to design an efficient exploration algorithm that can strategically make trade-offs between analysis comprehensiveness and computational overhead.

4 Methodology

This section presents CAM, an automated causality-based analysis framework for analyzing intermediate outputs in MACGS. We begin by establishing a systematic categorization method to model the complex intermediate outputs into different categories with structured semantic fields, and construct causal graph for causal analysis (Sec. 4.1). Subsequently, we simulate realistic errors of intermediate outputs through systematic counterfactual interventions to reflect their robustness issues (Sec. 4.2). Finally, we present our automated algorithm for important feature identification and aggregate feature responsibility to measure the importance of features (Sec. 4.3).

4.1 Causal Modeling of Intermediate Outputs

To address the challenge of causal modeling discussed in Sec. 3.3, we propose a fine-grained categorization method that categorizes the intermediate outputs of MACGS, and systematically constructs causal graph to prevent circular dependencies.

Intermediate Outputs Categorization. We leverage the output organization of MetaGPT [25] as our foundation. As a representative MACGS inspiring numerous subsequent implementations [13, 27, 43, 77], MetaGPT includes structured intermediate outputs across different agents. However, MetaGPT's organization scheme presents two key limitations. First, the intermediate outputs of MetaGPT are organized by standard operating procedure [4] (SOP), which is not inherently reliable and may introduce biases or inconsistencies [69]. Second, this SOP-based organization lacks generalizability—MACGS with distinct architectures cannot be readily adapted to this scheme.

To overcome these limitations, we propose a function-oriented categorization method which reorganizes the intermediate outputs based on their functional roles. Specifically, we first enumerate structured semantic fields from the intermediate output of MetaGPT. Then, we establish four categories reflecting the functional role of these fields in the generation process. Finally, we map every field to its corresponding category, yielding 12 features that span the complete generation chain. As illustrated in Table 3, these categories capture distinct aspects of the code generation process, including problem context (*Specification*), requirement understanding (*Analysis*), architectural decisions (*Design*), and dependency management (*Dependency*). While our categorization is grounded

Table 3. Categorization of Intermediate Outputs.

Category	Description	Example
Specification (3)	This category is responsible for translating the customer’s high-level requirements, i.e., user queries in this context, into concrete and actionable documents like product requirement document (PRD) that guide the subsequent LLM agents in the MACGS.	Programming Language (Program_Lang)
Analysis (4)	This category focuses on systematically breaking down the specified requirements into structured analysis in the MACGS. It consolidates all identified needs, examines the problem’s objectives, and performs comparative analysis to inform the downstream agents with comprehensive understanding of what to build.	Requirement Analysis (Req_Anal)
Design (2)	This category primarily involves making architectural and algorithmic decisions that define the solution structure in the MACGS. It translates the analytical insights from previous agents into concrete implementation steps, specifying module hierarchies, data structures and function signatures.	Implementation Approach (Implement)
Dependency (3)	This category is tasked with managing external resources and inter-agent coordination in the MACGS. It identifies required external packages, specifies files to be produced, and facilitates knowledge exchange between collaborating agents to ensure coherent integration across the system.	Required Packages (Req_Pack)

in MetaGPT, we emphasize that it exhibits substantial generalizability to other MACGS [13, 27, 77], which will be discussed in Sec. 9.1.

Causal Graph Construction. After modeling the intermediate outputs as semantic fields of different categories, we model each field as a feature variable and construct the causal graph. As discussed in Sec. 3.3, the iterative refinement process in MACGS introduces cycles that violate the construction of valid causal graph. To resolve this challenge, we define each feature variable based on its final value in the complete execution trace, excluding mid-stage values generated during iterative refinements. This definition captures the ultimate contribution of each feature to the final output while abstracting away the iterative refinement process, thereby guaranteeing the acyclicity of the causal graph. To enforce temporal ordering, we construct the causal graph by encoding the workflow dependencies in MACGS—each node in the graph corresponds to a feature variable, and edges between nodes capture the information flow across agents: an edge from feature f_i to feature f_j indicates that the output of an upstream agent (containing f_i) influences the computation of a downstream agent (producing f_j). The directionality of these edges thus mirrors the sequential information flow of the MACGS pipeline, in which each downstream agent builds upon the outputs of its upstream predecessors. The constructed causal graph facilitates subsequent causal computation and systematic analysis of feature importance.

4.2 Realistic Error Simulation

As mentioned in Sec. 3.1, robustness of intermediate outputs is crucial for the overall reliability of MACGS. To understand the robustness of intermediate outputs, we simulate realistic errors of different features by systematically conducting counterfactual interventions. However, as discussed in Sec. 3.3, the intervened value of a feature should represent a semantically coherent but incorrect instantiation of that feature—one that an adversary might inject or that might arise from model errors. To achieve this, we first consider several straightforward approaches but find them inadequate: (1) random perturbation may break semantic coherence and produce unrealistic errors; (2) rule-based modification lacks flexibility and struggles to provide systematic interventions; (3) direct text deletion causes information loss without reflecting plausible failures.

Leveraging strong capabilities of LLM in reflecting real-world scenarios [56, 64], we employ an LLM as counterfactual intervention engine. Specifically, given the original features generated by the MACGS, we prompt an LLM to produce modified versions that introduce realistic errors or misunderstandings while maintaining surface-level coherence. To reflect plausible system errors,

we design category-specific instructions in the intervention prompt to generate diverse error types that adapt to different feature categories and problem contexts. To satisfy the requirement of intervention validity discussed in Sec. 3.3, interventions are implemented as direct value substitutions of corresponding features, eliminating potential confounding factors of output changes. Therefore, the changes in final outputs can be attributed to the interventions rather than stochasticity.

This LLM-based intervention approach offers two key advantages: (1) it generates semantically coherent interventions that remain contextually appropriate within the MACGS workflow, and (2) it adapts interventions to problem-specific contexts, producing diverse and realistic errors. To confirm the effectiveness of our LLM-based intervention approach, we manually inspect a random sample of 100 generated counterfactual interventions, confirming that 99 are semantically coherent with the original context while introducing realistic errors.

4.3 Identification of Important Features

After establishing the causal graph and counterfactuals, we now identify important features through automated causal analysis. Specifically, we present an efficient algorithm leveraging influence sets, and aggregate feature responsibility to measure the relative importance of different features.

Problem Formulation. For each coding problem p that MACGS successfully solve in its original execution, our goal is to identify the important features for the success. Formally, let $\mathcal{G}$ denote the causal graph extracted from problem p , with original values of feature nodes leading to correct code. Through systematically conducting counterfactual interventions on feature values, we seek to discover the *important features set*, denoted as $\mathcal{S}_p = \{S_1, S_2, \dots, S_k\}$. Each $S_j \in \mathcal{S}_p$ is an *important feature combination*, defined as a minimal combination of features whose simultaneous intervention leads to failure, leveraging the definition of actual cause introduced in Sec. 2.2. However, as mentioned in Sec. 3.3, it is infeasible to exhaustively test all possible combinations of features, which emphasizes the need for an efficient search algorithm.

Influence Set. Leveraging the unique error propagation properties of MACGS, we introduce the concept of *influence sets*. Unlike traditional configurable software systems [9, 60, 61, 67] where errors propagate deterministically through data dependencies, MACGS exhibit error containment properties due to their agents' ability to self-correct [11]. Consequently, errors on upstream features do not necessarily corrupt all downstream features. Therefore, measuring the actual influence of features facilitates better understanding of their importance. Given a feature combination S , we define the influence set of S as:

$$E(S) = \{f_j \in \mathcal{G} \setminus S : \text{sim}(f_j, f'_j) < \theta\} \quad (2)$$

where f_j and f'_j represent the original feature content and the content observed after intervening on S . $\text{sim}(\cdot, \cdot)$ and θ determine whether a feature is semantically influenced. In essence, $E(S)$ represents the set of features whose semantics are influenced when S is intervened. Influence set provides two key advantages for our analysis. First, computing influence sets reflects the specific characteristics of different problems, facilitating analysis of diverse tasks. Second, features with larger influence sets ($|E(S)|$) potentially exert broader impact on the generation process, as intervention on them affects more downstream components, providing a useful heuristic for our search strategy.

