

Requirements Development and Formalization for Reliable Code Generation: A Multi-Agent Vision

Xu Lu

*ICTT and ISN Lab
Xidian University
Xi'an, China
xlu@xidian.edu.cn*

Weisong Sun*

*Nanyang Technological University
Singapore
weisong.sun@ntu.edu.sg*

Yiran Zhang

*Nanyang Technological University
Singapore
yiran002@e.ntu.edu.sg*

Ming Hu

*Singapore Management University
Singapore
ecnu_hm@163.com*

Cong Tian

*ICTT and ISN Lab
Xidian University
Xi'an, China
ctian@mail.xidian.edu.cn*

Zhi Jin

*Wuhan University
Wuhan, China
zhijin@whu.edu.cn*

Yang Liu

*Nanyang Technological University
Singapore
yangliu@ntu.edu.sg*

Abstract—Automated code generation has long been considered the holy grail of software engineering. The emergence of Large Language Models (LLMs) has catalyzed a revolutionary breakthrough in this area. However, existing methods that only rely on LLMs remain inadequate in the quality of generated code, offering no guarantees of satisfying practical requirements. They lack a systematic strategy for requirements development and modeling. Recently, LLM-based agents typically possess powerful abilities and play an essential role in facilitating the alignment of LLM outputs with user requirements. In this paper, we envision the first multi-agent framework for reliable code generation based on Requirements Development and Formalization, named ReDeFo. This framework incorporates three agents, highlighting their augmentation with knowledge and techniques of formal methods, into the requirements-to-code generation pipeline to strengthen quality assurance. The core of ReDeFo is the use of formal specifications to bridge the gap between potentially ambiguous natural language requirements and precise executable code. ReDeFo enables rigorous reasoning about correctness, uncovering hidden bugs, and enforcing critical properties throughout the development process.

Index Terms—automated code generation, formal specification, large language model, multi-agent

I. INTRODUCTION

Automated Code Generation (ACG) [1] has long been regarded as an elusive yet highly desirable goal in software engineering, promising to dramatically improve productivity and reduce development effort. The recent rise of Large Language Models (LLMs) has cast a new light on this vision, offering unprecedented capabilities in understanding, reasoning, and synthesizing code that align with human intent [2], [3]. LLM-based or LLM-agent-based ACG techniques accelerate software development by an end-to-end transformation from

high-level specifications, which are typically expressed in Natural Language Requirements (NLRs) to executable code [4]–[7]. However, the quality and correctness of generated code often remain uncertain. The primary challenge is the lack of guarantees about the consistency between high-level specifications and code, which can lead to unreliable software and elevated maintenance costs [8]. Therefore, the need for a sound transition from ambiguous NLRs to precise executable code is highlighted, with a systematic strategy for requirements development and modeling being especially essential.

Formal methods provide a rigorous approach to establishing software correctness by mathematically proving that a program adheres to its formal specifications [9]. Applications of formal methods in the broader domain of program verification have achieved notable success, including detecting subtle bugs, enforcing safety and security guarantees, and supporting high-assurance software development [10]. Naturally, this leads us to consider whether ACG can be integrated with formal methods, as such integration could revolutionize the way we produce software by ensuring that the generated code is not only syntactically correct but also semantically faithful to the given requirements.

In this paper, we propose a multi-agent framework, Requirements Development and Formalization (ReDeFo), to address this challenge. The framework begins with NLRs and proceeds through formal specification [11] to verified code generation. It leverages distinct types of agents together that work collaboratively to bridge the gap between informal human intent and provably correct software artifacts. To the best of our knowledge, this is the first multi-agent framework for ACG with formal correctness guarantees.

II. ReDeFo FRAMEWORK

Figure 1 depicts a multi-agent pipeline of ReDeFo. The overall process involves three functionally specialized agents, i.e., Analyst, Formalizer, and Coder who are responsible for

*Weisong Sun is the corresponding author. This research is supported by the National Natural Science Foundation of China (62192734, 62192731, 62372347), the Fundamental Research Funds for the Central Universities (ZYTS25076), and the National Research Foundation, Singapore, and DSO National Laboratories under the AI Singapore Programme (AISG Award No: AISG4-GC-2023-008-1B).

