

# Metadata of the chapter that will be visualized in SpringerLink

Book Title	New Trends in Disruptive Technologies, Tech Ethics and Artificial Intelligence	
Series Title		
Chapter Title	Applying Large Language Models to Software Development: Enhancing Requirements, Design and Code	
Copyright Year	2025	
Copyright HolderName	The Author(s), under exclusive license to Springer Nature Switzerland AG	
Author	Family Name	<b>Santos</b>
	Particle	
	Given Name	<b>Gonçalo</b>
	Prefix	
	Suffix	
	Role	
	Division	
	Organization	Instituto Politécnico da Guarda
	Address	Guarda, Portugal
	Email	1700800@sal.ipg.pt
	ORCID	<a href="http://orcid.org/0009-0007-9322-9582">http://orcid.org/0009-0007-9322-9582</a>
Corresponding Author	Family Name	<b>Silveira</b>
	Particle	
	Given Name	<b>Clara</b>
	Prefix	
	Suffix	
	Role	
	Division	
	Organization	Instituto Politécnico da Guarda
	Address	Guarda, Portugal
	Email	mclara@ipg.pt
	ORCID	<a href="http://orcid.org/0000-0003-2809-4208">http://orcid.org/0000-0003-2809-4208</a>
Author	Family Name	<b>Santos</b>
	Particle	
	Given Name	<b>Vitor</b>
	Prefix	
	Suffix	
	Role	
	Division	
	Organization	Nova IMS, Universidade Nova de Lisboa
	Address	Lisbon, Portugal
	Email	vsantos@novaims.unl.pt
	ORCID	<a href="http://orcid.org/0000-0002-4223-7079">http://orcid.org/0000-0002-4223-7079</a>
Author	Family Name	<b>Santos</b>
	Particle	

Given Name **Arnaldo**  
Prefix  
Suffix  
Role  
Division  
Organization Universidade Aberta, INESC TEC  
Address Lisbon, Portugal  
Email arnaldo.santos@uab.pt  
ORCID <http://orcid.org/0000-0001-5139-6728>

---

Author

Family Name **Mamede**  
Particle  
Given Name **Henrique**  
Prefix  
Suffix  
Role  
Division  
Organization Universidade Aberta, INESC TEC  
Address Lisbon, Portugal  
Email jose.mamede@uab.pt  
ORCID <http://orcid.org/0000-0002-5383-9884>

---

Abstract

This paper explores the potential of Large Language Models (LLM) to optimize various stages of the software development lifecycle, including requirements elicitation, architecture design, diagram creation, and implementation. The study is grounded in a real-world case, where development time and result quality are compared with and without LLM assistance. This research underscores the possibility of applying prompt patterns in LLM to support and enhance software development activities, focusing on a B2C digital commerce platform centered on fashion retail, designated LUNA. The methodology adopted is Design Science, which follows a practical and iterative approach. Requirements, design suggestions, and code samples are analyzed before and after the application of language models. The results indicate substantial advantages in the development process, such as improved task efficiency, faster identification of requirement gaps, and enhanced code readability. Nevertheless, challenges were observed in interpreting complex business logic. Future work should explore the integration of LLM with domain-specific ontologies and business rule engines to improve contextual accuracy in code and model generation. Additionally, refining prompt engineering strategies and combining LLM with interactive development environments could further enhance code quality, traceability, and explainability.

---

Keywords  
(separated by '-')

Large Language Models - Prompt Engineering - Software Development Lifecycle

---



# Applying Large Language Models to Software Development: Enhancing Requirements, Design and Code

Gonçalo Santos<sup>1</sup> , Clara Silveira<sup>1</sup>  (✉), Vitor Santos<sup>2</sup> , Arnaldo Santos<sup>3</sup> ,  
and Henrique Mamede<sup>3</sup> 

<sup>1</sup> Instituto Politécnico da Guarda, Guarda, Portugal  
1700800@sal.ipg.pt, mclara@ipg.pt

<sup>2</sup> Nova IMS, Universidade Nova de Lisboa, Lisbon, Portugal  
vsantos@novaims.unl.pt

<sup>3</sup> Universidade Aberta, INESC TEC, Lisbon, Portugal  
{arnaldo.santos, jose.mamede}@uab.pt



**Abstract.** This paper explores the potential of Large Language Models (LLM) to optimize various stages of the software development lifecycle, including requirements elicitation, architecture design, diagram creation, and implementation. The study is grounded in a real-world case, where development time and result quality are compared with and without LLM assistance. This research underscores the possibility of applying prompt patterns in LLM to support and enhance software development activities, focusing on a B2C digital commerce platform centered on fashion retail, designated LUNA. The methodology adopted is Design Science, which follows a practical and iterative approach. Requirements, design suggestions, and code samples are analyzed before and after the application of language models. The results indicate substantial advantages in the development process, such as improved task efficiency, faster identification of requirement gaps, and enhanced code readability. Nevertheless, challenges were observed in interpreting complex business logic. Future work should explore the integration of LLM with domain-specific ontologies and business rule engines to improve contextual accuracy in code and model generation. Additionally, refining prompt engineering strategies and combining LLM with interactive development environments could further enhance code quality, traceability, and explainability.

