

Enhancing Software Development through Prompt Engineering A Study on Large Language Models for Code Generation and De- veloper Productivity

**¹Jimit Patel, ²Meet Bipinchandra Patel, ³Nishil Sureshkumar Prajapati, ⁴Rahul Rathi,
⁵Raghavendra Kamarthi Eranna, ⁶Pratikkumar Prajapati, ⁷Krishna Chaitanya Chittoor**

Submitted:04/11/2024

Accepted:17/12/2024

Published:27/12/2024

Abstract: Large Language Models (LLMs) are increasingly changing Software Development with capabilities to generate code snippets, debug, etc., and towards design work as well. Successful outcomes using LLMs is heavily reliant on prompt engineering. Well-designed prompts influence the quality of generated code, improve developer workflows, and build effective human-computer interactivity in the use of LLM models. This study examines prompt engineering in improving developer productivity via a designed process of exploration of prompting strategies to generate code. A taxonomy of potential prompt engineering techniques is introduced conceptualizing four experimental approaches for the coding task: instruction-based prompts, example-based prompts, chain-of-thought prompts, and hybrid prompts. The study focuses on developer-oriented productivity metrics beyond technical quality. Productivity metrics include a reduction in overall development time, reduced errors and better readability, e.g. improved structure of codes, and found improvements to tools used to develop software, e.g. integrated development environments, collaborative coding tools. The comparative evaluation of prompt patterns identifies how differentiating prompt patterns can create variable impact on code quality, but also variable experiences for developers. This suggests that prompt engineering can influence the continuing problem of debugging

¹Staff Software Engineer At Very Good Security

jimit7patel@gmail.com

²Senior Manager, Data and AI/ML engineering

Meet61@gmail.com

³Lead Cloud Development Engineer,

nishilp017@gmail.com

⁴BI- Manager

Rathirahul53@gmail.com

⁵System Analyst

keraghu@gmail.com

⁶Senior Manager, Data Engineering

pratik.prajapati020@gmail.com

⁷Principal Data Engineer

chaitueie17@gmail.com

and support more rapid delivery of software to clients. The paper describes barriers to adoption in practice, such as prompt sensitivity, context limitations and reusability limitations and offer a roadmap for integrating adaptive prompting systems directly in developer environments. By relating LLM capabilities to productivity outcomes, this work provides a new perspective in bridging prompt engineering research with real-world software development pipelines.

Keywords: *Prompt engineering, code generation, software development productivity, large language models (LLMs), create tools, AI in software engineering, intelligent programming assistants*

1. Introduction

The explosion of artificial intelligence has fundamentally affected the way software is developed, tested and maintained. Among the many developments, large language models (LLMs) clearly emerged as intelligent flexible assistants (that can produce source code, document auto-magically, and reduce cognitive load on developers). While these developments are obviously a huge advancement of software engineering practices, the actual productivity gains to developers' efficiency are heavily reliant on the design and structure of prompts. For this reason, prompt engineering has become somewhat of a new competency in bridging the divide between what LLMs can produce on the theoretical plane, and what a developer is looking for in real world work. Even though tools such as code assistants and automated programming interfaces are developing quickly, there remains a lack of knowledge as to how particular prompt patterns directly relate to software development productivity. This lack of knowledge has paradoxically created new opportunities and obstacles to the research community and industry stakeholders alike, highlighting the importance of studying prompt engineering not only as a technical optimization, but as a productivity enabler as well.

Recent research has documented the expanding use of large language models for software engineering functions, such as code generation, program repair, and software testing [1–3]. Despite the large language models showing promising performance and accuracy when generating syntactically correct content, there remain questions around their effectiveness in improving developer productivity [4–6]. The variability in outcomes produced from similar queries, the challenge of prompt building, and other risk factors will dictate whether LLMs will provide productivity or frustration. This dependency introduces further challenges that will require developers to learn the art of prompt engineering and use the large-language model programming environment effectively [7–9].

The discussion of developer productivity has always been associated with measurable factors (e.g., code quality, defect reduction, time-to-solution, maintainability) [10–12]. Although, in the literature, there are conceptual discussions of efficiency or correctness of algorithms, there has been very little research which considers practical productivity measures that fit with everyday developer workflows [13–15]. There is very little universal agreement on how to validly measure the real contribution of LLM-based code assistants to software engineering general practices [16,17]. Prompt engineering is an important factor in these practices, as it is the engagement layer between human intent and machine output. This study will engage with this dimension of prompt engineering to enact a sense of frameworks that connect prompt design and factors of productivity through a transparent and measurable context across the software development lifecycle [18,19].

Although the results were optimistic, many of the challenges remain unsolved. There is uncertainty in the results due to prompt sensitivity, token limitations restrict usability for larger projects and the learning curve associated with constructing effective prompts is high [20–22]. In addition, it is not clear how productivity gains can be compared between prompt engineering methods in different coding environments because there are no standardized benchmarks to determine productivity in coding [23,24]. These limitations however highlight the necessity of further research that considers the productivity role of LLMs for developers, rather than just their technical correctness.

This research attempts to fill the above gaps by systematically investigating how prompt engineering impacts both the outcome (quality) and efficiency of code generation by an LLM. This work focuses on developer-centric outcomes (time savings, error avoidance, improved workflow), as opposed to existing literature focusing on algorithmic accuracy. This paper contributes in 3 ways:

- The paper offers a taxonomy of prompt engineering techniques that are relevant to software development.

- The paper proposes developer productivity measures to evaluate the impact of LLM-based assistants.
- The paper provides a case study of how LLMs are incorporated into software engineering pipelines in practice, demonstrating both opportunities and barriers to adoption.

2. Literature Review

Large Language Models (LLMs) represent powerful new actors in the software engineering space, ranging from code creation, to automated testing and drafter of documentation. Prompt engineering is increasingly attracting attention, due to the observation that the quality, form, and specificity of prompts have important effects on the outputs from LLM-based systems. For this reason, both academic research and industry reports have begun to examine how prompt engineering might be used to positively engage with developer productivity and efficiency when the software is developed in context. This literature review will consider the state of the research in four major areas: (1) LLMs and software engineering, (2) prompt engineering methods, (3) empirical evidence of improvements in productivity, and (4) challenges and limitations in terms of adoption.

2.1. Large Language Models in Software Engineering

The application of LLMs into software engineering has developed across multiple domains: code generation, bug finding, program repair, and software testing. Multiple surveys and systematic reviews have mentioned several applications. Initial applications demonstrated LLMs could generate code, at the function level, with decent accuracy, but it was difficult to depend on the results or to explain why or how the LLM produced a given answer [51, 52, 63, 70]. Later reviews referenced the increasing use of LLMs in more specialized areas, such as automated program repair [67, 72] and model-driven engineering [75].

In addition to capability mapping, researchers have also documented scenarios where LLMs were scalable enough for large software projects, where it was found that, while generative outputs were useful for prototyping, production-grade software still had limits on correctness and maintainability [58, 65]. These findings, along with those from earlier chapters, highlight the importance of combining LLMs with well-designed prompts to maximize the productivity of their use, rather than being reliant on generation alone.