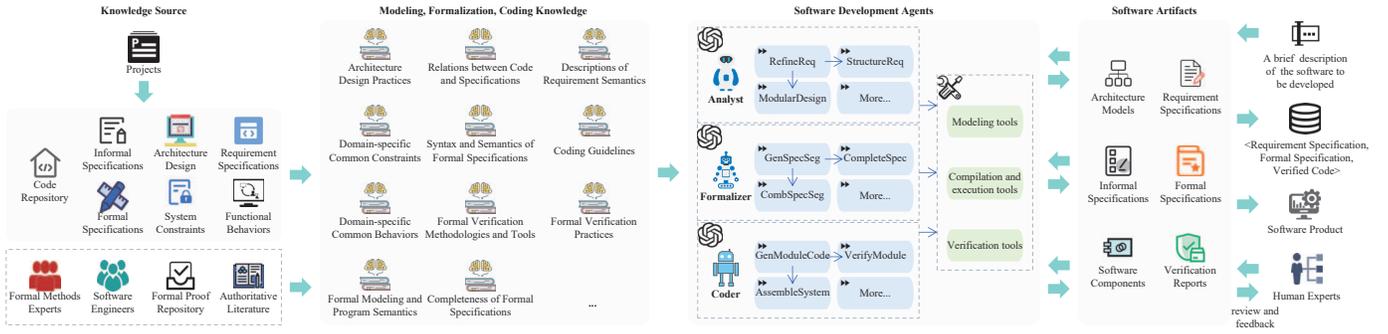


Fig. 1. Overview of the proposed framework REDeFo

completing the requirements development, formalization, and code generation tasks, respectively.

Upon receiving user-provided NLRs, the **Analyst** autonomously analyzes the text to refine each functionality into more detailed sub-requirements [12], [13]. Based on this understanding, it further extracts or infers system constraints and functional behaviors, identifies architectural components, and constructs an initial architecture model. The aforementioned results are then organized and summarized into informal, yet structured specifications. Further, the **Formalizer** is responsible for translating the informal specifications into formal specifications using domain-specific logic, formal semantics, and modeling guidelines. Finally, the **Coder** uses both the refined NLRs and formal specifications to generate code modules that meet the given requirements. Formal verification is then performed on the architecture model and generated code to guarantee correctness with respect to the specifications.

REDeFo advocates human-in-the-loop interaction by incorporating expert judgment at critical points and facilitating more active expert participation beyond post-hoc review. Based on feedback produced throughout the pipeline, human experts can either review intermediate artifacts or respond to clarification questions and confirmation requests posed by the agents. The feedback may comprise verification failures, unmet constraints, or inconsistencies between artifacts, etc. Through a modular and collaborative workflow, the framework produces a comprehensive set of artifacts, including requirement specifications, architecture model, informal and formal specifications, verified architecture model/software components, and verification reports, thus culminating in a trustworthy software product and, more importantly, a high-quality corpus of aligned data.

A. Agents in REDeFo

1) **Analyst**: The Analyst is responsible for interpreting and structuring the NLRs. To fulfill its role effectively, the Analyst performs the following actions: (1) Refinement and slicing of requirements, where each function is analyzed in detail to uncover implicit constraints or behaviors and express them as self-contained, formalization-ready sub-requirements. (2) Elaboration of vague or underspecified requirements, refining ambiguous or high-level descriptions into concrete, actionable statements. (3) Identification of system constraints or behaviors

that are directly formalizable. (4) Functional decomposition and architectural design, whereby the overall system is broken down into modular components that reflect logical or domain-specific divisions.

The above actions are supported by a range of necessary domain knowledge and engineering expertise: (1) Basic understanding of formal specification languages and their capabilities, enabling the identification of formalizable content within informal text. (2) Methods for extracting semantic meaning from NLRs, including structured techniques in requirements engineering [14] such as use case analysis, goal modeling, and scenario-based modeling. (3) Skills in drafting requirements that closely mirror formal semantics, thereby easing the translation to formal specifications in the next phase. (4) Expertise in system architecture and modular design, allowing the agent to produce a viable and maintainable component layout.