However, computing influence sets requires actually executing interventions and observing the results, which incurs substantial computational overhead. To extend the influence sets for further interventions, we introduce *collective influence sets*. Leveraging the insight that a feature is likely to be influenced by S if it is influenced by subsets of S , we define the collective influence set of S as:

$$\hat{E}(S) = \bigcup_{S' \subseteq S} E(S') \quad (3)$$

Algorithm 1: Computing important feature set for one coding problem

Input: MACGS M , Problem p , Causal Graph $\mathcal{G}$, Max Combination Length L_{max} , Similarity Threshold θ , Maximum MACGS Executions Per Question N , Early Stop Patience k

Output: Important Feature Set S_p

```

1  $S_p \leftarrow \emptyset$  // Initialize  $S_p$ 
2 for feature  $f \in \mathcal{G}$  do
3    $result \leftarrow M_f(p)$  // Execute MACGS with intervention on  $f$ 
4   if  $result = \text{False}$  then
5      $S_p \leftarrow S_p \cup \{f\}$ ;
6   else
7      $E(\{f\}) \leftarrow \text{ComputeInfluenceSet}(\{f\}, \mathcal{G}, \theta)$  // Compute influence set of  $f$ 
8 for  $l = 2$  to  $L_{max}$  do
9    $Comb \leftarrow \text{Combinations}(\mathcal{G}, l)$ 
10  for all ( $S \in Comb$  and  $\exists S' \in S_p \wedge S' \subseteq S$ ) do
11     $Comb \leftarrow Comb \setminus S$  // Prune based on minimality of important feature combination
12   $S \leftarrow \arg \max_{S \in Comb} \hat{E}(S)$ ;
13   $result \leftarrow M_S(p)$  // Execute MACGS with intervention on  $S$ 
14  if  $result = \text{False}$  then
15    if  $\text{CheckMinimal}(S)$  then
16       $S_p \leftarrow S_p \cup \{S\}$  // Find an important feature combination
17  else
18     $E(S) \leftarrow \text{ComputeInfluenceSet}(S, \mathcal{G}, \theta)$ 
19    for all ( $S'$  where  $S' \in \text{Combinations}(E(S), l)$ ) do
20       $Comb \leftarrow Comb \setminus S'$  // Prune based on property of influence set
21    if  $\text{Consecutive\_failures} \geq k$  then
22      continue // Early stop, shift to next length
23  if  $\text{Consumed\_queries} \geq N$  then
24    break // Budget exhausted
25 return  $S_p$ ;

```

where $E(S')$ denotes the influence set of subset S' that has been previously computed. The collective influence set $\hat{E}(S)$ approximates the influence of S without executing MACGS, which facilitates prioritized search strategies and better resource allocation.

Algorithm. Building upon influence set, we propose a novel algorithm which integrates two key components: (1) a greedy selection strategy that prioritize features with substantial influence and (2) pruning techniques that leverage the properties of important features and influence sets to reduce the unpromising search space.

As shown in Alg. 1, our algorithm explores feature combinations in order of increasing length (i.e., number of features in the combination), as shorter combinations can better reflect the importance of individual features and provide valuable guidance for identifying longer combinations. Specifically, we first initialize S_p and then execute MACGS with intervention on each individual feature (lines 1–3). If the intervention on f leads to system failure, we identify it as an important feature and add it to S_p (line 5). If the intervention does not result in failure (i.e., MACGS still generates correct code), we compute its influence sets for subsequent exploration (line 7). Subsequently, we iteratively explore feature combinations of increasing length from 2 to L_{max} (line 8). For each length ℓ , we first prune all combinations that include any previously discovered important feature combinations of length $\leq \ell - 1$ as a subset, since the important feature combinations are minimal by definition (lines 10–11). Then, leveraging the insight that features with substantial influence are more likely to induce

system failures, we apply greedy selection strategy which prioritizes S with the largest collective influence set (line 12). If the intervention on S succeeds in inducing failure and S is confirmed as a minimal combination, we add S to $\mathcal{S}_p$ (lines 14–16). Otherwise, we can prune not only S but also all length- ℓ feature combinations within its influence sets (lines 19–20), since intervening on S already perturbs all features in $E(S)$ transitively. Therefore, any length- ℓ feature combinations from $E(S)$ result in a less comprehensive perturbation, which cannot induce failure if S itself was insufficient. During the search process, if k consecutive interventions fail to discover new important combinations, we shift to the next length to avoid searching unpromising regions (lines 21–22). We continue the above process until we reach predefined resource constraints. This algorithm balances exploration and exploitation: the greedy selection exploits collective influence set to find promising combinations, while the pruning techniques aggressively eliminate unpromising regions to conserve computational resources.

Aggregating Feature Importance. After identifying $\mathcal{S}_p$ for individual problems, we aggregate the results and measure the importance of each feature. Inspired by the concept of responsibility in actual causality (see Sec. 2.2), we measure the contribution of each feature to understand its importance. Leveraging the insight that shorter feature combinations indicate stronger necessity of each feature, we define the *feature responsibility* (FR) of a feature f_i as:

$$FR(f_i) = \frac{1}{|\mathcal{P}|} \sum_{p \in \mathcal{P}} \sum_{S \in \mathcal{S}_p, f_i \in S} \left(\frac{1}{|S|} \right)^2 \quad (4)$$

where $\mathcal{P}$ denotes the problem set and S is an important feature combination for problem p . We normalize the FR by the number of problems to mitigate the influence of dataset size on the results. The squared inverse of combination length assigns substantially higher weights to features appearing in shorter combinations, thereby amplifying the distinction between features with varying importance [59]. By ranking features according to their FR, we obtain a quantitative assessment of their importance to MACGS correctness, providing insights for further analysis.

5 Experimental Setup

MACGS. We instantiate CAM on MetaGPT [25], one of the most representative and widely-adopted open-source MACGS that inspires numerous subsequent implementations [13, 27, 43, 77]. We further demonstrate CAM’s extensibility to Self-Collab [13], PairCoder [77] and MapCoder [27], which are representative MACGS with distinct architectures, in Sec. 9.1.

Datasets. We evaluate CAM on four established benchmarks [13, 27, 40]: HumanEval-ET [7] (HumanEval) and MBPP-ET [2] (MBPP), which offer foundational, well-constructed tasks; CodeContest [38], featuring higher-difficulty competitive programming challenges; and CoderEval [75], a recent benchmark drawn from real-world GitHub repositories for realistic MAS assessment. This selection spans classical to contemporary real-world scenarios, enabling comprehensive analysis across diverse programming environments. For computational efficiency, we use MBPP’s sanitized subset, CodeContest’s test partition, and CoderEval’s standalone-level tasks.

Backend LLMs. Current MACGS implementations employ the same backend LLM for all constituent agents. We employ three mainstream LLMs as backend: GPT-4o-mini [1], DeepSeek-Coder-V2-Instruct [83] (DS-Coder in short) and Qwen-2.5-Coder-14B-Instruct [26] (Qwen in short). All these LLMs are widely adopted and demonstrate strong capabilities in code generation [16, 32, 81]. This selection covers both proprietary and open-source options, reflecting the diverse landscape of LLMs employed in contemporary MACGS implementations.

Parameters. We allocate a maximum of 100 MACGS executions per coding problem ($N = 100$) to balance comprehensiveness with computational feasibility, and cap the feature-combination search

length at five ($L_{max} = 5$), as longer combinations become less interpretable and actionable. Following existing works [15, 54, 63], semantic similarity in Eq. 2 is computed via Sentence-Transformer [53], with threshold $\theta = 0.5$, to capture semantic changes while filtering noise. For each length ℓ , we apply early stopping after 10 consecutive non-failure-inducing interventions ($k = 10$) to avoid exhaustive search of unpromising regions. Sec. 10.1 provides detailed justification for these choices. **Stochasticity Mitigation.** To mitigate the effect of stochasticity, we configure all backend LLMs with a temperature of zero to reduce sampling-induced variance. Moreover, FR aggregates results across all dataset problems (Eq. 4), which mitigates problem-level variance. High CAM-human correlations further confirm the reliability of the results (see Sec. 6).