2.2. Prompt Engineering Strategies

Prompt engineering has developed into a key avenue to improve the interaction between developers and LLMs. Scholars have established different ways prompts function, including instruction-based prompting, chain-of-thought reasoning, and prompted templates, that offer different levels of impacts on output quality [53, 56, 64]. As shown in systematic studies, even minor changes in wording or structure can have large differences on correctness and code efficacy. For example, prompts that are structured in a concrete way (e.g., specifying language, variables, performance requirements) produce more consistent outputs in CRUD and pedagogical-based coding tasks [53, 60, 62]. Additionally, more advanced prompting strategies, such as prompt optimization frameworks [77], prompt chaining, and prompt refactoring have been analyzed for reducing hallucinations and improving consistency. The studies reviewed provide evidence that prompt engineering is not only a user practice, but a research-driven optimization technique that affects productively developer outcomes.




Figure 1. Prompt Engineering Strategies in LLM-Driven Software Development

2.2.1. Zero-Shot Prompting

Zero-shot prompting allows large language models (LLMs) to produce code from natural language input without examples. It may be used in several ways in software engineering such as boilerplate generation, implementation of basic algorithms, and iterating

prototypes very fast. An example may be provided where a developer could simply state "Write a Python function for factorial using recursion" and receive the entire function in one shot! This is helpful for automating the mundane, speeding up experimentation, however it comes with some limitations such as inconsistent quality and project standards.




Figure 2. Zero-Shot Prompting Workflow in Software Engineering

The developer query is the starting point, where intent is conveyed in natural language without examples. Accuracy depends on how clearly functionality, constraints, and expected behaviour are defined, as ambiguity here leads to ambiguous outputs. In the prompt generation phase, the query is structured into a prompt, where phrasing, tone, and detail (like language, style, or error handling) guide the outcome.

During LLM processing, the model leverages its pre-trained knowledge and generalization abilities to interpret the query and predict solutions without prior examples. This results in the code generation phase, where code snippets or scripts are produced varying in correctness, efficiency, and readability. Finally, in developer evaluation, the output is tested, refined, and iteratively improved. This feedback loop is essential, as zero-shot outputs are not always directly usable, but

refinement leads to better productivity and reduced coding effort.

2.2.2. Few-Shot Prompting

Few-shot prompting, as the term implies, is where the LLM is given a few examples along with a query from the developer. Instead of relying, as zero-shot prompting does, on the model inferring everything from a single instruction, few-shot prompting provides demonstrations with which to help the model produce better results in terms of accuracy and contextual relevance. As examples give the model reference points with

which to understand the structure, logic and coding style to conjecturally produce the desired response. When it comes to software engineering tasks - again, typically prompting and training the models on specific domains, few-shot prompting can greatly improve accuracy and coverage when generating unit tests, debugging code or formatting projects in a consistent manner. Few-shot prompting achieves a reasonable compromise between flexibility and accuracy, making it one of the easiest and most useful ways of augmenting productivity of real-world developers.




Figure 3. Few-Shot Prompting Workflow

The process begins with the developer prompt plus examples, where illustrative inputs and outputs provide context, reduce ambiguity, and guide generalization. In the prompt writing stage, both the task and examples are structured together for instance, showing 2–3 Python function examples before requesting a similar one.

During LLM processing, the model interprets the prompt and examples, recognizing patterns, logic, and formatting, which improves alignment with developer intent. This leads to contextualized code generation, where outputs reflect the same style and structure, enhancing accuracy and trust. Finally, in developer evaluation, the programmer verifies the output against examples and requirements, refining prompts or examples if needed for better results.

2.2.3. Chain-of-Thought Prompting

Prompting Chain-of-Thought (CoT) is an organized approach where you instruct the LLM to describe its reasoning step by step before its final output. CoT prompting allows the model to produce intermediate reasoning steps instead of jumping to a conclusion. This style of reasoning simulates the typical process a human user would follow for completing a task. CoT is effective for all software engineering tasks, such as debugging, designing an algorithm, and working on a complicated coding task. Explicitly constructing the model's reasoning allows developers to have greater transparency into the process the model uses to arrive at a solution, which lowers the chance of hidden errors and provides increased confidence with LLM-generated code. This method also enhances accuracy, while also allowing developers to be aware of other paths/solutions considered by the model.




Figure 4. Chain-of-Thought Prompting Workflow

The process begins with the **developer query**, where tasks may involve reasoning such as algorithm design or debugging. In the **prompt construction phase**, the developer instructs the LLM to “think step by step,” prompting it to display its reasoning rather than just the final code.

During **LLM processing**, the model breaks the problem into smaller reasoning steps identifying flows, conditions, or pseudocode which enhances transparency and helps detect logical errors. The **final output** is then generated from these steps, and in the **developer verification phase**, the solution is confirmed for both correctness and explainability. This makes CoT prompting highly effective for improving developer productivity.

2.2.4. Instruction-Based Prompting

Instruction-Based Prompting is probably the safest and most commonly used prompting technique, where developers provide direct natural language instructions to the LLM. The clarity and specificity of these instructions determines the quality of the output, although in general this technique is less strict about examples and reasoning traces due to reliance on the fundamental training of the model to execute commands. In software engineering, these instructions are commonly applied to such operations as converting requirements into code, producing documentation, or acquiring refactoring of an existing function. It is simple and flexible in the context of improving developer productivity.




Figure 5. Instruction-Based Prompting

The process begins when a developer issues an explicit instruction such as “Generate a Python function to parse JSON files” or “Refactor this Java code to improve readability” which is then interpreted by the

LLM through its natural language reasoning and programming knowledge. The model produces an output that may include code blocks, unit tests, or documentation, depending on the request. The developer

subsequently validates the output against the intended functionality, performance, and task requirements, refining it, if necessary, by editing the instruction or adding new constraints. This establishes an iterative feedback loop that rapidly converges on high-quality solutions, demonstrating how instruction-based prompting seamlessly integrates into developer workflows while reducing both time and cognitive effort in repetitive coding tasks.

2.2.5. Role/Persona-Based Prompting

Role or Persona-Based Prompting occurs when the developer specifies a role for the LLM, asking it to act as some kind of expert. The prompt might include requests to "act like a senior Python developer," "behave as a code reviewer," or "take the role of a software architect." Due to the definition of the persona, the output will be more context-sensitive, associated with the norms of professional developer workflows. The role/persona-based approach is most suitable for tasks that benefit from domain knowledge, coding standards compliance, or context-based reasoning while developing software.




Figure 6. Role/Persona-Based Prompting

In this technique, developers first define a specific role or persona for the language model, such as a code reviewer, software architect, or security auditor. The prompt is then enriched with task details and contextual knowledge of the project, along with constraints like coding standards or security rules. The model interprets the role-based intent and adapts its reasoning, tone, and expertise accordingly. This results in primary outputs such as code suggestions, while alternative interpretations may also be generated to reflect different perspectives. Developers review and validate these outputs, providing refinements or re-prompts when necessary. Over time, the continuous feedback loop helps fine-tune role definitions and enhances overall effectiveness,

making the interaction more reliable and closer to real-world software engineering workflows.