2) **Formalizer**: The Formalizer undertakes the task of translating informal requirements into precise, unambiguous formal specifications. It engages in several key actions to accomplish the task: (1) Selection of suitable formalisms and generation of formal specification segments from informal requirement slices [15]–[18]. (2) Composition of specification segments by merging similar or overlapping statements that describe the identical or logically entailed underlying semantics, thus preserving coherence and reducing redundancy. (3) Completion of specifications, which serves to infer and explicitly describe implicit semantics not fully captured in the informal requirements. (4) Consistency assessment, whereby the agent evaluates the logical soundness and completeness of the specifications to detect contradictions, gaps, or ambiguities.

The actions rely on a solid foundation of domain-specific formal knowledge: (1) Proficiency in formal specification syntax and semantics, including understanding of popular formalisms used in model checking [19], theorem proving [20], static analysis [21], etc. (2) Methodologies for formally describing system properties cover both functional and non-functional aspects, such as invariants [22], pre/post-conditions [23], and temporal logic assertions [24]. (3) Strategies for evaluating soundness and completeness of specifications [25], e.g., structural checks, behavioral coverage analysis, and traceability/consistency verification. (4) Semantic interpretation skills, supporting the detection of misalignments

between the informal intent and its formalized representation.

3) *Coder*: The Coder is tasked with programming the specifications into verified, executable code. The Coder organizes a series of coordinated actions: (1) Module code generation, where the agent synthesizes executable code based on the requirements and formal specifications for each module [26]. (2) Code & model verification, whereby the generated modules and models are checked against the formal specifications using appropriate formal verification techniques. (3) System assembly, in which individual modules are integrated into a complete software.

The Coder requires profound technical and formal understanding: (1) Professional practices in formal verification methodologies as well as their respective strengths, limitations, and applicable contexts. (2) Mapping strategies from formal specification to code, possibly including language-specific implementation patterns and automated code synthesis techniques [27]. (3) Expertise in modular composition and integration, a useful knowledge that system-level functionality and correctness are preserved during module integration.

B. Knowledge Acquisition for Agents

It is essential to endow the agents with deep and broad background knowledge in order to fully leverage their strong reasoning and generation capabilities.

1) *Real-World Software Projects*: Real-world software projects serve as the richest and most pragmatic source of knowledge for agents within the framework. We are confident that LLMs at the heart of agents have already internalized diverse programming paradigms, coding idioms, and implementation strategies used across different domains and contexts.

In contrast, NLRs in such projects are often sparse, outdated, or entirely missing. We can reverse-engineer requirements from the code itself, including identifying implicit functional behavior, reconstructing user stories, and deducing system-level constraints based on program logic and structure [28]. Similarly, we can analyze module boundaries, call hierarchies, and data flow patterns to infer the underlying software architecture [29]. Furthermore, informal and formal specifications are rarely present in real-world codebases. A feasible way is to extract informal specifications by summarizing source code and comments [28], and even attempt to derive formal specifications by identifying verifiable properties, preconditions, and invariants directly from implementations [30], [31]. By engaging in the above process, agents can gradually develop a comprehensive understanding of how requirements, architecture, specifications, and code are interrelated.

2) *Authoritative Literature*: Classic textbooks and reference works in software engineering offer systematic knowledge across core areas such as requirements engineering, software architecture, coding practices, testing, and project management. These materials define the best practices and design principles that help agents reason about software quality, scalability, and maintainability.

Equally important are texts that bridge formal methods with software engineering. Agents should learn the methodologies

for specifying, modeling, and verifying key properties using mathematical rigor. Topics such as model checking [32], deductive verification [33], theorem proving [34], abstract interpretation [35], and SMT (Satisfiability Modulo Theories) [36] are essential to enable the Formalizer and Coder to produce reliable, provably correct artifacts.

3) *Professional Experts*: Obviously, two major categories of experts are relevant to ReDeFo, i.e., software engineers and formal methods experts. Software engineers contribute by validating functional requirements, refining architecture designs, and promoting best practices in code generation. Their experience provides grounded insight into practical development processes, common pitfalls, and robust design patterns. Formal methods experts, on the other hand, play a crucial role in verifying specification quality, guiding the formulation of formal properties, and applying proof techniques. They possess the know-how to select and use formal verification tools [37]–[40] effectively, interpret verification outcomes, and ensure that specifications are both expressive and provable. By simulating the knowledge and feedback from these experts (e.g., Fine-tuning [41], CoT [42] and RAG [43]), agents can operate with greater accuracy and adaptability, marrying theoretical rigor with practical feasibility.