6 Pilot Study: Validation of Causal Analysis

To validate the reliability and accuracy of CAM, we conduct a pilot study comparing feature importance ranking identified by CAM against manually annotated ranking.

For each setting with different dataset-LLM combination (12 in total), we randomly sample 15% of coding problems. Two experienced annotators (5+ years of programming experience and no conflicts of interest) independently examine each problem, provided with the original problem description, all intermediate outputs from the MACGS pipeline, and the final generated code.

For each problem, annotators identify all features responsible for system failure and rank them by their impact on the final result, establishing a feature importance ranking. Prior to annotation, both annotators are trained on the categorization scheme to ensure sufficient understanding of different feature types. During annotation, they have no access to CAM's results to prevent any potential bias. The two annotators work independently, with no communication during the annotation process, which requires approximately 10 person-hours. The quadratic weighted Cohen's Kappa score [10] between the two annotators is 0.87, indicating substantial agreement [36]. After that, they resolve disagreements through discussions and unresolved cases are adjudicated by a third reviewer. The rankings from all sampled problems are then aggregated to produce the overall feature importance ranking.

For each setting, we evaluate the similarity of feature importance rankings between CAM and manual annotations using Kendall's correlation coefficient [34], which measures ordinal association between two ranked lists and handles tied ranks appropriately. Table 4 presents the quantitative agreement metrics between CAM and manual annotations. The results demonstrate consistently high correlation across settings, with values ranging from 0.76 to 0.91, indicating strong agreement between CAM and human judgments. Therefore, the results produced by CAM are reliable and can be trusted for further causal analysis.

Table 4. Kendall's correlation between CAM and human annotated results.

Dataset	GPT-4o-mini	DS-Coder	Qwen
HumanEval	0.85	0.91	0.79
MBPP	0.88	0.82	0.76
CodeContest	0.82	0.85	0.76
CoderEval	0.82	0.85	0.82

7 Empirical Analysis

In this section, we present a comprehensive analysis of feature importance identified by CAM. Specifically, we investigate how feature importance patterns vary across different settings (RQ1), the influence of backend LLM choices (RQ2), and the impact of dataset characteristics (RQ3).

7.1 RQ1: Overall Feature Importance

We investigate the relative importance of different features by examining their FR distribution. Table 5 presents the appearance count of each feature in top-5 FR rankings across all 12 dataset-LLM combinations. From a categorical perspective, we observe a pronounced hierarchy in feature importance. *Specification* and *Design* features collectively dominate the top positions. Specifically,

Table 5. Feature appearance count in top-5 FR-rankings across all dataset-LLM settings.

Count	Feature	Category	Description
12	Req_Stat	Specification	Problem statement or specific task definition
12	Data_Struct	Design	Core data structures and their detailed definitions
10	Implement	Design	Technical implementation approach and algorithm design
10	Language	Specification	Natural language used for documentation
8	Program_Lang	Specification	Programming language chosen for code implementation
3	Req_Anal	Analysis	Detailed analysis and examination of requirement statements
3	Compet_Anal	Analysis	Comparative analysis of similar problems or questions
2	Req_Pack	Dependency	Required library packages and dependencies
0	Req_Pool	Analysis	Comprehensive requirement pool of requirements
0	Logic_Anal	Analysis	Breakdown of requirement logic flow
0	File_List	Dependency	Complete list of output files needed
0	Share_Know	Dependency	Common shared information across different modules

Table 6. Distribution of important feature combinations across different lengths. Values represent the percentage of combinations (%) at each length to all combinations containing the specified feature.

Feature	L-1	L-2	L-3	L-4	L-5	Feature	L-1	L-2	L-3	L-4	L-5
Req_Stat	97.1	1.6	1.3	0.0	0.0	Data_Struct	53.9	22.5	12.0	10.7	1.0
Logic_Anal	74.3	18.6	3.1	1.8	2.2	Req_Pack	46.2	46.2	1.0	2.9	3.7
Implement	65.7	15.7	4.5	6.4	7.8	Req_Anal	34.0	40.8	15.1	7.2	3.0
Share_Know	65.6	27.3	2.4	2.1	2.6	Compet_Anal	28.2	43.7	12.5	9.2	6.3
Req_Pool	63.6	23.8	3.5	4.0	5.1	Language	25.1	63.2	7.3	3.8	0.5
File_List	55.9	37.3	1.0	3.5	2.2	Program_Lang	21.2	45.6	15.6	15.3	2.4

Req_Stat and Data_Struct both appear in top-5 rankings across all 12 settings, followed by Implement and Language with 10 appearances each. This dominance reflects an intuitive yet empirically validated result: modifications to fundamental problem characteristics and architectural decisions exert the greatest influence on final outcomes. In contrast, *Analysis* features exhibit moderate rankings, with Req_Anal and Compet_Anal appearing in only 3 settings, while *Dependency* features occupy the lowest positions, rarely or never reaching the top-5. This difference suggests that users seeking to reduce computational costs while maintaining performance may strategically prune or simplify *Analysis* and *Dependency* features, which will be discussed in Sec. 8.

We analyze the distribution across different lengths of important feature combinations. Table 6 presents the proportion of different combination lengths in which selected features appear. Overall, the majority of important feature combinations have length ≤ 3 , indicating that in many cases, a few critical features could determine the correctness of MACGS. Intriguingly, we uncover a previously overlooked phenomenon: *context-dependent features*—features whose importance manifests primarily through interactions with other features. For example, among all important feature combinations including Program_Lang, only 21.2% are identified when intervening on Program_Lang itself; in the remaining 78.8% of cases, its importance emerges when simultaneously intervened with other features. This finding reveals a critical insight for improving MACGS: *failures arise not only from individual feature errors, but may also stem from subtle incompatibilities between seemingly correct intermediate features*. The context-dependency of Program_Lang suggests that when combined with other features (e.g., data structure), it can create semantic inconsistencies that lead to system failures—errors that would remain undetected if each feature were validated in isolation. From a practical perspective, this implies that quality assurance for MACGS should move beyond module-level validation to incorporate *cross-feature consistency checks*, ensuring that intermediate outputs are not only individually correct but also mutually compatible. This finding also validates the comprehensiveness of CAM: prior approaches [13, 27, 77] evaluating different modules solely in isolation would fail to identify such synergistic effects.




Fig. 2. Feature responsibility (FR) distribution on CodeEval. (1) – (3) represent top-3 FR-ranked features.

Table 7. Pass@1 comparison of uniform and hybrid backend.

Backend LLM	MBPP	CodeContest	CodeEval	LiveCodeBench
GPT-4o-mini (uniform)	0.6979	0.1178	0.3478	0.8148
DS-Coder (uniform)	0.6812	0.1126	0.3298	0.8052
Hybrid (GPT-4o-mini + DS-Coder)	0.6904	0.1264	0.3645	0.8216

Findings: Our analysis reveals a clear hierarchy in feature importance where *Specification* and *Design* features occupy top rankings. Quality assurance for MACGS should move beyond unit-level validation to incorporate cross-feature consistency checks.

7.2 RQ2: The Influence of Backend LLM Choices

To understand the impact of backend LLMs on feature importance, we compare the FR distributions across different LLMs. Fig. 2 illustrates the FR distributions with different colors representing different feature categories. Due to space constraints, we focus on CodeEval as a representative example; similar patterns are observed across other datasets (results available in our artifact).

As shown in Fig. 2, different backend LLMs exhibit markedly distinct feature importance patterns, reflecting varying capabilities in completing feature-specific subtasks. Specifically, DS-Coder and Qwen exhibit elevated importance on *Design* features compared to GPT-4o-mini. When using DS-Coder or Qwen as the backend, *Data_Struct* surpasses *Req_Stat* in ranking; similar trends are observed across other datasets, where *Data_Struct* achieves FR comparable to *Req_Stat*. In contrast, GPT-4o-mini shows higher reliance on *Specification* features, with *Req_Stat* consistently ranking as most important. These distinct importance patterns reveal that different models possess varying capabilities in generating high-quality outputs for specific features, which may stem from model-specific training objectives and corpus compositions. DS-Coder and Qwen, as specialized code LLMs pretrained on extensive code repositories [22, 26], demonstrate strong capability in producing well-formed, high-quality designs, thereby enhancing their ability to solve complex programming tasks. Conversely, GPT-4o-mini, as part of the general-purpose GPT family, is trained to excel across diverse tasks by prioritizing core semantic features (e.g., requirement statements) from heterogeneous contexts. Therefore, different models prioritize different features for task completion. This suggests that optimal MACGS configurations could leverage DS-Coder and Qwen’s superior design generation capabilities while using GPT-4o-mini’s strength in semantic feature understanding for task specification.