2.2.6. Test-Case Generation Prompting

The technique involves the use of LLMs to automatically create test cases from the requirements, specifications, or even code snippets. By embedding prompts in instructions e.g., "Generate unit tests for this function," or "Provide edge-case tests for login validation," the model produces a comprehensive list of test cases which often include positive, negative, and boundary test cases. This technique reduces developer effort, provides better coverage of tests, and results in higher software reliability.




Figure 7. Test-Case Generation Prompting

In this method, developers provide either source code, requirements, or a specification as input with instructions in the prompt to generate related test cases. The LLM uses contextual project knowledge, together with a specified testing framework to generate test cases. The model generates multiple types of test cases; unit tests, integration tests, negative tests, boundary tests, etc. The next step is for developers to review the generated tests for correctness and if more coverage is necessary, ask the model to iterate and improve. This feedback loop allows for the continuously improving relevancy and completion of generated test cases while decreasing the amount of work that the

developers have to do and speeding up the entire testing cycle.

2.2.7. Debugging/Error-Fixing Prompting

This method utilizes LLMs to identify, explain, and correct errors in code. Developers input code snippets with mistakes, asking for debugging along the lines of "Debug this function and provide updates". The model will analyze and hallucinate the syntax, logical issues, and even possible runtime issues, it will create code with fixes and offer an explanation as to why. This process decreases the time required for debugging in a manual way and creates much needed speed.




Figure 8. Debugging/Error-Fixing

This block diagram illustrates the end-to-end debugging pipeline. The process begins by having the developer provide erroneous code and pair it with a debugging prompt. Next, the LLM evaluates the error types, identifies the bugs (if any), and provides a rationale for the reason. Then, the LLM will generate

code suggestions with fixes. The developer will evaluate the outputs and provide feedback. If needed, the cycle could continue through iterative or feed.

2.2.8 Requirement-to-Code Prompting

Requirement-to-code prompting simply means turning natural language requirements into executable code using LLMs. In the conventional sense, the developer writes detailed specifications using formal languages. Most practitioners today write functionality a user needs in plain English (or any natural language). The model then parses the requirements, maps them to a structured format like

pseudo-code or a template and finally generates production-ready code. This capability improves productivity immensely by mitigating the cumbersome process of translating user requirements, creating a bridge for non-technical stakeholders with developers and improving the software development life cycle overall.




Figure 9: Requirement-to-Code

Table 1: Comparative Summary of Prompt Engineering Technique

Prompt Technique	Primary Application in Software Engineering	Strengths	Limitations
Zero-Shot Prompting	Quick code snippets, Simple automation tasks	Fast, no training data needed, easy to apply	May produce vague/inaaccurate results without context
Few-Shot Prompting	Code generation with specific style or format	Increases accuracy, adapts to coding style	Requires carefully chosen examples; scalability issues
Chain-of-Thought	Algorithm explanation, debugging logic	Improves reasoning quality, enhances interpretability	Slower, may generate verbose answers
Instruction-Based Prompting	Generating boilerplate code, API integration	Easy to design, highly flexible	Highly sensitive to wording, small changes affect output
Role/Persona-Based Prompting	Acting as a code reviewer, tutor or system architect	Produces context aware, role-specific responses	May overfit persona, sometimes inconsistent

Test-Case Generation Prompting	Automated test creation for software validation	Saves developer time, improves code reliability	Quality of tests depends on clarity of prompt
Debugging/Error-Fixing Prompting	Syntax correction, logical error fixing	Reduces debugging effort, improves productivity	May miss subtle context-specific bugs
Requirement to-Code Prompting	Rapid prototyping, requirement-driven coding	Bridges gap between client requirements & implementation	Risk of misinterpretation of ambiguous requirements

In this technique, software requirements expressed in natural language are systematically refined and transformed into structured prompts, enabling clarity and precision. The LLM then interprets these prompts to generate executable code or pseudocode aligned with the intended functionality. The correctness and reliability of the produced code are subsequently validated through rigorous testing mechanisms. Following validation, developers review and refine the output to ensure compliance with project standards and professional practices. This horizontally structured workflow illustrates the essence of requirement-driven prompting, effectively bridging client needs with practical implementation and demonstrating its value in real-world software engineering contexts.

2.3 Empirical Evidence of Improvements in Developer Productivity

Recent empirical studies have demonstrated that carefully engineered prompts lead to measurable gains in developer productivity when leveraging LLMs for software engineering tasks. These improvements manifest across multiple dimensions:

- **Time-to-Solution Reduction:** Developers using few-shot and instruction-based prompts reported up to 30-50% reduction in task completion time for code generation and debugging tasks [25]. The structured guidance embedded in the prompts reduced the cognitive overhead of rephrasing or rewriting code queries.

- **Code Quality and Correctness:** Chain-of-Thought (CoT) and Self-Consistency prompting showed measurable improvements in unit test pass rates, with some benchmarks reporting increases of 15-20% correctness compared to zero-shot baselines [31]. This suggests that reasoning-oriented prompts enhance logical soundness and prevent superficial, syntactic solutions.

- **Reusability and Scalability:** Retrieval-Augmented Generation (RAG) demonstrated strong performance in enterprise settings by leveraging API documentation, internal repositories, and domain-specific datasets. In experiments with enterprise-level repositories, developers experienced 40% fewer manual interventions when LLMs were supplemented with retrieval-enhanced prompts [46].

- **Developer Experience & Usability:** Role-based prompting and hybrid strategies improved developer satisfaction and trust. Controlled user studies revealed that developers found role-based prompts easier to align with real-world tasks such as code review and mentoring, thereby reducing frustration associated with prompt sensitivity [52].

2.4 Challenges and Limitations in Adoption

Despite the demonstrated advantages, the real-world adoption of prompt engineering within software development pipelines faces several barriers:

- **Prompt Sensitivity:** A single change in wording or example selection may drastically alter the LLM’s output [61]. This unpredictability can lead to developer frustration, especially in time-sensitive workflows.
- **Token and Context Constraints:** Many LLMs have strict context window limitations, restricting the number of examples or instructions that can be embedded. This makes it difficult to scale few-shot or hybrid approaches in large codebases [68].
- **Generalization vs. Specialization Trade-off:** Zero-shot and few-shot methods often fail in highly domain-specific scenarios, while RAG methods are heavily dependent on external retrieval systems. This creates a trade-off between breadth of application and depth of accuracy [71].
- **Reproducibility Issues:** Unlike traditional software engineering techniques, LLM outputs can vary across runs due to stochastic sampling methods. This lack of deterministic behavior raises concerns for mission-critical applications [77].
- **Learning Curve and Usability:** While role-based and hybrid prompting offer high potential, they often require expert knowledge in crafting optimal prompts, limiting accessibility for novice developers.
- **Integration Challenges:** Embedding LLM-driven prompting strategies within continuous integration/continuous deployment (CI/CD) pipelines is non-trivial. Latency, token costs, and external API dependencies remain practical obstacles for enterprise-scale adoption.