III. CASE STUDY

In this section, we demonstrate the potential capabilities of ReDeFo through a practical example. The selected case is typical control software in the safety-critical aerospace domain, which imposes strict constraints on functional correctness, compliance with domain standards, and hardware-dependent operations. The system is named Solar Search (SS), which aims to accomplish the functionality of solar acquisition. SS is composed of several subsystems, with the mode-switching component as the central one. The process of solar acquisition involves four separate stages, namely modes. The four modes switch from each other based on specific conditions. ReDeFo is applied to generate the mode-switching component with intermediate outputs to show the stepwise progression (Figure 2).

The coarse NLRs clearly define four modes, along with switch conditions between them. The refined NLRs are built upon the coarse ones, enriched with more detailed inputs, outputs and functional descriptions of the four modes. Upon reviewing the refined NLRs, the Analyst produces a component diagram and a state diagram to represent the architecture. The content extracted from the system-related NLRs is transformed into system-level properties, capturing the high-level behavioral expectations of the system. The content extracted from the mode-specific NLRs is mapped to mode-level properties, focusing on the low-level input, output and behavior within individual operational modes.

The semantic meaning of the requirements is further interpreted by the Formalizer, and necessary constraints are applied to eliminate any remaining ambiguities. Here the high-level informal specifications are transformed into LTL (Linear Temporal Logic) specifications [44], and the low-level

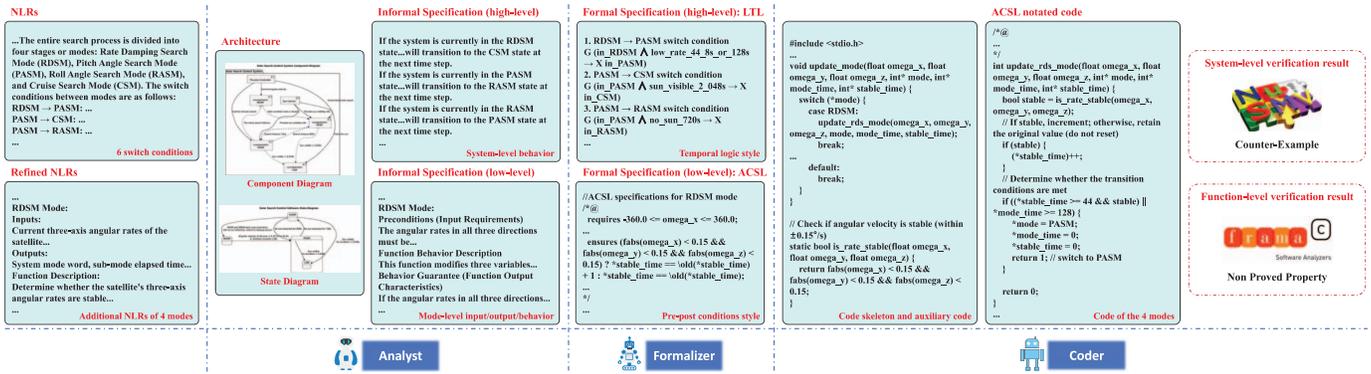


Fig. 2. Main artifacts of ReDeFo on a practical example

informal specifications into ACSL (ANSI/ISO C Specification Language) specifications [45].

The Coder generates C code for each mode according to the low-level formal specifications and the refined NLRs. The Coder also generates a code skeleton to integrate the four modes. Subsequently, the Coder performs the verification on two aspects of the system. First, model checking the behavior model of the system (state diagram) against the LTL specifications using the model checker NuSMV [46]. Second, deductive reasoning is applied to each mode function annotated with ACSL specifications using the program analyzer Frama-C [47].

IV. CHALLENGES AND OPPORTUNITIES

A. Completeness and Consistency of Formal Specifications

One of the key challenges is how to generate formal specifications that are sufficiently complete. What we mean by “complete” is that formal specifications are capable of specifying all essential aspects of the NLRs. Incomplete specifications may lead to a divergence between what the system is expected to do and what the verification process can guarantee. The Formalizer encounters difficulties in inferring all relevant system properties, especially those that are implicit or only loosely defined. Therefore, specification completeness requires not only the ability to parse and formalize explicit requirements but also to infer hidden assumptions, identify missing properties, and reason about potential corner cases.