**Keywords:** Large Language Models · Prompt Engineering · Software Development Lifecycle

## 1 Introduction

As Large Language Models (LLM) continue to advance, researchers are increasingly exploring their capabilities across specialized domains, including software engineering. These models have demonstrated significant effectiveness in tasks such as code generation and identifying software vulnerabilities [1]. While LLM have demonstrated

potential in basic programming tasks, their role in more complex tasks requiring semantic understanding and contextual reasoning remains less explored. This research examines how LLM can support different stages of the software development lifecycle, from requirements engineering to design and implementation. This research uses the LUNA (a previously developed B2C digital commerce platform focused on fashion retail) as a case study. It compares outputs produced with LLM assistance to those without, focusing on quality, standards compliance, semantic accuracy, and development time. The main research problem is whether LLM can meaningfully contribute beyond code generation, especially in critical activities like requirements elicitation and architectural modeling. This work weaves into an expanding corpus body of research on the use of LLM in software engineering. Beyond development acceleration, this research investigates how LLM can maintain or improve the quality and precision of software artifacts. The research gives handy insight into the application of LLM in complex software projects, particularly addressing gaps in semantic modeling and code quality assessment.

## 2 Related Work

The concept of pattern is often referenced in the works of Alexander [2] who collected and documented patterns in the field of architecture. A solution pattern is a generic solution to a specific problem in a way that it can be repeatedly applied in different contexts [3]. The idea of patterns was extrapolated to software, with patterns being developed for object-oriented modeling [4,5], for object-oriented design and programming [6], for exploring techniques, strategies, and applications [7] and for writing software requirements [8], among many other studies.

Since 2020, prompt patterns have become a vital area of research for the evolution of generative AI systems. As models grow in complexity and capability, increased attention has been directed toward how prompts can be designed, refined, and optimized. Prompt patterns are analogous to software patterns, offering structured techniques to improve interaction with LLM [9]. Indeed, as LLM like ChatGPT become an integral part of software engineering, the need for prompt engineering emerges, similar to the development of software design patterns. This transition underscores the importance of systematic approaches to enhancing the effectiveness of prompts in software development. Recent research continues to explore new strategies to enhance human-machine interaction. Studies suggest that the use of structured prompt patterns with models such as ChatGPT can lead to a 41% increase in software development efficiency compared to traditional methods [10].

The taxonomy of prompt design dimensions proposed by Braun et al. [11] offers a valuable framework for future research on how prompt construction influences user interaction with AI systems. Similarly, the AutoDev framework introduced by Tufano et al. [12] presents an AI-driven structure that automates software engineering tasks, enhancing both efficiency and security. Although the study does not explicitly analyze ChatGPT prompt patterns, AutoDev's capabilities simplify complex development tasks and improve productivity in real-world software engineering applications. Empirical studies applying LLM in software engineering environments have produced important lessons. Schmidt et al. [13] emphasize the need to move beyond ad hoc prompting practices, advocating for more systematic and reusable approaches. Prompt patterns serve as

foundational components of prompt engineering, providing proven solutions to recurring problems and accelerating task resolution across multiple phases of the software development lifecycle. Future work in this field of research should aim to enhance the robustness of these models, reduce hallucinations, and improve their handling of complex modeling scenarios [13].

From Kim’s perspective [14], prompt-based software engineering streamlines development by automating repetitive tasks and generating high-quality outputs. Liang et al. [15] take this idea further, arguing that certain types of prompts can be considered programs in themselves. They propose that prompt development constitutes a distinct phenomenon in programming, coining the term prompt programming to describe this emerging paradigm. Patil & Gudivada [16] present a comprehensive evaluation of LLM performance in tasks such as code generation, refactoring, and debugging. Their findings underscore the strengths of tools like OpenAI’s ChatGPT and GitHub Copilot, particularly in producing syntactically correct code with relatively low error rates. Nonetheless, the study highlights several limitations, including hallucinations, limited semantic depth, and challenges in maintaining context throughout extended interactions. In the realm of software modeling, Cámara et al. [17] investigate the use of ChatGPT for generating Unified Modeling Language (UML) class diagrams with Object Constraint Language (OCL). While results show that LLM are proficient in tasks with structural similarity to programming languages—such as OCL expressions, the models struggle with complex modeling constructs, including association classes and multiple inheritance. Moreover, variability in responses to identical prompts raises concerns regarding consistency and repeatability—critical factors in model-based development. The quality of LLM-assisted output is also affected by factors such as prompt formulation and model configuration. While research such as [8] emphasizes the role of hyperparameter tuning in optimizing model behavior, tools like ChatGPT do not provide access to such parameters, instead relying on prompt design to guide output behavior. In summary, LLM hold substantial promises in software engineering tasks, particularly in code-centric activities [16]. However, their application to software modeling still presents limitations in reliability, structural consistency, and semantic precision [17].