The observation of distinct importance patterns across LLMs motivates hybrid multi-backend architectures for MACGS, where different backend LLMs are strategically assigned to specific sub-tasks based on their relative strengths indicated by FR. To demonstrate the real-world implications of this approach, we conduct an illustrative experiment using GPT-4o-mini and DS-Coder.

Setup. We apply GPT-4o-mini as the base backend, replacing only the design stage with DS-Coder to leverage its superior capability in generating high-quality designs. We derive model-specific

feature importance patterns via FR analysis on the datasets described in Sec. 5 (MBPP-ET sanitized, CodeContest test, CoderEval standalone tasks). The hybrid backend is then evaluated on disjoint held-out partitions to avoid circularity: for MBPP, we use the samples present in the full MBPP-ET dataset but absent in the sanitized version; for CodeContest, we use the validation split; for CoderEval, we use the plib-runnable subset. We further include the easy-difficulty LeetCode problems from LiveCodeBench [28] as an additional benchmark for comprehensive evaluation.

Results. As shown in Table 7, this hybrid configuration outperforms both uniform backends on three out of four datasets, achieving pass@1 improvements of up to 7.3%. Intriguingly, these improvements are most pronounced on datasets requiring complex structural reasoning. Specifically, CodeContest, which demands more sophisticated algorithmic designs, benefit substantially from the hybrid approach (7.3% improvement). In contrast, MBPP—which features simpler problems that rely less on algorithm design—shows marginal degradation under the hybrid configuration. This observation aligns with our feature importance analysis: DS-Coder’s design-centric strength better serves tasks where design quality is significant for success. Furthermore, these results suggest a critical insight for MACGS optimization: *the optimal backend assignment should match model-specific feature importance patterns with stage-specific quality requirements*. Future work could explore more fine-grained assignments, such as dynamically selecting backends based on task characteristics detected at runtime. This opens a promising direction for practical MACGS optimization through systematically exploiting the complementary strengths of diverse LLMs.

Findings: Backend LLMs exhibit varying capabilities in different subtasks. The optimal backend assignment should match model-specific importance patterns with stage-specific requirements.

7.3 RQ3: The Impact of Dataset Characteristics

We analyze how problem difficulty and domain focus influence feature importance patterns, revealing systematic variations that facilitates context-aware MACGS optimization strategies.

To comprehensively analyze importance patterns, we compute the standard deviation [62] (STD) of FR distributions. Since FR results are normalized by dataset size,

STD facilitates comparison of the uniformity of FR across different settings. As illustrated in Table 8, we observe that challenging datasets like CodeContest (average pass@1 10.2%) exhibit relatively uniform FR distributions across all LLMs (STD ranging from 0.1137 to 0.1773), indicating that multiple features become comparably critical. In contrast, simpler datasets like HumanEval and MBPP (average pass@1 $\geq 60\%$) display higher STD, indicating distributions with dominant high-importance features. This contrast may stem from differences in how MACGS processes problems of varying complexity. For simpler problems, MACGS relies on a small subset of critical features to derive solutions, resulting in substantially greater importance. Conversely, complex problems necessitate the synthesis of information across multiple features. Consequently, less important features provide substantial information that contributes to coding process, yielding more uniform distributions. These findings suggest that optimal resource allocation strategies should adapt to target problem difficulty. When deploying MACGS for simple tasks (e.g., basic algorithm design), users can concentrate optimization efforts on top-ranked features. For complex tasks (e.g., competitive programming), a more balanced approach that maintains quality across all feature categories becomes necessary.

We investigate how domain characteristics influence importance patterns by comparing FR distributions across different datasets. Fig. 3 presents results on DS-Coder, while similar patterns

Table 8. Standard deviation of FR distributions (lower is more uniform).

Dataset	GPT-4o-mini	DS-Coder	Qwen
HumanEval	0.2139	0.1787	0.2121
MBPP	0.2243	0.1946	0.2226
CodeContest	0.1773	0.1137	0.1323
CoderEval	0.2216	0.1277	0.1722




Fig. 3. Feature responsibility (FR) distributions across different datasets on DS-Coder.

emerge with other backend LLMs. Notably, we observe that task characteristics exert remarkable influence on importance patterns. In CodeEval, a dataset reflecting realistic engineering scenarios with complex module interactions and external dependencies, the importance of *Dependency* features is remarkably elevated. Specifically, *Dependency* features surpass *Analysis* features in FR ranking, demonstrating similar importance to *Specification* features. This shift aligns with the nature of engineering-oriented problems, where diverse libraries are inherently integrated, making errors in dependency resolution easily manifest as system failures. These findings suggest that practitioners applying MACGS to domain-specific tasks should prioritize optimization on corresponding features, potentially incorporating adaptive resource allocation strategies for optimized deployment.

Findings: Problem difficulty and domain characteristics modulate feature importance patterns, suggesting adaptive resource allocation strategies for optimal MACGS deployment.

8 Downstream Applications

Leveraging the FR rankings in Sec. 7, we further demonstrate two practical applications to address real-world challenges in MACGS deployment.

8.1 Causality-Guided Failure Repair

When MACGS produces incorrect code, developers face the challenge of diagnosing and correcting failures efficiently [42]. Exhaustive examination of all intermediate outputs incurs substantial human effort [48]. Our analysis suggests a causality-guided strategy: with high-FR ranked features demonstrating elevated importance to MACGS success, targeted optimization of these features may effectively resolve these failures.

Setup. For each of the 12 dataset-LLM settings, we split the questions by MACGS’s initial correctness: questions MACGS originally fails serve as repair targets, while questions it solves correctly are used to derive FR rankings. From each setting’s failure set, we randomly sample 18 cases and optimize the top- n FR-ranked features through a three-step process: (1) understanding the problem requirements, (2) assessing each feature’s semantic clarity and completeness, and (3) augmenting each feature with detailed interpretations and explanations. The refined features are reintegrated into the execution pipeline, and the resulting code is re-evaluated on the original test suite. To assess whether CAM-derived feature importance is actionable, we compare causality-guided repair against three baselines: *random selection* (uniformly sampling n features), *temporal-first* (prioritizing the earliest n generated features), and *length-based* (selecting the n features with the longest outputs). All selected features receive equal refinement effort across strategies. Given the substantial manual effort required, we set $n = 3$ to balance effectiveness against cost.

Table 9. Effectiveness of target enhancement of feature quality ($n = 3$).

Method	Pass Rate
Random-select	27.8%
Temporal-first	36.6%
Length-based	39.4%
Causality-guided	73.6%

Table 10. Impact of feature pruning on model performance across benchmarks.