Another critical issue lies in maintaining consistency between the informal requirements and the corresponding formal specifications. Misalignment may occur due to semantic drift during translation, misunderstanding of domain-specific language, or the inherent difficulty of preserving intent across different levels of abstraction. Such inconsistency undermines the trustworthiness of the entire pipeline. Even if the code is formally verified against the specification, any mismatch between the specification and the original intent still renders the system incorrect from the user’s perspective.

B. High-Quality Code from Formal Specifications

The challenge pertains to the Coder’s ability to synthesize high-quality, verifiable code from formal specifications. Although formal specifications provide a clear semantic target,

the actual process of generating robust, maintainable, and idiomatic code that meets this target is non-trivial. Similar to informal requirements, the semantic gap between formal specifications and executable code still poses a fundamental difficulty. Some specifications may admit many correct implementations, and choosing the most suitable one requires understanding performance constraints, usability factors, and deployment contexts. Moreover, verification tools vary in capability and may impose restrictions on how code must be structured to be verifiable.

C. Practical Deployment

The real-world deployment of our framework raises several practical challenges. As mentioned earlier, the availability of high-quality training data for each agent remains a concern. Informal and formal specifications are often missing or under-documented in real-world software projects, making it difficult to pretrain or fine-tune the agents effectively. The human-in-the-loop review process, although beneficial for quality assurance, introduces additional complexity. It is non-trivial to define when and how expert feedback should be requested, how it is incorporated into agent decisions, and how to balance automation with human judgment. Scalability and performance pose another important engineering concerns. As the system scales to larger projects with numerous components and deeper formal properties, the computational cost of verification and the complexity of code synthesis may grow substantially.

V. CONCLUSION

This paper proposes ReDeFo that combines software engineering and formal methods to achieve verifiable ACG. By dividing responsibilities among Analyst, Formalizer, and Coder agents, the framework creates a structured pipeline from NLRs to verified executable code, supporting consistency, traceability, and correctness throughout the software development lifecycle. Aiming to improve each agent’s powerful reasoning and generation capabilities, we emphasize the significance of deep background knowledge. While promising, the framework still faces key challenges such as ensuring specification completeness, maintaining semantic alignment, and producing requirement-compliant verified code.