Despite LLM advances in automating software development tasks, including code synthesis, program repair and test generation [19], it appears that comparative analysis in relation to traditional development methods remains insufficiently studied.

### 3 Methodology

The methodology used in this work – Design Science – is a research methodology suited to the fields of Information Systems and Software Engineering for the creation of an artifact, Design Science is the design and investigation of artifacts [18], serving as the theoretical foundation supporting the scientific validity of the work. This approach was selected because it enables systematic evaluation through iterative development cycles while maintaining practical relevance through continuous evaluation. Specifically, the study seeks to answer the following research questions:

1. What is the impact of LLM on the efficiency and quality of requirements elicitation, design, and coding tasks?

2. How does LLM-generated UML compare to human-created diagrams in terms of syntax, semantics, and practical value?
3. What limitations emerge when LLM are applied to domain-specific logic and complex software modeling?



**Fig. 1.** Wieringa Engineering activities

Figure 1, presents the activities underlying the Wieringa Engineering cycle [18] involved in the artifact creation process. In the problem investigation activity, the needs related to the research questions under study were identified; treatment design: interaction between artifact and context; treatment validation allows identification of changes in the artifact and/or context; treatment implementation: transfer to practice; implementation evaluation involves checking whether the designed artifacts support the initial assumptions.

The manual LUNA software development process, serving as the baseline, adopted a traditional agile Scrum methodology. It involved iterative development cycles with manual coding, debugging, and deployment of components such as a web-based management platform (using React.js for frontend and Node.js with Express for backend) and a mobile commerce application (using Flutter). Each feature was manually planned, coded, reviewed, and tested through conventional user acceptance tests and functional usability tests. This manual approach relied heavily on direct developer input for coding standards adherence and semantic accuracy, as detailed in previous internal documentation and iterative sprint retrospectives.

To systematically define qualitative aspects such as ‘enhanced code readability and semantic accuracy’, we implemented a structured evaluation protocol. Code readability was assessed using automated metrics (e.g., SonarQube scores) complemented by peer review checklists ensuring consistency. Semantic accuracy in UML modelling was evaluated through a defined rubric where evaluators systematically categorized diagrams according to explicit criteria, such as the specificity of model elements, clarity of relationships, and risk of semantic misunderstanding. Inter-rater reliability was ensured through calibration sessions among evaluators, mitigating subjectivity risks.

## 4 Comparison with the LLM-Assisted Approach

This section provides benchmark data on the effort and accuracy of the various elements of the LUNA software development process, allowing a direct comparison between manual work and work assisted by LLM. The manual LUNA software development process

served as a basis for comparison with the LLM-assisted approach, covering requirements gathering, architectural design, UML modeling and coding activities. These activities followed traditional practices based on expert knowledge, standard design patterns, and iterative validation through peer review. The source code was written from scratch, and testing was fully manual.

#### 4.1 Requirements Engineering

Requirements are the base of all software products, and consequently, requirements engineering plays an important role in system development [20]. To explore the application of LLM in requirements elicitation, a tailored prompt was developed based on IEEE 830–1998 requirements specification guidelines and the Software Engineering Body of Knowledge requirements engineering practices. The prompt design incorporated key principles of effective prompt engineering [10]: Contextualization: Clearly defined project scope, including mobile app and back-office functionalities. Explicit Intent: Clear objective to identify and categorize functional and non-functional requirements. Constraints: Responses limited to relevant requirements, categorized as mandatory or optional. Specificity: Targeted questions on security, scalability, usability, and social media integration. Logical Structure: Use of numbered lists and bullet points to enhance analysis and coherence. The prompt was applied to the LUNA platform. The output, Fig. 2, was systematically structured into three major categories: Functional Requirements include user authentication, sales processing, inventory control, and social media integration. Non-Functional Requirements: covering security, performance, scalability, and usability—ensuring compliance with industry standards and good engineering practices. Behavioral Insights: Requirements tied to social media data collection and analysis, enabling personalized marketing strategies and behavioral analytics.

The model's response not only organized these elements coherently but also validated the technical and strategic feasibility of the proposed system.

The use of LLM significantly reduced the time required for requirements gathering and structuring, Table 1. In a controlled study, empirical results show that the traditional approach required approximately 7 h, while the AI-assisted method completed the task in 3 h, representing a 42.86% reduction in time. This improvement is due to the model's ability to rapidly synthesize information, detect patterns, and provide organized outputs. Moreover, the structured prompt strategy ensured that the results were relevant and aligned with the project scope.