Table 2: Empirical Benefits vs. Adoption Challenges

Dimension	Empirical Gains	Adoption Challenges
Task Completion Time	30–50% faster with few-shot and instruction-based prompts	Sensitive to wording; not robust across domains
Code Correctness	15–20% higher test pass rates with CoT & self-consistency	Non-deterministic outputs across runs
Enterprise Integration	40% fewer manual interventions using RAG	Latency, dependency on retrieval systems
Developer Satisfaction	Higher trust with role-based prompting	Requires expertise in crafting effective prompts
Scalability	Hybrid prompting improves large-project workflows	Token/context window limitations

3. Methodology

We explored a variety of prompt engineering methodologies to improve software development productivity using LLMs, from various experimental and theoretical perspectives within a conventional

dynamic architecture and aligned to the real-world process of developing software. The methodology that we follow consists of delineating prompting techniques, establishing our experiments, composing a hyperplaned architecture, creating mathematical

models of the metrics, exploring validity concerns. One key feature is the synthesis of the iterative use of multiple prompt functions (e.g., zero-shot, chain-of-thought, request-code) with developer reflections, unit tests, and context markers. This sets up an expectation of improvement in developer productivity over time while maintaining the basic aspects of quality, reusability, and explainability - melding all together to be reproduced, extendable, and measurable.

3.1 Paper Selection Process

The paper selection process was carried out systematically to ensure that only the most relevant and high-quality resources were included for the analysis of prompt engineering in software development. The process began with the identification of three core dimensions:

1. **Large Language Models (LLMs)** – such as GPT, LLaMA, PaLM, and Codex, which represent the current state-of-the-art in code generation.
2. **Prompt Designs** – covering diverse strategies including zero-shot, few-shot, chain-of-thought, instruction-based, role-based, test-case generation, debugging, and requirement-to-code techniques.
3. **Code Tasks** – practical programming assignments, debugging challenges, requirement translation, and unit test generation, which reflect real-world developer workflows.

To ensure methodological rigor, this study applied a strict inclusion-exclusion criterion whereby papers and datasets were incorporated only if they (i) reported on LLM-based code generation or productivity, (ii) proposed or evaluated prompt strategies applicable to software engineering, and (iii) presented empirical results suitable for cross-benchmark comparison. Works that were speculative, irreproducible, or limited to non-software domains were excluded. The resulting pool of sources provided both diversity in tasks, models, and methodologies, and a balance between academic novelty and industrial practicality. This carefully curated selection not only strengthened the experimental design but also directly informed the

development of the proposed architecture presented in subsequent sections.

3.2 Experimental Setup

The experimental setup was designed to evaluate the impact of prompt engineering techniques on software development tasks using state-of-the-art LLMs. A multi-step process was followed to ensure consistency, reproducibility, and practical applicability.

Dataset Selection:

A diverse set of datasets was employed to represent typical developer workflows. These included HumanEval [1], MBPP - Mostly Basic Programming Problems [2], CodeXGLUE [3], and task-specific repositories curated from GitHub [4]. The datasets collectively covered algorithmic challenges, debugging scenarios, test-case generation, and requirement-to-code tasks, thereby ensuring wide coverage across the software engineering lifecycle.

LLMs Used:

Experiments were conducted on leading LLMs that are widely adopted in both academia and industry, including OpenAI Codex [5], GPT-3.5/4 [6], Google PaLM [7], and Meta's LLaMA series [8]. These models were chosen due to their demonstrated strengths in code synthesis, reasoning, and adaptability across different programming languages.

Baseline Prompts:

To provide a fair comparison, baseline prompts were designed to simulate real-world developer instructions without advanced engineering strategies. For example, a simple instruction such as “Write a Python function to calculate factorial” was used as a baseline against enhanced prompts (e.g., zero-shot, few-shot, or chain-of-thought). This allowed for a controlled measurement of productivity improvements when applying structured prompting techniques [9].

Productivity Proxies:

Developer productivity was measured using a set of **quantitative proxies** that align with practical software engineering outcomes. These proxies not only reflect traditional metrics of software quality but

also capture the **real-world value of prompt engineering** in reducing time, effort, and error rates during development.

- **Time-to-Solution:** Average time taken by the LLM to generate the correct or near-correct code.

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

Where t_i is the time taken by the LLM to generate a solution for task i , and N is the total number of tasks.

- **Correctness Rate:** Percentage of generated programs passing functional requirements.

$$CR = \frac{C}{N} * 100$$

Where C is the number of correctly generated solutions and N is the total number of tasks.

- **Test Pass Rate:** Fraction of automatically generated test cases successfully passed by the code.

$$TPR = \frac{\sum_{i=1}^N P_i}{\sum_{i=1}^N T_i} \times 100$$

Where P_i is the number of passed test cases for task i , and T_i is the total number of test cases for task i .

- **Bug Density:** Number of logical or syntactic errors per solution.

$$BD = \frac{\sum_{i=1}^N b_i}{\sum_{i=1}^N LOC_i}$$

Where b_i is the number of bugs in task i , and LOC_i is the lines of code generated for that task.

- **Developer Effort Reduction:** Measured as the reduction in manual corrections required for generated code.

$$DER = \left(1 - \frac{M}{M_{base}}\right) \times 100$$

Where M is the number of manual corrections required with prompt engineering, and M_{base} is the corrections required with baseline prompts.

3.3. Proposed Architecture

The proposed architecture introduces a Prompt-Driven Development Framework (PDDF) designed to seamlessly integrate Large Language Models (LLMs)

into the software engineering workflow with a primary focus on developer productivity enhancement. Unlike conventional approaches where LLMs are used as isolated assistants, PDDF treats LLMs as an embedded component of the development lifecycle, orchestrated through systematic prompt engineering strategies.

The architecture is structured into five interlinked layers, each contributing to the translation of natural language developer inputs into optimized, high-quality code artifacts:

1. **Input Layer (Developer Intent Capture):**
At this stage, the developer's requirements, expressed in natural language, are processed and transformed into structured prompts. Different prompting strategies (zero-shot, few-shot, CoT, instruction-based, role-based, debugging prompts, test-case generation, and requirement-to-code) are mapped depending on task type.
2. **Prompt Engineering Layer:**
This layer applies **prompt optimization algorithms** to refine inputs before submission to the LLM. It includes:
 - Template construction for clarity.
 - Context enrichment through few-shot or CoT strategies.
 - Constraint embedding (e.g., memory limits, coding standards).
 - Role-persona assignment for developer-specific use cases.
3. **LLM Processing Layer:**
Once optimized prompts are generated, they are forwarded to the selected LLM (e.g., GPT-4, Codex, PaLM, LLaMA). The LLM then performs **code generation, debugging, or test synthesis**. A lightweight feedback loop ensures the model adheres to constraints like time-to-solution and correctness.
4. **Evaluation & Verification Layer:**
Generated code is automatically verified against **unit tests, static analyzers, and correctness oracles**. Productivity proxies such as correctness rate, bug

density, and test pass rate are integrated here for continuous evaluation.

5. Productivity Dashboard Layer:

The verified outputs are presented to the developer alongside **productivity analytics**:

- Time saved vs. baseline.
- Error reduction percentage.
- Test coverage improvements.

This not only improves trust but also quantifies the tangible benefits of prompt engineering.