REFERENCES

- [1] A. Church, "Logic, arithmetic, and automata," in *Proceedings of the international congress of mathematicians*, vol. 1962, 1962, pp. 23–35.
- [2] X. Bi, D. Chen, G. Chen, S. Chen, D. Dai, C. Deng, H. Ding, K. Dong, Q. Du, Z. Fu *et al.*, "Deepseek llm: Scaling open-source language models with longtermism," *arXiv preprint arXiv:2401.02954*, 2024.
- [3] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat *et al.*, "Gpt-4 technical report," *arXiv preprint arXiv:2303.08774*, 2023.
- [4] J. Jiang, F. Wang, J. Shen, S. Kim, and S. Kim, "A survey on large language models for code generation," *arXiv preprint arXiv:2406.00315*, 2024.
- [5] S. Hong, X. Zheng, J. Chen, Y. Cheng, J. Wang, C. Zhang, Z. Wang, S. K. S. Yau, Z. Lin, L. Zhou *et al.*, "Metagpt: Meta programming for multi-agent collaborative framework," *arXiv preprint arXiv:2308.00352*, vol. 3, no. 4, p. 6, 2023.
- [6] C. Qian, W. Liu, H. Liu, N. Chen, Y. Dang, J. Li, C. Yang, W. Chen, Y. Su, X. Cong *et al.*, "Chatdev: Communicative agents for software development," *arXiv preprint arXiv:2307.07924*, 2023.
- [7] B. Wei, "Requirements are all you need: From requirements to code with llms," in *2024 IEEE 32nd International Requirements Engineering Conference (RE)*. IEEE, 2024, pp. 416–422.
- [8] J. Liu, C. S. Xia, Y. Wang, and L. Zhang, "Is your code generated by chatgpt really correct? rigorous evaluation of large language models for code generation," *Advances in Neural Information Processing Systems*, vol. 36, pp. 21 558–21 572, 2023.
- [9] E. M. Clarke and J. M. Wing, "Formal methods: State of the art and future directions," *ACM Computing Surveys (CSUR)*, vol. 28, no. 4, pp. 626–643, 1996.
- [10] J. Woodcock, P. G. Larsen, J. Bicarregui, and J. Fitzgerald, "Formal methods: Practice and experience," *ACM computing surveys (CSUR)*, vol. 41, no. 4, pp. 1–36, 2009.
- [11] N. Nissanke, *Formal Specification: Techniques and Applications*. Springer Science & Business Media, 2012.
- [12] J. Huang, D. Jin, W. Sun, Y. Liu, and Z. Jin, "Envisioning intelligent requirements engineering via knowledge-guided multi-agent collaboration," in *Proceedings of the 40th International Conference on Automated Software Engineering - New Ideas and Emerging Results Track*. Seoul, South Korea: IEEE Computer Society, Sun 16 - Thu 20 November 2025, pp. 1–6.
- [13] M. Ataei, H. Cheong, D. Grandi, Y. Wang, N. Morris, and A. Tessier, "Elicitor: An llm agent-based simulation framework for design requirements elicitation," *arXiv preprint arXiv:2404.16045*, 2024.
- [14] P. A. Laplante and M. Kassab, *Requirements engineering for software and systems*. Auerbach Publications, 2022.
- [15] C. Hahn, F. Schmitt, J. J. Tillman, N. Metzger, J. Siber, and B. Finkbeiner, "Formal specifications from natural language," *arXiv preprint arXiv:2206.01962*, 2022.
- [16] F. Fuggitti and T. Chakraborti, "NL2LTL – a python package for converting natural language (NL) instructions to linear temporal logic (LTL) formulas," in *AAAI*, 2023, system Demonstration.
- [17] J. Cao, Y. Lu, M. Li, H. Ma, H. Li, M. He, C. Wen, L. Sun, H. Zhang, S. Qin *et al.*, "From informal to formal—incorporating and evaluating llms on natural language requirements to verifiable formal proofs," *arXiv preprint arXiv:2501.16207*, 2025.
- [18] Y. Wu, A. Q. Jiang, W. Li, M. Rabe, C. Staats, M. Jamnik, and C. Szegedy, "Autoformalization with large language models," *Advances in Neural Information Processing Systems*, vol. 35, pp. 32 353–32 368, 2022.
- [19] C. Baier and J.-P. Katoen, *Principles of model checking*. MIT press, 2008.
- [20] D. W. Loveland, *Automated theorem proving: A logical basis*. Elsevier, 2016.
- [21] X. Rival and K. Yi, "Introduction to static analysis," 2020.
- [22] A. F. Donaldson, L. Haller, D. Kroening, and P. Rümmer, "Software verification using k-induction," in *Static Analysis: 18th International Symposium, SAS 2011, Venice, Italy, September 14-16, 2011. Proceedings 18*. Springer, 2011, pp. 351–368.
- [23] C. A. R. Hoare, "An axiomatic basis for computer programming," *Communications of the ACM*, vol. 12, no. 10, pp. 576–580, 1969.
- [24] N. Rescher and A. Urquhart, *Temporal logic*. Springer Science & Business Media, 2012, vol. 3.
- [25] C. Menghi, E. Balai, D. Valovcin, C. Stickse, and A. Rajhans, "Completeness and consistency of tabular requirements: An smt-based verification approach," *IEEE Trans. Software Eng.*, vol. 51, no. 2, pp. 595–620, 2025. [Online]. Available: <https://doi.org/10.1109/TSE.2025.3530820>