In addition to time savings, the AI-assisted process enhanced the quality of the requirement specification by reducing ambiguity and improving precision. These findings highlight the potential of LLM as effective tools for critical software engineering processes and open avenues for further research into AI-assisted requirements engineering.

#### 4.2 Architecture

The widely used three-layer architectural model served as the foundation for the original LUNA system design. It was organized to distinguish between the data layer (PostgreSQL), the application layer (Node.js with Express), and the presentation layer, which

Category	Sub-Category	Examples
Functional Requirement – Mobile App	User Management	Registration, Login, Profile Editing
Functional Requirement – Mobile App	Product Catalog & Search	Browsing, Filtering, Product Details
Functional Requirement – Mobile App	Shopping Cart & Checkout	Cart Operations, Payment Methods, Order Confirmation
Functional Requirement – Commercial Platform	Inventory & Order Management	Stock Management, Order Tracking, Sales Dashboard
Functional Requirement – Commercial Platform	Social Media Integration	Social Login, Behavior Data Collection, Campaign Tracking
Non-Functional Requirement	Security	GDPR, Encryption, Secure Payments
Non-Functional Requirement	Performance	Fast Response Time, API Performance, High Availability
Non-Functional Requirement	Scalability	Auto-scaling, Cloud Infrastructure
Non-Functional Requirement	Usability	Responsive Design, Accessibility, Intuitive Navigation
Social Media Integration Requirement	Data Analytics	Sentiment Analysis, Influencer Detection
Social Media Integration Requirement	Marketing Automation	Campaign Scheduling, Performance Analytics
Social Media Integration Requirement	User Behavior Insights	Trends Tracking, Personalized Recommendations

**Fig. 2.** Functional and Non-Functional Requirements identified, generated by ChatGPT.

**Table 1.** Comparison of time spent on requirements elicitation using traditional methods versus the AI-assisted approach.

Approach	Time Required (hours)	Time Saved	% Improvements	Qualitative Notes
Manual (Traditional)	7	—	—	Standard elicitation with interviews
AI-Assisted	3	4h	42.86%	Structure, low ambiguity, faster synthesis

comprised a web platform (React.js) and a mobile application (Flutter). In accordance with established practices for software engineering, this strategy sought to guarantee modularity, scalability, and maintainability. But as AI-assisted development techniques were investigated, it became clear that including LLM might inevitably result in a unique architectural vision, especially when it came to automation, modular code creation, and more infrastructure flexibility, Fig. 3.

The layered structure is maintained while major alterations are introduced in the architectural concept covered in this study, which is based on AI-driven development processes. The substitution of React.js for Next.js in the web front-end was one of the primary changes. This decision was driven by Next.js's support for server-side rendering and API routes, which complement the "API-first" methodology promoted by LLM-assisted code generation. The substitution of Express for NextJS in the application layer was another significant change. Express was adequate for the manual development stage,

Layer	Technologies / Tools	Description
Presentation Layer (Web)	Next.js (React Framework)	API-first web front-end with SSR and SSG support, optimized for scalability and performance
Presentation Layer (Mobile)	Flutter	Cross-platform development framework for mobile applications
Application Layer	Node.js with NestJS	Modular backend with dependency injection, promoting clean architecture and scalability
AI Layer	OpenAI API, LangChain	Dedicated layer for prompt management, code generation, automated testing and documentation
Database Layer	PostgreSQL (Cloud RDS)	Robust relational database for structured data storage
Infrastructure	AWS Lambda, API Gateway	Serverless architecture to ensure scalability, cost-efficiency, and event-driven execution
DevOps & Quality	GitHub Actions, SonarQube	Automated CI/CD pipeline with integrated code quality analysis
Monitoring & Observability	Sentry, Prometheus	Real-time error tracking, system monitoring, and performance metrics
API Management	AWS API Gateway or Kong	API lifecycle management, security, and analytics
LLM Assistants	GitHub Copilot, ChatGPT API	AI-assisted code completion, documentation generation, testing support

**Fig. 3.** AI-Assisted architecture proposed for LUNA

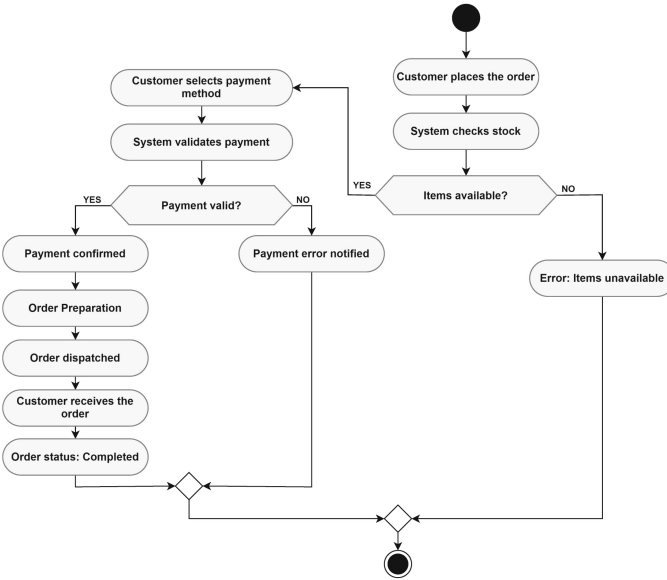
but NestJS provides more organized, flexible, and scalable framework—features that are especially crucial when some code is generated or optimized with AI support.