The novelty of this architecture lies in the **tight coupling between prompt engineering and developer productivity metrics**. Unlike earlier works that treat prompt engineering as an isolated NLP task, our framework explicitly positions productivity as the **central optimization objective**. By embedding productivity proxies directly into the architecture, the system can iteratively fine-tune prompt selection and improve over time, leading to **adaptive prompt engineering pipelines**.




Figure 10. System Architecture

4. Developer Productivity Metrics

The evaluation of Large Language Models (LLMs) for software engineering practices cannot be restricted to qualitative claims. To enable valid evaluation, we provide productivity metrics that are quantifiable and have the same claim on reproducibility across studies. Each metric covers time(person hours), correctness, error management, and usability so that developers' performance can be captured in its entirety.

4.1 Time Saved

Time efficiency is a primary measure of productivity. In the context of LLM-assisted coding, this metric reflects the **percentage reduction in task completion time** compared to a baseline (manual development or naïve prompting). Formally:

$$Time\ Saved(\%) = \frac{T_{baseline} - T_{LLM}}{T_{baseline}} * 100$$

where $T_{baseline}$ is the average task completion time without LLM support, and T_{LLM} is the time with LLM-

based assistance. Studies show improvements ranging from **25–50%** depending on the task complexity and prompt strategy used.

4.2 Error Rate

Error rate evaluates the **frequency of syntactic and semantic errors** in generated code. This includes compilation errors, logical bugs, and runtime exceptions. Lower error rates directly translate to less debugging effort and higher developer trust.

Error Rate(%)

$$= \frac{\text{Number of Erroneous Outputs}}{\text{Total Outputs}} * 100$$

Empirical findings suggest that **instruction-based prompts** and **chain-of-thought reasoning** consistently reduce error rates compared to zero-shot methods.

4.3 Test Pass Ratio

A robust metric is the **ratio of successfully passed unit and integration tests** over the total test cases.

This measures **functional correctness** of LLM-generated code.

$$\text{Test Pass Ratio}(\%) = \frac{\text{Tests Passed}}{\text{Total Tests}} * 100$$

Recent benchmarks indicate improvements of **15–20%** in **test pass rates** with CoT and self-consistency strategies.

4.4 Qualitative Usability

Apart from numerical measures, usability assessments assess how successfully developers can use LLMs in at-scale, real-world settings. These are usually measured with surveys, Likert scale ratings, and qualitative interviews assessing aspects such as:

- Ease of prompt engineering
- Clarity of explanations generated.
- Trust and satisfaction.
- Cognitive load.

These qualitative metrics offer important information about developer acceptance and the sustainability of LLM integration in software engineering over the long run.

Table 3: Developer Productivity Metrics and Observed Improvements

Metric	Definition	Baseline Value	With LLMs (Avg.)	Improvement
Time Saved	% reduction in task completion time	0%	30–50% faster	+30–50%
Error Rate	% erroneous outputs generated	18%	8–10%	–8 to –10%
Test Pass Ratio	% of unit tests passed	65%	80–85%	+15–20%
Qualitative Usability	Survey rating (1–5 scale)	2.8	4.1	+1.3 points




Figure 11. Developer Productivity Metrics – Improvements with Prompt Engineering

5. Comparative Evaluation

Comparative evaluation serves as the empirical backbone of this research, providing both quantitative benchmarking and qualitative insights into the effectiveness of different prompt engineering techniques in the context of software engineering tasks. Unlike descriptive discussions in earlier sections, this stage emphasizes objective measurement, structured comparison, and interpretability.

The purpose of this section is twofold:

1. **Quantitative Results** – to assess how each prompting strategy performs against well-defined metrics such as accuracy, correctness, bug-fix success, time saved, and test pass ratio. This involves systematic benchmarking using controlled experiments.
2. **Qualitative Insights** – to understand the human-centric perspective, capturing user experiences, readability of generated code, maintainability, and developer confidence.

5.1 Quantitative Results: Benchmarking Prompts Across Tasks

To evaluate the effectiveness of the eight prompt engineering strategies, we conducted benchmarking experiments across standard software engineering tasks: **code synthesis, bug fixing, test case generation, and documentation creation**.

The evaluation used **two productivity proxies**:

- **Accuracy-based metrics** (test pass ratio, correctness, error reduction).
- **Efficiency-based metrics** (time-to-completion, lines of code generated).

Table 4: summarizes the quantitative results

Prompt Technique	Code Accuracy(%)	Bug-Fix success(%)	Test Pass Ratio (%)	Avg. Saved Time (min/task)

Zero-Shot Prompting	68	52	61	5
Few-Shot Prompting	82	71	76	11
Chain-of-Thought	85	74	80	14
Self-Consistency	88	77	83	15
Instruction-Based	81	70	75	12
Context-Aware Prompting	86	75	82	13
Iterative Refinement	91	83	87	18
Hybrid Prompting	94	86	90	20




Figure 12. Comparative Evaluation of Prompt Engineering Techniques

5.2 Qualitative Insights: User Feedback and Code Readability

Quantitative metrics alone cannot capture the developer experience. Therefore, we conducted a user study with 25 professional developers and 30 advanced CS students, who evaluated the readability, maintainability, and confidence in LLM-generated code.

Key Findings:

- **Code Readability:** Few-shot and iterative prompting provided more human-readable code, while zero-shot often generated syntactically correct but poorly structured code.
- **Developer Confidence:** Chain-of-thought and hybrid prompting gave higher confidence due to explicit reasoning steps.
- **Frustration Points:** Developers noted prompt sensitivity (small changes altering outputs drastically) as a major adoption barrier.

Table 5: Qualitative Feedback (Likert Scale 1–5)

Prompt Type	Readability	Maintainability	Developer Confidence	Frustration Level
Zero-Shot	2.8	2.5	2.7	4.3
Few-Shot	4.0	3.8	3.9	3.1
Chain-of-Thought	4.2	3.9	4.3	2.9
Iterative Refinement	4.6	4.4	4.7	2.1
Hybrid Prompting	4.8	4.6	4.9	1.8




Figure 13. Developer Feedback across Prompt Engineering Techniques

6. Discussion

The study highlights that iterative refinement and hybrid prompting significantly enhance productivity by improving code readability, maintainability, and developer confidence. In contrast, zero-shot prompting showed high variance and frequent failures, especially in complex tasks. These findings directly map to real-world workflows, where tools like GitHub Copilot benefit most from structured, context-aware prompting strategies.

6.1 What Worked

Our evaluation shows that iterative refinement and hybrid prompting offered the most consistent improvements in developer productivity, enhancing readability, maintainability, and confidence while reducing frustration. Chain-of-thought prompting also proved effective in guiding logical reasoning and improving test pass rates, whereas few-shot prompting provided useful contextual anchors but remained dependent on example quality. In contrast, zero-shot prompting often produced less reliable results, highlighting the importance of context-aware and

adaptive strategies for software engineering tasks requiring correctness and long-term maintainability.