Backend LLM	# Pruned	HumanEval		MBPP		CoderEval		CodeContest	
		$\Delta\text{Pass@1}$	ΔTokens	$\Delta\text{Pass@1}$	ΔTokens	$\Delta\text{Pass@1}$	ΔTokens	$\Delta\text{Pass@1}$	ΔTokens
GPT-4o-mini	2	+5.08%	-6.38%	+3.96%	-15.79%	-1.53%	-8.52%	+1.82%	-9.25%
	4	+3.63%	-25.42%	-0.67%	-21.05%	+2.06%	-31.85%	-9.09%	-22.15%
	6	+2.54%	-45.07%	-3.08%	-44.01%	-1.46%	-48.24%	-18.18%	-30.63%
	8	-4.24%	-57.98%	-2.64%	-60.91%	-3.87%	-62.52%	-36.36%	-46.84%
DS-Coder	2	-2.08%	-2.36%	+2.37%	-4.18%	+0.85%	-15.68%	-7.45%	-7.50%
	4	-2.69%	-6.12%	+3.32%	-7.65%	+5.88%	-33.58%	-17.27%	-14.89%
	6	-5.21%	-39.14%	-0.08%	-43.64%	-11.76%	-47.03%	-26.36%	-28.73%
	8	-6.24%	-56.60%	-12.80%	-55.82%	-17.65%	-56.47%	-45.45%	-36.11%
Qwen	2	+0.28%	-2.88%	+1.02%	-2.31%	+5.36%	-5.53%	-9.09%	-5.87%
	4	-3.08%	-14.81%	-0.51%	-24.81%	+7.69%	-27.32%	-18.18%	-19.43%
	6	+2.02%	-37.15%	+2.54%	-43.84%	+3.77%	-35.40%	-27.27%	-32.56%
	8	-7.56%	-60.81%	-5.08%	-66.80%	-23.08%	-57.60%	-45.45%	-39.72%

Results. As shown in Table 9, causality-guided repair achieves an overall successful repair rate of 73.6%, significantly outperforming other baseline approaches, which only resolve 27.8%–39.4% of the failures, indicating that quality of high FR-ranked features are crucial for system correctness. These results validate CAM’s ability to identify genuinely important features. Moreover, the concentration of feature importance suggests practical failure repair: when MACGS produces incorrect code, developers should prioritize optimization of top FR-ranked features with elevated importance for system correctness, rather than exhaustively reviewing all intermediate outputs. In practice, developers can implement a hybrid workflow that combines automated validation with selective manual inspection: allocate manual review capacity exclusively to high-FR features (e.g., top-3 ranked) while accepting automated validation for lower-ranked components.

8.2 Causality-guided Feature Pruning

The computational overhead of MACGS constitutes a significant barrier to widespread adoption [25, 50]. Recent studies report that complex problems can consume hundreds of thousands of tokens across multiple agents [68]. Our causal analysis provides a new perspective for efficient MACGS deployment: if certain features exhibit low importance on system correctness, pruning these features should reduce computational costs without proportionate performance degradation.

Setup. We systematically remove low-FR features and measure the impact on both code generation performance and computational efficiency. For each configuration, we disable the bottom- n FR-ranked features from MetaGPT execution, from $n = 2$ to $n = 8$ in increments of 2, covering a spectrum from conservative pruning to aggressive pruning. For each pruning level, we compute the relative performance change $\Delta\text{Pass@1} = \frac{\text{Pass@1}_{\text{pruned}} - \text{Pass@1}_{\text{original}}}{\text{Pass@1}_{\text{original}}}$ and intermediate output token reduction $\Delta\text{Tokens} = \frac{\text{Tokens}_{\text{pruned}} - \text{Tokens}_{\text{original}}}{\text{Tokens}_{\text{original}}}$.

Results. Table 10 presents the results of performance-efficiency trade-offs. We observe that conservative pruning (removing 2–4 features) not only reduces token consumption substantially but may also improve pass@1 performance. For instance, with GPT-4o-mini, pruning two features yields performance improvements across many datasets while achieving token reductions of 6.38%–15.79%, and even more aggressive 4-feature pruning maintains positive gains on HumanEval (+3.63%) and CoderEval (+2.06%). This counterintuitive improvement may be attributed to low-FR features introducing redundant or inconsistent information that distracts MACGS from critical decision points—removing these features may enhance focus on essential reasoning pathways. However, aggressive pruning (removing 8 features) induces consistent performance degradation, confirming that high-FR features exert substantial importance. These results establish quantitative guidelines for MACGS optimization: for production deployments prioritizing cost efficiency, pruning 2–4

low-FR features achieves moderate token reduction of up to 33.6% with negligible or even positive performance impact. Moreover, CAM supports dynamic adaptive pruning strategies based on different configurations where future MACGS implementations could selectively enable or disable features based on task complexity.

9 Generalizability and Extensibility

9.1 Generalizability to other MACGS

To demonstrate the generalizability of CAM, we instantiate CAM on another prominent MACGS, Self-Collaboration Code Generation [13] (Self-Collab in short). We further discuss the extensibility of our categorization method to PairCoder [77] and MapCoder [27].

First, as illustrated in Table 3, we map the intermediate outputs of Self-collab to different categories and features. Specifically, the *subproblems* of Self-Collab involves problem characteristics and requirement decomposition, which can be decomposed into *Specification* and *Analysis* categories. Similarly, the *high-level steps* maps to *Design* categories and the output for agent coordination maps to *Dependency* categories. Then, we execute CAM on Self-Collab. Fig. 4 demonstrates the FR distributions on representative setting.

These results exhibit patterns consistent with MetaGPT while revealing architecture-specific characteristics. *Specification* features maintain dominance, with *Req_Stat* achieving the highest FR. Certain features, such as *Logic_Anal*, demonstrate elevated importance, reflecting Self-Collab's reliance on coding logic. Due to space limits, full results for Self-Collab are available in our artifact.

Beyond Self-Collab, our categorization method readily extends to other MACGS. For example, for PairCoder [77], output for *promising plans proposal* can be mapped into *Specification* and *Analysis*, while *optimal plan selection* and *plan switch* can be decomposed into *Design* and *Dependency*. Similarly, MapCoder [27]'s diverse intermediate outputs—problem analysis, exemplar retrieval, and refinement plans—categorize into *Specification* (problem context), *Analysis* (exemplar reasoning), *Design* (coding plan), and *Dependency* (inter-agent knowledge sharing).

These analyses validate CAM's generalizability and adaptability. Although different MACGS implementations employ different workflow structures, their intermediate outputs can be systematically aligned to our categories through appropriate semantic alignment.




Fig. 4. Results of Self-Collab.

9.2 Generalizability to Reasoning Models

While backend models selected in Sec. 7 represent the mainstream configurations employed in predominant MACGS [13, 25, 27, 77], reasoning-capable models represent an increasingly prominent paradigm in LLM development [12]. However, due to their substantially higher overhead [20], reasoning models are not yet the mainstream configuration in current MACGS deployments. Nevertheless, to empirically validate the generalizability of CAM to reasoning-capable backends, we instantiate CAM with DeepSeek-R1 [12], a representative reasoning model, as the backend LLM and evaluate across all four benchmarks. The FR distributions are presented in Fig. 5.

Our categorization method remains directly applicable to the intermediate outputs produced by reasoning models. Although reasoning models generate extended internal reasoning traces, the final outputs generated by them still maintain the same structured semantic fields as those generated by non-reasoning models [76]. This is consistent with our design choice of defining each feature variable based on its final value in the execution trace (see Sec. 4.1), which abstracts away reasoning traces and focuses on inter-agent information flow. Consequently, the causal graph




Fig. 5. Feature responsibility (FR) distributions across different datasets on DeepSeek-R1.

construction, counterfactual intervention, and important feature identification procedures of CAM require no modification when applied to reasoning model backends.

As illustrated in Fig. 5, the FR distributions exhibit patterns broadly consistent with other backend LLMs. Specifically, *Specification* and *Design* features continue to dominate the top FR rankings across all datasets, reaffirming the finding that fundamental problem characteristics and architectural decisions exert the greatest influence on system correctness. Additionally, the task-modulated importance patterns identified in Sec. 7.3 are consistent, where CodeContest displays the most uniform FR distribution (STD=0.1377). This consistency suggests that the task-modulated patterns are not specific to non-reasoning models but also hold for reasoning models, reflecting that complex problems necessitate the synthesis of information across multiple features for correct solutions.

These results demonstrate that CAM generalizes effectively to reasoning-capable model backends. The key empirical findings from Sec. 7 remain consistent with mainstream reasoning models, establishing CAM as a broadly applicable analysis tool for MACGS.

Table 11. Average contribution (%) for each important feature combination length to FR.