Furthermore, a specialized AI layer is incorporated into the new design, which was absent from the original version. This layer was created to centralize prompt management, automatic code scaffolding, test creation, and documentation—elements that naturally arose during the AI model experimentation—in addition to consuming LLM APIs. The specialized AI layer introduced in the LUNA architecture acts as an intermediary facilitating interaction between frontend/backend layers and external AI services. This layer includes clearly defined API endpoints for task-specific AI interactions (e.g., code generation, model evaluation), robust error handling mechanisms, and an auditing module for monitoring AI-generated outputs. Detailed architectural diagrams and implementation guidelines for this AI layer have been documented to assist other practitioners aiming to integrate similar AI capabilities into their software systems. The shift to a serverless approach based on AWS Lambda and API Gateway is a logical progression from an infrastructure standpoint, emphasizing scalability and operational simplicity, particularly in situations where development speed is crucial. Although these architectural changes seek to utilize LLM-enhanced development features, it is recognized that every replacement like switching to Next.js or NestJS brings trade-offs regarding learning curve, maturity of the ecosystem, and complexity of integration, warranting additional empirical assessment.

### 4.3 UML Modeling

UML modeling plays a foundational role in visualizing and validating software structure prior to implementation. Traditionally, this task is handled by experienced engineers who manually derive diagrams such as class, use case, sequence, and component diagrams from requirements. With the emergence of LLM, we explored AI-assisted diagram generation by crafting detailed prompts and using the PlantUML tool to convert the LLM-generated textual descriptions into visual diagrams. PlantUML was selected due

to its open-source nature and compatibility with textual UML generation across multiple diagram types (<https://plantuml.com/>). Prompts were designed to describe system architecture in detail, requesting UML outputs from models such as ChatGPT. The generated text was then processed by PlantUML to produce diagrams, Fig. 4. The process revealed several critical limitations, including syntax issues, lack of deep contextualization, structural weaknesses, and static outputs. In contrast, manual modeling demonstrated significant advantages, such as higher accuracy, better structure, and iterative refinement.



**Fig. 4.** Automatically generated flowchart from LLM prompt.

These findings reinforce that while LLM are useful in kickstarting UML modeling, final models still require human judgement, critical thinking, and architectural expertise to ensure robustness and accuracy. To consolidate these observations, the Table 2 presents a comparative summary between manual and AI-assisted UML modeling.

To provide a structured comparison between human-driven and AI-assisted UML modeling, this study adopted a set of quality criteria inspired by established literature on model quality assessment [22–24]. These dimensions reflect common attributes used to evaluate conceptual models in Model-Driven Engineering and software design. Table 2 summarizes the key criteria applied in this study.

#### 4.4 Coding

This section evaluates the quality of the source code generated by LLM, comparing it to code manually written by developers in the development of LUNA. The evaluation focuses on essential software quality characteristics aligned with the ISO/IEC

**Table 2.** Comparative Evaluation of UML Models: Manual vs. LLM-Assisted Approaches

Criteria	Manual Modeling	LLM-Assisted Modeling
Accuracy	High – reflects system logic and domain requirements precisely	Medium – prone to semantic errors and generic representations
Syntax Reliability	Always syntactically valid	Often requires syntax correction (e.g., PlantUML errors)
Adaptability	Supports iterative refinement based on feedback	Requires full re-prompting for each adjustment
Semantic Context	Strong – interprets business logic and specific requirements	Weak – struggles with complex or domain-specific interactions
Speed (Initial Draft)	Slower for initial creation	Faster generation of basic structure
Final Usability	High – directly applicable in development workflows	Limited – requires human review and refinement

25010 standard [25], namely: maintainability, reliability, and efficiency. The objective is to analyze and compare code clarity, adherence to best practices, structural quality, and development productivity. The analysis focused particularly on the Maintainability characteristic, which includes sub-characteristics such as Analyzability, Modifiability, Modularity, and Testability. To operationalize these quality dimensions, specific metrics widely adopted in the software engineering field were selected and applied using SonarQube—a static code analysis tool. The evaluated metrics are summarized in Table 3.