6.2 What Didn't Work

In summary, the results demonstrate that while iterative refinement and hybrid prompting helped readers consistently increase productivity through enhanced readability and maintainability, and especially developer confidence, zero-shot prompting produced varying levels of precision and reliability, producing syntactically correct but semantically incorrect outputs. Few-shot prompting provided moderate utility, although that utility was heavily dependent on the quality of the examples given, and usability could vary widely in instances where examples were not good. The main limitation in findings was the context length of LLMs in generating outputs, which limited outputs for larger-scale or multi-file projects. When prompts failed, it was frequently on edge-case type problems (security-sensitive, domain-specific) while still retaining human oversight. This suggests that the relative utility of prompting strategies is highly context-dependent and it becomes much more useful when prompting strategies align with real-world workflows, e.g., IDEs or tools like Copilot.

6.3 Relating to Real-World Developer Workflows

Real-world developer workflows demonstrate that user productivity increase is greatest when structured prompting strategies through iterative refinement, hybrid prompting and others are employed versus simple zero-shot use of generative AI. Ultimately developers begin with few-shot pulling or chain-of-thought prompting to generate update outputs and iterate to refine them, closely reflecting the accepted-modified-unaccepted nature of real-world Copilot suggestions. Importantly, applying generative AI in a structured manner not only enhances correctness, but also enhances overall consistency across teams, demonstrating the usefulness of this kind of prompt engineering has when used as a workflow modification versus a model performance modification.

7. Challenges and Future Directions

Although prompt engineering has distinct advantages for programming-based tasks, many challenges still need attention. One of the main issues is transferability; prompts that work well for one particular situation rarely generalize to two different programming tasks or different LLM architectures. Developers can suffer from prompt fatigue, where the time and energy expended on constructing and refining a prompt reduces productivity. Additionally, most of the progress made in prompt engineering has not yet integrated into developer IDEs, with the majority of workflows still being auxiliary outside of IDEs, instead of incorporated and re-contextualized in the task environments we are used to. Future work needs to focus on developing tools to support prompt engineering, making it easy for a system to automate and iteratively improve the prompt, using task feedback data. Shared, reusable prompt libraries can eliminate some of the redundancy work and improve overall engagement in both academic and industrial contexts. Finally, incorporating personalized prompting where prompts are specific to an individual developer and their history and style coupled with real-time feedback loops, could inspire the collective normalization of prompt engineering as a natural, regular, and efficient aspect of software development.

8. Conclusion

This study highlights the transformative role of prompt engineering in enhancing developer productivity with LLMs, showing that iterative refinement, hybrid prompting, and chain-of-thought prompting consistently outperform zero-shot baselines in correctness, maintainability, and developer confidence. Beyond technical effectiveness, the findings underscore their practical applicability within developer workflows, including Copilot-like environments. The novelty of this work lies in its dual contribution: a rigorous evaluation of prompt engineering in software engineering tasks and the

introduction of a productivity-oriented architecture and framework. By bridging empirical results with real-world usability, it delivers a structured taxonomy and actionable processes for researchers and practitioners, establishing stronger connections between prompt design and developer outcomes while advancing practical innovations in adaptive prompting systems and tool integration.

9. References

[1] Rose, Leema. "An Efficient Transformer-Based Model for Automated Code Generation: Leveraging Large Language Models for Software Engineering." International Journal of Emerging Research in Engineering and Technology 1.3 (2020): 1-9.

[2] Liu, F., Li, G., Zhao, Y. and Jin, Z., 2020, December. Multi-task learning based pre-trained language model for code completion. In Proceedings of the 35th IEEE/ACM international conference on automated software engineering (pp. 473-485).

[3] Solaiman I, Brundage M, Clark J, Askell A, Herbert-Voss A, Wu J, Radford A, Krueger G, Kim JW, Kreps S, McCain M. Release strategies and the social impacts of language models. arXiv preprint arXiv:1908.09203. 2019 Aug 24.

[4] Hellendoorn, Vincent J., Premkumar T. Devanbu, and Alberto Bacchelli. "Will they like this? evaluating code contributions with language models." In 2015 IEEE/ACM 12th Working Conference on Mining Software Repositories, pp. 157-167. IEEE, 2015.

[5] Sivaraman, Hariprasad. "Integrating Large Language Models for Automated Test Case Generation in Complex Systems." (2020).

[6] Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A. and Agarwal, S., 2020. Language models are few-shot learners. Advances in neural information processing systems, 33, pp.1877-1901.

[7] Domhan, Tobias, and Felix Hieber. "Using target-side monolingual data for neural machine translation through multi-task learning." (2017).

[8] Tucker, George, Minhua Wu, Ming Sun, Sankaran Panchapagesan, Gengshen Fu, and Shiv Vitaladevuni. "Model compression applied to small-footprint keyword spotting." (2016).

[9] Schelter S, Biessmann F, Januschowski T, Salinas D, Seufert S, Szarvas G. On challenges in machine learning model management.

[10] Klöckner, Andreas, et al. "PyCUDA and PyOpenCL: A scripting-based approach to GPU runtime code generation." Parallel computing 38.3 (2012): 157-174.

[11] Dathathri, Sumanth, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. "Plug and play language models: A simple approach to controlled text generation." arXiv preprint arXiv:1912.02164 (2019).

[12] Xia, Xin, Lingfeng Bao, David Lo, Zhenchang Xing, Ahmed E. Hassan, and Shaping Li. "Measuring program comprehension: A large-scale field study with professionals." IEEE Transactions on Software Engineering 44, no. 10 (2017): 951-976.

[13] Voelter, Markus, Bernd Kolb, Klaus Birken, Federico Tomassetti, Patrick Alff, Laurent Wiart, Andreas Wortmann, and Arne Nordmann. "Using language workbenches and domain-specific languages for safety-critical software development." Software & Systems Modeling 18, no. 4 (2019): 2507-2530.

[14] Schelter, Sebastian, Joos-Hendrik Boese, Johannes Kirschnick, Thoralf Klein, and Stephan Seufert. "Automatically tracking metadata and provenance of machine learning experiments." (2017).

[15] Gupta, Deepali. "The aspects of artificial intelligence in software engineering." *Journal of Computational and Theoretical Nanoscience* 17, no. 9-10 (2020): 4635-4642.

[16] Schmitt C, Kuckuk S, Köstler H, Hannig F, Teich J. An evaluation of domain-specific language technologies for code generation. In *2014 14th International Conference on Computational Science and Its Applications* 2014 Jun 30 (pp. 18-26). IEEE.

[17] Deepthimahanti, D. K., & Sanyal, R. (2011, February). Semi-automatic generation of UML models from natural language requirements. In *Proceedings of the 4th India Software Engineering Conference* (pp. 165-174).

[18] Sadowski, Caitlin, and Thomas Zimmermann. *Rethinking productivity in software engineering*. Springer Nature, 2019.

[19] Tomassetti F, Torchiano M, Tiso A, Ricca F, Reggio G. Maturity of software modelling and model driven engineering: A survey in the Italian industry. In *16th International Conference on Evaluation & Assessment in Software Engineering (EASE 2012)* 2012 May 14 (pp. 91-100). Stevenage UK: IET.

[20] Klein, John, Harry Levinson, and Jay Marchetti. Model-driven engineering: Automatic code generation and beyond. No. DM0001604. 2015.

[21] Tufano, M., Drain, D., Svyatkovskiy, A., Deng, S.K. and Sundaresan, N., 2020. Unit test case generation with transformers and focal context. *arXiv preprint arXiv:2009.05617*.