Length	1	2	3	4	5
Contribution	81.6	15.9	1.5	0.7	0.2

10 Discussion

10.1 Configuration of CAM

Maximum Combination Length ($L_{max} = 5$). Setting $L_{max} = 5$ balances comprehensiveness with interpretability. Table 11 presents the average contribution for each combination length to final FR. While individual features (length 1) contribute the majority of FR, combinations of length ≥ 2 contribute a significant part (18.4%), indicating that limiting analysis to the importance of individual features yields incomplete understanding. However, our experiment results show that FR values stabilize significantly and transition from length four to five produces minimal FR changes (only 0.2% contribution), suggesting convergence. This configuration facilitates sound identification of important features without sacrificing interpretability through excessively complex combinations.

Early Stopping Patience ($k = 10$). The parameter $k = 10$ prevents premature termination while avoiding excessive exploration of unpromising search spaces. Table 12 presents the average number of important feature combinations identified across different k values. We observe that $k < 10$ leads to insufficient exploration, missing 21.4% combinations. For $k \in [10, 20]$, the number of identified combinations remains stable. Values $k > 20$ lead to excessive exploration of unpromising regions. Thus, $k = 10$ represents an effective balance.

Table 12. Average identified combinations across different k .

k	5	10	15	20	25
# Comb	659.2	839.4	822.7	825.2	715.0

10.2 Threats to Validity

We identify and address three threats to validity. First, regarding representativeness of experimental setup, we conduct comprehensive evaluation across three mainstream LLMs from different families

and four diverse benchmarks, and demonstrate CAM's extensibility to other MACGS in Sec. 9.1. Second, to mitigate subjectivity in manual feature importance analysis, two software developers independently annotated feature rankings for sampled problems, with all disagreements resolved. Third, to address stochasticity, we set all backend LLMs to temperature = 0 and validate stability across three runs, confirming consistent feature importance and rankings.

10.3 Cost Analysis

The API consumption of CAM is around 25M tokens per dataset-LLM configuration. The time cost is around 8-12 hours for each configuration. We note that this overhead is comparable with other analysis works [5, 66] and reasonable for real-world deployment. As a one-time analysis tool, CAM follows an offline analysis, online deployment paradigm [29, 44]: the FR rankings are derived offline, then fixed for all subsequent deployments with zero additional inference-time overhead.

11 Related Work

MACGS. MACGS have emerged as a promising paradigm for automated software development [6, 13, 25, 77, 78]. MetaGPT [25] is one of the most widely-adopted systems which simulates a software development team with specialized agents in distinct roles. Self-collab [13] employs three agents for planning, coding, and testing. Other systems include PairCoder [77] with clustering-based plan selection, MapCoder [27] with automated plan exploration, and CAMEL [37] with role-playing conversations. Despite their success, MACGS generate substantial intermediate outputs whose importance to the system correctness remains opaque, hindering further optimization of MACGS.

Causality Analysis in Software Engineering. Causality analysis has been extensively applied in software engineering to identify causal relationships for debugging [18] and root cause analysis [33]. Recently, formal notions of actual causality [23] have enabled precise characterization of necessary and sufficient conditions for program behaviors [3, 14, 45, 72]. However, these techniques have not been systematically applied to MACGS. We bridge this gap by presenting CAM, a causality-based analysis framework specifically designed for MACGS, which systematically quantifies the contribution of different intermediate features for system correctness. Unlike predictive feature attribution methods like Shapley values [58], which primarily capture marginal contributions based on observational data without strict causal structures, CAM performs causal analysis on an explicit causal graph and captures the causal relation between intermediate outputs and system correctness through comprehensive feature categorization and realistic error simulation.

12 Conclusion

In this paper, we present CAM, the first causality-based framework for MACGS, which identifies how intermediate outputs causally influence final code correctness. Through comprehensive experiments, we reveal critical insights that provide actionable guidance for MACGS optimization and deployment. We demonstrate CAM's practical utility through two applications, establishing causality analysis as a powerful approach for understanding and improving MACGS.

Acknowledgments

We sincerely thank the anonymous reviewers for their valuable feedback. This work was supported in part by the National Natural Science Foundation of China (Grant No. 92582201), the Hong Kong SAR Research Grants Council General Research Fund (Ref. No. 16206524), the Hong Kong SAR Research Grants Council Theme-based Research Scheme (Ref. No. T41-517/25-N), a grant from the Research Grants Council of the Hong Kong Special Administrative Region, China, through HKUST (No. C6004-25G), and an ITF grant under the contract No. ITS/161/24FP.

Data Availability

We release our code and data to facilitate future research on MACGS at <https://github.com/zongyiLyu/CAM>.

References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774* (2023).
- [2] Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. 2021. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732* (2021).
- [3] Christel Baier, Clemens Dubslaff, Florian Funke, Simon Jantsch, Rupak Majumdar, Jakob Piribauer, and Robin Ziemek. 2021. From Verification to Causality-Based Explanations. In *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021*. 1–1.
- [4] R Meredith Belbin and Victoria Brown. 2022. *Team roles at work*. Routledge.
- [5] Islem Bouzenia and Michael Pradel. 2025. Understanding software engineering agents: A study of thought-action-result trajectories. *arXiv preprint arXiv:2506.18824* (2025).
- [6] Guangyao Chen, Siwei Dong, Yu Shu, Ge Zhang, Jaward Sesay, Börje Karlsson, Jie Fu, and Yemin Shi. 2024. AutoAgents: a framework for automatic agent generation. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*. 22–30.
- [7] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374* (2021).
- [8] Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chi-Min Chan, Heyang Yu, Yaxi Lu, Yi-Hsin Hung, Chen Qian, et al. 2023. Agentverse: Facilitating multi-agent collaboration and exploring emergent behaviors. In *The Twelfth International Conference on Learning Representations*.
- [9] Philipp Chrszon, Clemens Dubslaff, Sascha Klüppelholz, and Christel Baier. 2018. ProFeat: feature-oriented engineering for family-based probabilistic model checking. *Formal Aspects of Computing* 30, 1 (2018), 45–75.
- [10] Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement* 20, 1 (1960), 37–46.
- [11] Yanbo Dai, Zhenlan Ji, Zongjie Li, Kuan Li, and Shuai Wang. 2025. Disabling Self-Correction in Retrieval-Augmented Generation via Stealthy Retriever Poisoning. *arXiv preprint arXiv:2508.20083* (2025).
- [12] DeepSeek-AI. 2025. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning. *CoRR abs/2501.12948* (2025). doi:10.48550/ARXIV.2501.12948 arXiv:2501.12948
- [13] Yihong Dong, Xue Jiang, Zhi Jin, and Ge Li. 2024. Self-collaboration code generation via chatgpt. *ACM Transactions on Software Engineering and Methodology* 33, 7 (2024), 1–38.
- [14] Clemens Dubslaff, Kallistos Weis, Christel Baier, and Sven Apel. 2022. Causality in Configurable Software Systems. In *Proceedings of the 44th International Conference on Software Engineering (ICSE)*. ACM, 325–337. doi:10.1145/3510003.3510200
- [15] Ábel Elekes, Martin Schäler, and Klemens Böhm. 2017. On the various semantics of similarity in word embedding models. In *2017 ACM/IEEE Joint Conference on Digital Libraries (JCDL)*. IEEE, 1–10.
- [16] Sarah Fakhoury, Aaditya Naik, Georgios Sakkas, Saikat Chakraborty, and Shuvendu K Lahiri. 2024. Llm-based test-driven interactive code generation: User study and empirical evaluation. *IEEE Transactions on Software Engineering* (2024).
- [17] Sarah Fakhoury, Aaditya Naik, Georgios Sakkas, Saikat Chakraborty, Madan Musuvathi, and Shuvendu Lahiri. 2024. Exploring the effectiveness of llm based test-driven interactive code generation: User study and empirical evaluation. In *Proceedings of the 2024 IEEE/ACM 46th International Conference on Software Engineering: Companion Proceedings*. 390–391.
- [18] Anna Fariha, Suman Nath, and Alexandra Meliou. 2020. Causality-guided adaptive interventional debugging. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*. 431–446.
- [19] Zhangyin Feng, Daya Guo, Duyu Tang, Nan Duan, Xiaocheng Feng, Ming Gong, Linjun Shou, Bing Qin, Ting Liu, Daxin Jiang, et al. 2020. CodeBERT: A Pre-Trained Model for Programming and Natural Languages. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. 1536–1547.
- [20] Jiaxuan Gao, Shu Yan, Qixin Tan, Shusheng Xu, Wei Fu, Zhiyu Mei, Kaifeng Lyu, Yi Wu, et al. 2026. How far are we from optimal reasoning efficiency? *Advances in Neural Information Processing Systems* 38 (2026), 10916–10946.