All source code, both LLM-generated and human-written, was processed using SonarQube to detect issues such as bugs, code smells, duplicated code blocks, and documentation coverage. The results are shown in Table 4, in which the values reflect the averages of three LUNA modules: User Registration, Order Processing, and Inventory Management.

The results indicate that LLM-generated code is generally syntactically correct and logically consistent, but tends to exhibit higher complexity in certain situations, increased code smells, and lower documentation coverage. Despite these limitations, LLM-generated code was considered usable after minor adjustments. Beyond code quality, development productivity was also evaluated, considering the time spent creating and refining each module. Prompt engineering involved iterative refinement of input prompts tailored for various development tasks. For example, initial prompts for code generation were overly generic, yielding suboptimal code. Through successive refinement (e.g., specifying task context, expected coding standards, and framework-specific guidelines), prompt effectiveness improved significantly. A specific instance included refining a React.js component prompt from a generic ‘create a UI component’ to a detailed ‘create a responsive login form using Tailwind CSS with validation logic in React Hook Form’, dramatically enhancing generated code quality. The times spent are detailed in Table 5.

**Table 3.** Code quality metrics selected for this study, aligned with ISO/IEC 25010 sub-characteristics.

Metric	ISO/IEC 25010 Sub-Characteristic	Purpose
Cyclomatic Complexity	Analyzability/Modifiability	Measures code complexity and potential difficulty in understanding and modifying code
Code Smells (SonarQube)	Analyzability/Modularity	Identifies bad practices affecting code clarity and maintainability
Bugs Detected (SonarQube)	Reliability	Detecting potential defects compromising system stability
Modularity Score	Modularity/Reusability	Evaluates component independence and cohesion
Maintainability Index	Maintainability (Composite Indicator)	Aggregates multiple aspects of code quality
Documentation Coverage	Analyzability	Evaluates clarity through comments and code documentation
Linting Compliance	Maintainability/Modifiability	Verifies adherence to coding standards (e.g., ESLint for JavaScript)

**Table 4.** Comparative results of code quality metrics between manually written and LLM-generated code. Values reflect averages across three core modules of the LUNA.

Metric	Human Code	LLM-Generated Code
Cyclomatic Complexity (avg)	4.2	3.1
Code Smells (per 1k LOC)	3	7
Bugs Detected (Number)	0	2
Modularity Score (0–10)	8.5	7.2
Maintainability Index (0–100)	82	74
Documentation Coverage (%)	65	40
Linting Compliance (%)	96	88

On average, the use of LLM reduced development time by approximately 34%, even considering the time required for reviewing and correcting the generated code. This result aligns with findings from other studies [10] reporting productivity gains between 30% and 40% when applying prompt engineering techniques to LLM-based development. In

**Table 5.** Time spent on manual vs LLM-assisted development across three functional modules. Post-edit time includes reviewing and refining AI-generated outputs.

Module	Human Time (min)	LLM Time (min)	Post-edit Time (min)	Net Time Saved (min)
User Registration	90	40	20	30 (33%)
Order Processing	100	42	25	33 (33%)
Inventory Management	110	45	27	38 (35%)
Average	100	42.3	24	~ 34 (34%)

practical terms, the LLM allows developers to dedicate less time to routine coding tasks and focus more on higher-level concerns such as architecture refinement, optimization, and testing strategies. However, it became evident that the effectiveness of LLM strongly depends on the quality of prompts and the developer’s ability to guide and review the model’s output.

The architectural decision to transition from React.js to Next.js and from Express to NestJS was driven by enhanced server-side rendering capabilities, improved modular architecture, and ease of integration with advanced AI functionalities. However, these shifts introduced trade-offs, including increased initial complexity and the necessity for developer re-training. Potential challenges included managing the learning curve associated with NestJS’s advanced decorators and modular dependency injection system, which were mitigated through targeted training sessions and comprehensive documentation.

## 5 Discussion of Results

This section discusses the results obtained from the development of the LUNA using both manual and LLM-assisted approaches. The evaluation focused on performance, efficiency and quality at different stages of the software development lifecycle.

Results indicate that the use of LLM, particularly ChatGPT, significantly accelerated early-stage tasks such as requirements elicitation and initial design modeling. The structured application of prompt patterns enabled a 42.86% reduction in the time required for requirements analysis, while improving clarity and organizational output. These results are in the same direction as [21], where they demonstrated the effectiveness of LLM, such as GPT-4 and CodeLlama, in generating and validating software requirements specifications with a high degree of accuracy and completeness. Nonetheless, manual validation remained necessary to ensure accuracy. In addition to requirements and design, LLM contributed to faster code generation, and documentation, allowing developers to focus more on architecture refinement and user experience improvements. These findings align with Hamdi & Lim [1], who reported productivity gains of approximately 41% in AI-assisted development using prompt patterns with ChatGPT-4. Despite these positive

outcomes, several challenges were identified. Maintaining context throughout complex tasks, especially during integration phases, remains difficult for LLM. Moreover, the type and structure of prompts used continue to play a critical role in output quality—a topic that warrants further research.