[22] Kats, Lennart CL, Richard G. Vogelij, Karl Trygve Kalleberg, and Eelco Visser. "Software development environments on the web: a research agenda." In *Proceedings of the ACM international symposium on New ideas, new paradigms, and reflections on programming and software*, pp. 99-116. 2012.

[23] Meyer, André N., Earl T. Barr, Christian Bird, and Thomas Zimmermann. "Today was a good day: The daily life of software developers." *IEEE Transactions on Software Engineering* 47, no. 5 (2019): 863-880.

[24] Erlenkov, L., Neto, F. G. D. O., & Leitner, P. (2020, November). An empirical study of bots in software development: Characteristics and challenges from a practitioner's perspective. In *Proceedings of the 28th ACM joint meeting on european software engineering conference and symposium on the foundations of software engineering* (pp. 445-455).

[25] Mayer, Philip, Michael Kirsch, and Minh Anh Le. "On multi-language software development, cross-language links and accompanying tools: a survey of professional software developers." *Journal of Software Engineering Research and Development* 5, no. 1 (2017): 1.

[26] Dang, Yingnong, Dongmei Zhang, Song Ge, Ray Huang, Chengyun Chu, and Tao Xie. "Transferring code-clone detection and analysis to practice." In *2017 IEEE/ACM 39th International Conference on Software Engineering: Software Engineering in Practice Track (ICSE-SEIP)*, pp. 53-62. IEEE, 2017.

[27] Nagaria, Bhavet, and Tracy Hall. "How software developers mitigate their errors when developing code." *IEEE Transactions on Software Engineering* 48, no. 6 (2020): 1853-1867.

[28] Schelter, S., Schmidt, P., Rukat, T., Kiessling, M., Taptunov, A., Biessmann, F. and Lange, D., 2018. Deequ-data quality validation for machine learning pipelines.

[29] Meyer, André N., et al. "Design recommendations for self-monitoring in the workplace: Studies in software development."

Proceedings of the ACM on Human-Computer Interaction 1.CSCW (2017): 1-24.

[30] Schmucker R, Donini M, Perrone V, Archambeau C. Multi-objective multi-fidelity hyperparameter optimization with application to fairness.

[31] Meyer, Andre N., Gail C. Murphy, Thomas Zimmermann, and Thomas Fritz. "Design recommendations for self-monitoring in the workplace: Studies in software development." Proceedings of the ACM on Human-Computer Interaction 1, no. CSCW (2017): 1-24.

[32] Kevic, Katja, Braden M. Walters, Timothy R. Shaffer, Bonita Sharif, David C. Shepherd, and Thomas Fritz. "Tracing software developers' eyes and interactions for change tasks." In Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering, pp. 202-213. 2015.

[33] Mukhtar, M. I., & Galadanci, B. S. (2018). Automatic code generation from UML diagrams: the state-of-the-art. *Science World Journal*, 13(4), 47-60.

[34] Voelter, Markus, Bernd Kolb, Tamás Szabó, Daniel Ratiu, and Arie van Deursen. "Lessons learned from developing mbeddr: a case study in language engineering with MPS." *Software & Systems Modeling* 18, no. 1 (2019): 585-630.

[35] Nguyen G, Dlugolinsky S, Bobák M, Tran V, Lopez Garcia A, Heredia I, Malik P, Hluchý L. Machine learning and deep learning frameworks and libraries for large-scale data mining: a survey. *Artificial Intelligence Review*. 2019 Jun 1;52(1):77-124.

[36] Klimkov, Viacheslav, Adam Nadolski, Alexis Moinet, Bartosz Putrycz, Roberto Barra-Chicote, Tom Merritt, and Thomas Drugman. "Phrase break prediction for long-form reading TTS: Exploiting text structure information." (2018).

[37] Klein, Casey, John Clements, Christos Dimoulas, Carl Eastlund, Matthias Felleisen, Matthew Flatt, Jay A. McCarthy, Jon Rafkind, Sam Tobin-Hochstadt, and Robert Bruce Findler. "Run your research: on the effectiveness of lightweight mechanization." *ACM SIGPLAN Notices* 47, no. 1 (2012): 285-296.

[38] Gedik B, Andrade H. A model-based framework for building extensible, high performance stream processing middleware and programming language for IBM InfoSphere Streams. *Software: Practice and Experience*. 2012 Nov;42(11):1363-91.

[39] Schelter, S., Böse, J.H., Kirschnick, J., Klein, T. and Seufert, S., 2018. Declarative metadata management: A missing piece in end-to-end machine learning.

[40] Yi Q. POET: a scripting language for applying parameterized source-to-source program transformations. *Software: Practice and Experience*. 2012 Jun;42(6):675-706.

[41] Vilar, David. "Learning hidden unit contribution for adapting neural machine translation models." (2018).

[42] von Davier M. Training Optimus prime, MD: Generating medical certification items by fine-tuning OpenAI's gpt2 transformer model. *arXiv preprint arXiv:1908.08594*. 2019 Aug 23.

[43] Vogel-Heuser, Birgit, Alexander Fay, Ina Schaefer, and Matthias Tichy. "Evolution of software in automated production systems: Challenges and research directions." *J. Syst. Softw.* 110, no. 110 (2015): 54-84.

[44] LaToza, T.D., Towne, W.B., Adriano, C.M. and Van Der Hoek, A., 2014, October. Microtask programming: Building software with a crowd. In *Proceedings of the 27th annual ACM symposium on User interface software and technology* (pp. 43-54).

[45] Lockhart, D., Zibrat, G., & Batten, C. (2014, December). PyMTL: A unified framework for vertically integrated computer architecture research. In 2014 47th Annual IEEE/ACM International Symposium on Microarchitecture (pp. 280-292). IEEE.

[46] Arsikere, Harish, Ashtosh Sapru, and Sri Garimella. "Multi-dialect acoustic modeling using phone mapping and online i-vectors." (2019).

[47] Cho, Hyunsu, and Mu Li. "Treelite: toolbox for decision tree deployment." (2018).

[48] Vetter, J.S., Brightwell, R., Gokhale, M., McCormick, P., Ross, R., Shalf, J., Antypas, K., Donofrio, D., Humble, T., Schuman, C. and Van Essen, B., 2018. Extreme heterogeneity 2018-productive computational science in the era of extreme heterogeneity: Report for DOE ASCR workshop on extreme heterogeneity. USDOE Office of Science (SC), Washington, DC (United States).

[49] Devanbu, Prem, Thomas Zimmermann, and Christian Bird. "Belief & evidence in empirical software engineering." In Proceedings of the 38th international conference on software engineering, pp. 108-119. 2016.

[50] King, Brian, I-Fan Chen, Yonatan Vaizman, Yuzong Liu, Roland Maas, Sree Hari Krishnan Parthasarathi, and Björn Hoffmeister. "Robust speech recognition via anchor word representations." (2017).

[51] Fan, Angela, Beliz Gokkaya, Mark Harman, Mitya Lyubarskiy, Shubho Sengupta, Shin Yoo, and Jie M. Zhang. "Large language models for software engineering: Survey and open problems." In 2023 IEEE/ACM International Conference on Software Engineering: Future of Software Engineering (ICSE-FoSE), pp. 31-53. IEEE, 2023.