- [21] Daya Guo, Shuai Lu, Nan Duan, Yanlin Wang, Ming Zhou, and Jian Yin. 2022. UniXcoder: Unified Cross-Modal Pre-training for Code Representation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 7212–7225.
- [22] Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao Bi, Yu Wu, YK Li, et al. 2024. DeepSeek-Coder: When the Large Language Model Meets Programming—The Rise of Code Intelligence. *arXiv preprint arXiv:2401.14196* (2024).
- [23] Joseph Y. Halpern. 2016. *Actual Causality*. MIT Press.
- [24] Junda He, Christoph Treude, and David Lo. 2024. LLM-Based Multi-Agent Systems for Software Engineering: Literature Review, Vision and the Road Ahead. *ACM Transactions on Software Engineering and Methodology* (2024).
- [25] Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiwu Zheng, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, Chenyu Ran, Lingfeng Xiao, Chenglin Wu, and Jürgen Schmidhuber. 2024. MetaGPT: Meta Programming for A Multi-Agent Collaborative Framework. In *The Twelfth International Conference on Learning Representations*.
- [26] Binyuan Hui, Jian Yang, Zeyu Cui, Jiaxi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Keming Lu, et al. 2024. Qwen2.5-coder technical report. *arXiv preprint arXiv:2409.12186* (2024).
- [27] Md Ashrafur Islam, Mohammed Eunus Ali, and Md Rizwan Parvez. 2024. MapCoder: Multi-Agent Code Generation for Competitive Problem Solving. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 4912–4944.
- [28] Naman Jain, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. 2025. Livecodebench: Holistic and contamination free evaluation of large language models for code. In *International Conference on Learning Representations*, Vol. 2025. 58791–58831.
- [29] Zhenlan Ji, Pingchuan Ma, Zongjie Li, Zhaoyu Wang, and Shuai Wang. 2025. Causality-Aided Evaluation and Explanation of Large Language Model-Based Code Generation. *Proceedings of the ACM on Software Engineering* 2, ISSTA (2025), 1374–1397.
- [30] Juyong Jiang, Fan Wang, Jiasi Shen, Sungju Kim, and Sunghun Kim. 2024. A survey on large language models for code generation. *arXiv preprint arXiv:2406.00515* (2024).
- [31] Zhijing Jin, Yuen Chen, Felix Leeb, Luigi Gresele, Ojasv Kamal, Zhiheng Lyu, Kevin Blin, Fernando Gonzalez Adauto, Max Kleiman-Weiner, Mrinmaya Sachan, et al. 2023. Cladder: Assessing causal reasoning in language models. *Advances in Neural Information Processing Systems* 36 (2023), 31038–31065.
- [32] Sathvik Joel, Jie Wu, and Fatemeh Fard. 2024. A survey on llm-based code generation for low-resource and domain-specific programming languages. *ACM Transactions on Software Engineering and Methodology* (2024).
- [33] Brittany Johnson, Yuriy Brun, and Alexandra Meliou. 2020. Causal testing: Understanding defects’ root causes. In *Proceedings of the ACM/IEEE 42nd international conference on software engineering*. 87–99.
- [34] Maurice G Kendall. 1938. A new measure of rank correlation. *Biometrika* 30, 1-2 (1938), 81–93.
- [35] Emre Kiciman, Robert Ness, Amit Sharma, and Chenhao Tan. 2023. Causal reasoning and large language models: Opening a new frontier for causality. *Transactions on Machine Learning Research* (2023).
- [36] J Richard Landis and Gary G Koch. 1977. The measurement of observer agreement for categorical data. *biometrics* (1977), 159–174.
- [37] Guohao Li, Hasan Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2023. Camel: Communicative agents for "mind" exploration of large language model society. *Advances in Neural Information Processing Systems* 36 (2023), 51991–52008.
- [38] Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, et al. 2022. Competition-level code generation with alphacode. *Science* 378, 6624 (2022), 1092–1097.
- [39] Fulin Lin, Shaowen Chen, Ruishan Fang, Hongwei Wang, and Tao Lin. 2025. Stop wasting your tokens: Towards efficient runtime multi-agent systems. *arXiv preprint arXiv:2510.26585* (2025).
- [40] Feng Lin, Dong Jae Kim, and Tse-Hsun Chen. 2025. SOEN-101: Code Generation by Emulating Software Process Models Using Large Language Model Agents. In *2025 IEEE/ACM 47th International Conference on Software Engineering (ICSE)*. IEEE, 1527–1539.
- [41] Junwei Liu, Kaixin Wang, Yixuan Chen, Xin Peng, Zhenpeng Chen, Lingming Zhang, and Yiling Lou. 2024. Large language model-based agents for software engineering: A survey. *arXiv preprint arXiv:2409.02977* (2024).
- [42] Zongyi Lyu, Songqiang Chen, Zhenlan Ji, Liwen Wang, Shuai Wang, Daoyuan Wu, Wenxuan Wang, and Shing-Chi Cheung. 2025. Testing and Enhancing Multi-Agent Systems for Robust Code Generation. *arXiv preprint arXiv:2510.10460* (2025).
- [43] M. A. Islam, M. E. Ali, and M. R. Parvez. 2025. CodeSim: Multi-Agent Code Generation and Problem Solving through Simulation-Driven Planning and Debugging. In *Findings of the Association for Computational Linguistics: NAACL 2025*. 5113–5139.