Preliminary code quality analysis showed that LLM-generated code is generally well-structured and readable. However, it often lacks robustness in error handling and domain-specific logic when compared to human-written code. Manual code demonstrated greater alignment with business rules and architectural constraints, despite taking longer to produce. Overall, these results reinforce the potential of LLM as valuable support tools in software development, particularly for accelerating repetitive or structured tasks. Nevertheless, their current limitations prevent them from fully replacing human expertise in critical or context-sensitive development activities [16, 17].

Several threats to validity were identified and mitigated during the study. Participant-related bias was addressed by selecting a diverse developer group with varying experience levels with LLM-assisted development. The learning effect associated with LLM interaction was minimized by ensuring all developers underwent standardized LLM familiarization sessions. Baseline bias, resulting from the manual development approach, was mitigated by thoroughly documenting manual processes and using consistent criteria across both development scenarios to maintain fairness and comparability.

## 6 Conclusion

This study evaluated the practical use of LLM guided by prompt engineering across different phases of software development, using the LUNA as a real-world case study. The results indicate that LLM can significantly improve development efficiency, particularly in elicitation requirements.

Structured prompts proved essential to obtaining coherent and usable outputs, although human validation and refinement remained necessary, especially in complex tasks.

Preliminary results comparing AI-generated and manually written code suggest that LLM produce clean and syntactically correct code but struggle with complex business logic, edge cases, and contextual adaptation, areas where human developers continue to excel. The most significant challenges observed were: limited handling of edge cases and input validation; less consistent naming conventions and modularity; reduced documentation and comments. As AI adoption expands in software engineering, prompt engineering emerges as a critical practice, bridging the gap between human intent and LLM capabilities.

While LLM demonstrated potential in accelerating routine tasks, their application in system architecture and UML modeling still presents limitations. Human expertise remains essential to ensure quality, correctness, and alignment with complex business domains. In conclusion, the AI-focused design illustrates a potential future evolution of the LUNA, whereas the initial architecture adhered to traditional engineering principles and was implemented entirely by hand. It was influenced by the opportunities and difficulties found during the research as well as by the experience of utilizing LLM. It is crucial to emphasize that this history shows how AI-assisted development might impact

the architectural choices of contemporary software systems rather than invalidating the original design.

Future work should explore more advanced LLM fine-tuning, domain-specific training corpora, and hybrid human-AI workflows to improve semantic accuracy, test coverage, and architectural consistency.

## References

1. Jin, H., Huang, L., Cai, H., Yan, J., Li, B., Chen, H.: From LLMs to LLM-based Agents for Software Engineering: A Survey of Current, Challenges and Future, 05 Aug. 05 (2024) . <https://doi.org/10.48550/arXiv.2408.02479>
2. Alexander, C., Alexander, C.: A Pattern Language: Towns, Buildings, Construction. in Center for Environmental Structure Series. Oxford University Press, Oxford (1978)
3. About Solution-Patterns – VV-Patterns. <https://vvpatterns.ait.ac.at/about-vv-patterns/>
4. Coad, P.: Object-oriented patterns. Commun. ACM **35**(9), 152–159 (1992). <https://doi.org/10.1145/130994.131006>
5. Gamma, E., Helm, R., Johnson, R., Vlissides, J.: Design patterns: elements of reusable object-oriented software. Addison-Wesley Longman Publishing Co., Inc, USA (1995)
6. Coplien, O., Schmidt, D.: Pattern Languages of Program Design: 9780201607345: Amazon.com: Books. 12 Apr (2025). <https://www.amazon.com/Pattern-Languages-Program-Design-Coplien/dp/0201607344>
7. Rising, L. (ed.) The patterns handbook: techniques, strategies, and applications. Cambridge, U.K. ; New York : Cambridge University Press ; [New York] : SIGS Books, 1998. Accessed: Apr. 12, (2025). <http://archive.org/details/patternshandbook0000unse>
8. Adolph, S., Cockburn, A., Bramble, P.: Patterns for Effective Use Cases. Addison-Wesley Longman Publishing Co., Inc, USA (2002)
9. White, J., et al.: A Prompt Pattern Catalog to Enhance Prompt Engineering with ChatGPT, 21 Feb. (2023) . <https://doi.org/10.48550/arXiv.2302.11382>
10. Hamdi, M., Kim, L.D.: A prompt-based approach for software development. In: 2023 International Conference on Computational Science and Computational Intelligence (CSCI), pp. 1612–1614 (Dec. 2023). <https://doi.org/10.1109/CSCI62032.2023.00267>
11. Braun, M., Greve, M., Kegel, F., Kolbe, L.M., Beyer, P.E.: Can (A)I have a word with you? a taxonomy on the design dimensions of AI prompts. In: Proceedings of the 57th Annual Hawaii International Conference on System Sciences, HICSS 2024, Hawaii International Conference on System Sciences (HICSS), pp. 559–568 (2024). Accessed 12 Apr 2025
12. Tufano, M., Agarwal, A., Jang, J., Moghaddam, R.Z., Sundaresan, N.: AutoDev: Automated AI-Driven Development, 13 Mar (2024). <https://doi.org/10.48550/arXiv.2403.08299>
13. Schmidt, D.C., Spencer-Smith, J., Fu, Q., White, J.: Towards a catalog of prompt patterns to enhance the discipline of prompt engineering. Ada Lett. **43**(2), 43–51 (2024). <https://doi.org/10.1145/3672359.3672364>
14. Kim, D.-K.: Prompted Software Engineering in the Era of AI Models, 7 Sep (2023). <https://doi.org/10.48550/arXiv.2311.03359>
15. Liang, J.T., Lin, M., Rao, N., Myers, B.A.: Prompts Are Programs Too! Understanding How Developers Build Software Containing Prompts (2024). <https://doi.org/10.48550/ARXIV.2409.12447>
16. Patil, R., Gudivada, V.: A review of current trends, techniques, and challenges in large language models (LLMs). Appli. Sci. **14**(5), Art. no. 5 (2024). <https://doi.org/10.3390/app14052074>