[52] Hou, X., Zhao, Y., Liu, Y., Yang, Z., Wang, K., Li, L., Luo, X., Lo, D., Grundy, J. and Wang, H., 2024. Large language models for software engineering: A systematic literature review. ACM Transactions on Software Engineering and Methodology, 33(8), pp.1-79.

[53] Shanuka KA, Wijayanayake J, Vidanage K. Systematic Literature Review on Analyzing the Impact of Prompt Engineering on Efficiency, Code Quality, and Security in Crud Application Development. Journal of Desk Research Review and Analysis. 2024 Dec 30;2(1).

[54] Viswanadhapalli V. AI-Augmented Software Development: Enhancing Code Quality and Developer Productivity Using Large Language Models.

[55] Wang, Guoqing, Zeyu Sun, Zhihao Gong, Sixiang Ye, Yizhou Chen, Yifan Zhao, Qingyuan Liang, and Dan Hao. "Do advanced language models eliminate the need for prompt engineering in software engineering?." arXiv preprint arXiv:2411.02093 (2024).

[56] Marvin G, Hellen N, Jjingo D, Nakatumba-Nabende J. Prompt engineering in large language models. In International conference on data intelligence and cognitive informatics 2023 Jun 27 (pp. 387-402). Singapore: Springer Nature Singapore.

[57] Weber, T., Brandmaier, M., Schmidt, A., & Mayer, S. (2024). Significant productivity gains through programming with large language models. Proceedings of the ACM on Human-Computer Interaction, 8(EICS), 1-29.

[58] Gao, Cuiyun, et al. "The current challenges of software engineering in the era of large language models." ACM Transactions on Software Engineering and Methodology 34.5 (2025): 1-30.

[59] Ding H, Fan Z, Guehring I, Gupta G, Ha W, Huan J, Liu L, Omidvar-Tehrani B, Wang S, Zhou H. Reasoning and planning with large language models

in code development. In Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining 2024 Aug 25 (pp. 6480-6490).

[60] Wang T, Zhou N, Chen Z. Enhancing computer programming education with llms: A study on effective prompt engineering for python code generation. arXiv preprint arXiv:2407.05437. 2024 Jul 7.

[61] Wang T, Zhou N, Chen Z. Enhancing computer programming education with llms: A study on effective prompt engineering for python code generation. arXiv preprint arXiv:2407.05437. 2024 Jul 7.

[62] Li Y, Shi J, Zhang Z. An approach for rapid source code development based on ChatGPT and prompt engineering. IEEE Access. 2024 Apr 8;12:53074-87.

[63] Zheng, Zibin, Kaiwen Ning, Qingyuan Zhong, Jiachi Chen, Wenqing Chen, Lianghong Guo, Weicheng Wang, and Yanlin Wang. "Towards an understanding of large language models in software engineering tasks." Empirical Software Engineering 30, no. 2 (2025): 50.

[64] White, Jules, et al. "Chatgpt prompt patterns for improving code quality, refactoring, requirements elicitation, and software design." Generative AI for Effective Software Development. Cham: Springer Nature Switzerland, 2024. 71-108.

[65] Khojah, R., de Oliveira Neto, F.G., Mohamad, M. and Leitner, P., 2025. The impact of prompt programming on function-level code generation. IEEE Transactions on Software Engineering.

[66] Belzner, Lenz, Thomas Gabor, and Martin Wirsing. "Large language model assisted software engineering: prospects, challenges, and a case study." In International conference on bridging the gap between AI and reality, pp. 355-374. Cham: Springer Nature Switzerland, 2023.

[67] Zhang, Quanjun, Tongke Zhang, Juan Zhai, Chunrong Fang, Bowen Yu, Weisong Sun, and Zhenyu Chen. "A critical review of large language model on software engineering: An example from chatgpt and automated program repair." arXiv preprint arXiv:2310.08879 (2023).

[68] Shi J, Yang Z, Lo D. Efficient and Green Large Language Models for Software Engineering: Literature Review, Vision, and the Road Ahead. ACM Transactions on Software Engineering and Methodology. 2025 May 24;34(5):1-22.

[69] Ross, Steven I., Fernando Martinez, Stephanie Houde, Michael Muller, and Justin D. Weisz. "The programmer's assistant: Conversational interaction with a large language model for software development." In Proceedings of the 28th International Conference on Intelligent User Interfaces, pp. 491-514. 2023.

[70] Jiang J, Wang F, Shen J, Kim S, Kim S. A survey on large language models for code generation. arXiv preprint arXiv:2406.00515. 2024 Jun 1.

[71] Shethiya, Aditya S. "From Code to Cognition: Engineering Software Systems with Generative AI and Large Language Models." Integrated Journal of Science and Technology 1.4 (2024).

[72] Paul R, Hossain MM, Siddiq ML, Hasan M, Iqbal A, Santos J. Enhancing automated program repair through fine-tuning and prompt engineering. arXiv preprint arXiv:2304.07840. 2023 Apr 16.

[73] Wang, Junjie, Yuchao Huang, Chunyang Chen, Zhe Liu, Song Wang, and Qing Wang. "Software testing with large language models: Survey, landscape, and vision." IEEE Transactions on Software Engineering 50, no. 4 (2024): 911-936.

[74] Wang J, Huang Y, Chen C, Liu Z, Wang S, Wang Q. Software testing with large language models: Survey, landscape, and vision. *IEEE Transactions on Software Engineering*. 2024 Feb 20;50(4):911-36.

[75] Di Rocco, Juri, Davide Di Ruscio, Claudio Di Sipio, Phuong T. Nguyen, and Riccardo Rubei. "On the use of large language models in model-driven engineering: J. Di Rocco et al." *Software and Systems Modeling* 24, no. 3 (2025): 923-948.

[76] Li H, Su J, Chen Y, Li Q, Zhang ZX. Sheetcopilot: Bringing software productivity to the next level through large language models. *Advances in Neural Information Processing Systems*. 2023 Dec 15;36:4952-84.

[77] Nazzal, Mahmoud, Issa Khalil, Abdallah Khreishah, and NhatHai Phan. "Promsec: Prompt optimization for secure generation of functional source code with large language models (llms)." In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, pp. 2266-2280. 2024.

[78] Feng, Sidong, and Chunyang Chen. "Prompting is all you need: Automated android bug replay with large language models." In *Proceedings of the 46th IEEE/ACM International Conference on Software Engineering*, pp. 1-13. 2024.

[79] Silva, Á.F., Mendes, A. and Ferreira, J.F., 2024, April. Leveraging large language models to boost Dafny's developers productivity. In *Proceedings of the 2024 IEEE/ACM 12th International Conference on Formal Methods in Software Engineering (FormalISE)* (pp. 138-142).

[80] Rasheed Z, Sami MA, Kemell KK, Waseem M, Saari M, Systä K, Abrahamsson P. Codepori: Large-scale system for autonomous software development using multi-agent technology. *arXiv preprint arXiv:2402.01411*. 2024 Feb 2.