- [44] Xinyin Ma, Gongfan Fang, and Xinchao Wang. 2023. Llm-pruner: On the structural pruning of large language models. *Advances in neural information processing systems* 36 (2023), 21702–21720.
- [45] Herong Meng and Dalal Alrajeh. 2025. Causal Detection in Linear Temporal Counterexamples. In *Proceedings of the 33rd ACM International Conference on the Foundations of Software Engineering*. 1418–1423.
- [46] Ari Morcos, Maithra Raghu, and Samy Bengio. 2018. Insights on representational similarity in neural networks with canonical correlation. *Advances in neural information processing systems* 31 (2018).
- [47] Changan Niu, Chuanyi Li, Vincent Ng, Jidong Ge, Liguo Huang, and Bin Luo. 2022. Spt-code: Sequence-to-sequence pre-training for learning source code representations. In *Proceedings of the 44th international conference on software engineering*. 2006–2018.
- [48] Melissa Z Pan, Mert Cemri, Lakshya A Agrawal, Shuyi Yang, Bhavya Chopra, Rishabh Tiwari, Kurt Keutzer, Aditya Parameswaran, Kannan Ramchandran, Dan Klein, et al. 2025. Why do multiagent systems fail?. In *ICLR 2025 Workshop on Building Trust in Language Models and Applications*.
- [49] Judea Pearl. 2009. *Causality*. Cambridge university press.
- [50] Chen Qian, Xin Cong, Cheng Yang, Weize Chen, Yusheng Su, Juyuan Xu, Zhiyuan Liu, and Maosong Sun. 2023. Communicative agents for software development. *arXiv preprint arXiv:2307.07924* 6, 3 (2023).
- [51] Chen Qian, Wei Liu, Hongzhang Liu, Nuo Chen, Yufan Dang, Jiahao Li, Cheng Yang, Weize Chen, Yusheng Su, Xin Cong, et al. 2024. Chatdev: Communicative agents for software development. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 15174–15186.
- [52] Maithra Raghu, Justin Gilmer, Jason Yosinski, and Jascha Sohl-Dickstein. 2017. Svcca: Singular vector canonical correlation analysis for deep learning dynamics and interpretability. *Advances in neural information processing systems* 30 (2017).
- [53] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 3982–3992.
- [54] Navid Rekasaz, Mihai Lupu, and Allan Hanbury. 2017. Exploration of a threshold for similarity based on uncertainty in word embedding. In *European conference on information retrieval*. Springer, 396–409.
- [55] Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. Beyond Accuracy: Behavioral Testing of NLP Models with CheckList. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 4902–4912.
- [56] Cedric Richter and Heike Wehrheim. 2022. Learning realistic mutations: Bug creation for neural bug detectors. In *2022 IEEE Conference on Software Testing, Verification and Validation (ICST)*. IEEE Computer Society, 162–173.
- [57] Mohamad Salim, Jasmine Latendresse, SayedHassan Khatoonabadi, and Emad Shihab. 2026. Tokenomics: Quantifying Where Tokens Are Used in Agentic Software Engineering. *arXiv preprint arXiv:2601.14470* (2026).
- [58] Lloyd S. Shapley. 1953. A Value for n-Person Games. In *Contributions to the Theory of Games II*. Princeton University Press, Princeton, NJ, 307–317.
- [59] Donald Shepard. 1968. A two-dimensional interpolation function for irregularly-spaced data. In *Proceedings of the 1968 23rd ACM National Conference*. ACM, 517–524.
- [60] Norbert Siegmund, Sergiy S Kolesnikov, Christian Kästner, Sven Apel, Don Batory, Marko Rosenmüller, and Gunter Saake. 2012. Predicting performance via automated feature-interaction detection. In *2012 34th International Conference on Software Engineering (ICSE)*. IEEE, 167–177.
- [61] Norbert Siegmund, Marko Rosenmüller, Christian Kästner, Paolo G Giarrusso, Sven Apel, and Sergiy S Kolesnikov. 2013. Scalable prediction of non-functional properties in software product lines: Footprint and memory consumption. *Information and Software Technology* 55, 3 (2013), 491–507.
- [62] Carolin Strobl, Anne-Laure Boulesteix, Thomas Kneib, Thomas Augustin, and Achim Zeileis. 2008. Conditional variable importance for random forests. *BMC bioinformatics* 9, 1 (2008), 1–11.
- [63] Nandan Thakur, Nils Reimers, Johannes Daxenberger, and Iryna Gurevych. 2021. Augmented SBERT: Data augmentation method for improving bi-encoders for pairwise sentence scoring tasks. In *Proceedings of the 2021 conference of the North American chapter of the association for computational linguistics: Human language technologies*. 296–310.
- [64] Frank Tip, Jonathan Bell, and Max Schäfer. 2025. Llmorpheus: Mutation testing using large language models. *IEEE Transactions on Software Engineering* (2025).
- [65] Gladys Tyen, Hassan Mansoor, Victor Cărbune, Yuanzhu Peter Chen, and Tony Mak. 2024. LLMs cannot find reasoning errors, but can correct them given the error location. In *Findings of the Association for Computational Linguistics: ACL 2024*. 13894–13908.
- [66] Saad Ullah, Mingji Han, Saurabh Pujar, Hammond Pearce, Ayse Coskun, and Gianluca Stringhini. 2024. LLMs cannot reliably identify and reason about security vulnerabilities (yet?): A comprehensive evaluation, framework, and benchmarks. In *2024 IEEE symposium on security and privacy (SP)*. IEEE, 862–880.

- [67] Alexander Von Rhein, Alexander Grebhorn, Sven Apel, Norbert Siegmund, Dirk Beyer, and Thorsten Berger. 2015. Presence-condition simplification in highly configurable systems. In *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering*, Vol. 1. IEEE, 178–188.
- [68] Ningning Wang, Xavier Hu, Pai Liu, He Zhu, Yue Hou, Heyuan Huang, Shengyu Zhang, Jian Yang, Jiaheng Liu, Ge Zhang, et al. 2025. Efficient agents: Building effective agents while reducing cost. *arXiv preprint arXiv:2508.02694* (2025).
- [69] Qian Wang, Tianyu Wang, Zhenheng Tang, Qinbin Li, Nuo Chen, Jingsheng Liang, and Bingsheng He. 2025. MegaAgent: A large-scale autonomous LLM-based multi-agent system without predefined SOPs. In *Findings of the Association for Computational Linguistics: ACL 2025*. 4998–5036.
- [70] Yue Wang, Qiuzhi Liu, Jiahao Xu, Tian Liang, Xingyu Chen, Zhiwei He, Linfeng Song, Dian Yu, Juntao Li, Zhuosheng Zhang, et al. 2025. Thoughts are all over the place: On the underthinking of o1-like llms. *arXiv preprint arXiv:2501.18585* (2025).
- [71] Yue Wang, Weishi Wang, Shafiq Joty, and Steven C. H. Hoi. 2021. CodeT5: Identifier-aware Unified Pre-trained Encoder-Decoder Models for Code Understanding and Generation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 8696–8708.
- [72] Kallistos Weis, Leopoldo Motta Teixeira, Clemens Dubslaff, and Sven Apel. 2024. Blackbox Observability of Features and Feature Interactions. In *Proceedings of the 39th IEEE/ACM International Conference on Automated Software Engineering*. 1120–1132.
- [73] Zhaofeng Wu, Linlu Qiu, Alexis Ross, Ekin Akyürek, Boyuan Chen, Bailin Wang, Najoung Kim, Jacob Andreas, and Yoon Kim. 2024. Reasoning or reciting? exploring the capabilities and limitations of language models through counterfactual tasks. Association for Computational Linguistics.
- [74] John Yang, Carlos E Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan, and Ofir Press. 2024. Swe-agent: Agent-computer interfaces enable automated software engineering. *Advances in Neural Information Processing Systems* 37 (2024), 50528–50652.
- [75] Hao Yu, Bo Shen, Dezhi Ran, Jiaxin Zhang, Qi Zhang, Yuchi Ma, Guangtai Liang, Ying Li, Qianxiang Wang, and Tao Xie. 2024. Codereval: A benchmark of pragmatic code generation with generative pre-trained models. In *Proceedings of the 46th IEEE/ACM International Conference on Software Engineering*. 1–12.
- [76] Han Yuan, Yue Zhao, Li Zhang, Wuqiong Luo, and Zheng Ma. 2026. Quantifying the Impact of Structured Output Format on Large Language Models through Causal Inference. In *Findings of the Association for Computational Linguistics: EACL 2026*. 1771–1795.
- [77] Huan Zhang, Wei Cheng, Yuhan Wu, and Wei Hu. 2024. A Pair Programming Framework for Code Generation via Multi-Plan Exploration and Feedback-Driven Refinement. In *Proceedings of the 39th IEEE/ACM International Conference on Automated Software Engineering*. 1319–1331.
- [78] Kechi Zhang, Jia Li, Ge Li, Xianjie Shi, and Zhi Jin. 2024. CodeAgent: Enhancing Code Generation with Tool-Integrated Agent Systems for Real-World Repo-level Coding Challenges. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 13643–13658.
- [79] Shaokun Zhang, Ming Yin, Jieyu Zhang, Jiale Liu, Zhiguang Han, Jingyang Zhang, Beibin Li, Chi Wang, Huazheng Wang, Yiran Chen, et al. 2025. Which agent causes task failures and when? on automated failure attribution of llm multi-agent systems. *arXiv preprint arXiv:2505.00212* (2025).
- [80] Yuntong Zhang, Haifeng Ruan, Zhiyu Fan, and Abhik Roychoudhury. 2024. Autocoderever: Autonomous program improvement. In *Proceedings of the 33rd ACM SIGSOFT International Symposium on Software Testing and Analysis*. 1592–1604.
- [81] Dewu Zheng, Yanlin Wang, Ensheng Shi, Ruikai Zhang, Yuchi Ma, Hongyu Zhang, and Zibin Zheng. 2024. Towards more realistic evaluation of LLM-based code generation: an experimental study and beyond. *arXiv preprint arXiv:2406.06918* (2024).
- [82] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in neural information processing systems* 36 (2023), 46595–46623.
- [83] Qihao Zhu, Daya Guo, Zhihong Shao, Dejian Yang, Peiyi Wang, Runxin Xu, Y Wu, Yukun Li, Huazuo Gao, Shirong Ma, et al. 2024. Deepseek-coder-v2: Breaking the barrier of closed-source models in code intelligence. *arXiv preprint arXiv:2406.11931* (2024).