- [26] W. Murphy, N. Holzer, F. Qiao, L. Cui, R. Rothkopf, N. Koenig, and M. Santolucito, "Combining llm code generation with formal specifications and reactive program synthesis," *arXiv preprint arXiv:2410.19736*, 2024.
- [27] R. Bloem, B. Jobstmann, N. Piterman, A. Pnueli, and Y. Sa'ar, "Synthesis of reactive (1) designs," *Journal of Computer and System Sciences*, vol. 78, no. 3, pp. 911–938, 2012.
- [28] W. Sun, Y. Miao, Y. Li, H. Zhang, C. Fang, Y. Liu, G. Deng, Y. Liu, and Z. Chen, "Source code summarization in the era of large language models," in *Proceedings of the 47th International Conference on Software Engineering*, no. 1. Ottawa, Ontario, Canada: IEEE Computer Society, 27 April-3 May 2025, pp. 419–431.
- [29] Y. Zhang, R. Li, P. Liang, W. Sun, and Y. Liu, "Knowledge-based multi-agent framework for automated software architecture design," in *Proceedings of the 33rd ACM International Conference on the Foundations of Software Engineering-Ideas, Visions and Reflections*. Trondheim, Norway: ACM, Mon 23 - Fri 27 June 2025, pp. 1–5.
- [30] C. Wen, J. Cao, J. Su, Z. Xu, S. Qin, M. He, H. Li, S.-C. Cheung, and C. Tian, "Enchanting program specification synthesis by large language models using static analysis and program verification," in *International Conference on Computer Aided Verification*. Springer, 2024, pp. 302–328.
- [31] L. Ma, S. Liu, Y. Li, X. Xie, and L. Bu, "Specgen: Automated generation of formal program specifications via large language models," *arXiv preprint arXiv:2401.08807*, 2024.
- [32] S. Merz, "The specification language tla+," *Logics of specification languages*, pp. 401–451, 2008.
- [33] P. Baudin, J.-C. Filliâtre, C. Marché, B. Monate, Y. Moy, and V. Prevosto, *ACSL: ANSI/ISO C Specification Language*. [Online]. Available: <http://frama-c.com/download/acsl.pdf>
- [34] Y. Bertot and P. Castéran, *Interactive theorem proving and program development: Coq/Art: the calculus of inductive constructions*. Springer Science & Business Media, 2013.
- [35] P. Cousot, *Principles of abstract interpretation*. MIT Press, 2021.
- [36] C. Barrett and C. Tinelli, "Satisfiability modulo theories," *Handbook of model checking*, pp. 305–343, 2018.
- [37] R. Cavada, A. Cimatti, M. Dorigatti, A. Griggio, A. Mariotti, A. Micheli, S. Mover, M. Roveri, and S. Tonetta, "The nuxmv symbolic model checker," in *Computer Aided Verification: 26th International Conference, CAV 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 18-22, 2014. Proceedings 26*. Springer, 2014, pp. 334–342.
- [38] L. Correnson, P. Cuoq, F. Kirchner, A. Maroneze, V. Prevosto, A. Puccetti, J. Signoles, and B. Yakobowski, *Frama-C User Manual*. [Online]. Available: <http://frama-c.com/download/frama-c-user-manual.pdf>
- [39] D. Beyer and M. E. Keremoglu, "Cpachecker: A tool for configurable software verification," in *Computer Aided Verification: 23rd International Conference, CAV 2011, Snowbird, UT, USA, July 14-20, 2011. Proceedings 23*. Springer, 2011, pp. 184–190.
- [40] L. De Moura and N. Bjørner, "Z3: An efficient smt solver," in *International conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 2008, pp. 337–340.
- [41] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, W. Chen *et al.*, "Lora: Low-rank adaptation of large language models," *ICLR*, vol. 1, no. 2, p. 3, 2022.
- [42] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou *et al.*, "Chain-of-thought prompting elicits reasoning in large language models," *Advances in neural information processing systems*, vol. 35, pp. 24 824–24 837, 2022.
- [43] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel *et al.*, "Retrieval-augmented generation for knowledge-intensive nlp tasks," *Advances in neural information processing systems*, vol. 33, pp. 9459–9474, 2020.
- [44] A. Pnueli, "The temporal logic of programs," in *18th annual symposium on foundations of computer science (sfcs 1977)*. IEEE, 1977, pp. 46–57.
- [45] P. Baudin, J.-C. Filliâtre, C. Marché, B. Monate, Y. Moy, and V. Prevosto, "Acsl: Ansi/iso c specification," *URL <https://frama-c.com/html/acsl.html>*, 2021.

- [46] A. Cimatti, E. Clarke, E. Giunchiglia, F. Giunchiglia, M. Pistore, M. Roveri, R. Sebastiani, and A. Tacchella, "Nusmv 2: An opensource tool for symbolic model checking," in *Computer Aided Verification: 14th International Conference, CAV 2002 Copenhagen, Denmark, July 27–31, 2002 Proceedings 14*. Springer, 2002, pp. 359–364.
- [47] F. Kirchner, N. Kosmatov, V. Prevosto, J. Signoles, and B. Yakobowski, "Frama-c: A software analysis perspective," *Formal aspects of computing*, vol. 27, no. 3, pp. 573–609, 2015.