17. Cámara, J., Troya, J., Burgueño, L., Vallecillo, A.: On the assessment of generative AI in modeling tasks: an experience report with ChatGPT and UML. *Softw. Syst. Model.* **22**(3), 781–793 (2023). <https://doi.org/10.1007/s10270-023-01105-5>
18. Wieringa, R.J.: The Design Cycle. In: Wieringa, R.J. (ed.) *Design Science Methodology for Information Systems and Software Engineering*, pp. 27–34 Springer, Berlin (2014). [https://doi.org/10.1007/978-3-662-43839-8\\_3](https://doi.org/10.1007/978-3-662-43839-8_3)
19. Xia, C.S., Deng, Y., Dunn, S., Zhang, L.: Agentless: Demystifying LLM-based Software Engineering Agents, 29 Oct (2024). <https://doi.org/10.48550/arXiv.2407.01489>
20. Schön, E.-M., Thomaschewski, J., Escalona, M.J.: Agile requirements engineering: a systematic literature review. *Compute. Standards Interfaces* **49**, 79–91 (2017). <https://doi.org/10.1016/j.csi.2016.08.011>
21. Hemmat, A., Sharbaf, M., Kolahdouz-Rahimi, S., Lano, K., Tehrani, S.Y.: Research directions for using LLM in software requirement engineering: a systematic review. *Front. Comput. Sci.* **7** (2025). <https://doi.org/10.3389/fcomp.2025.1519437>
22. Booch, G., Rumbaugh, J., Jacobson, I.: *The Unified Modeling Language User Guide*, 2nd edn. Addison-Wesley Professional, Upper Saddle River, NJ (2005)
23. Krogstie, J., Sindre, G., Jørgensen, H.: Process models representing knowledge for action: a revised quality framework. *Eur. J. Inf. Syst.* **15**(1), 91–102 (2006). <https://doi.org/10.1057/palgrave.ejis.3000598>
24. Lindland, O.I., Sindre, G., Solvberg, A.: Understanding quality in conceptual modeling. *IEEE Softw.* **11**(2), 42–49 (1994). <https://doi.org/10.1109/52.268955>
25. ISO/IEC 25010:2023, *Systems and software engineering — Systems and software Quality Requirements and Evaluation (SQuARE) — System and software quality models*, International Organization for Standardization (2023)



# Author Queries

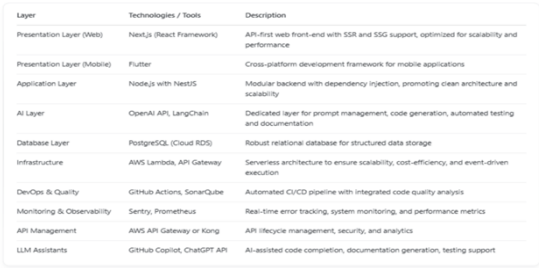

Chapter 20

---

Query Refs.	Details Required	Author's response
AQ1	This is to inform you that corresponding author has been identified as per the information available in the Copyright form.	

## Alternative Texts for Your Images, Please Check and Correct them if Required

Page no	Fig/Photo	Thumbnail	Alt-text Description
4	Fig1		<p>Flow chart illustrating a process with five sequential steps, each in a colored horizontal bar. From top to bottom: "Problem investigation" in orange, "Treatment design" in yellow, "Design validation" in light green, "Treatment implementation" in green, and "Implementation evaluation" in dark green. Each step represents a phase in a systematic approach.</p>
6	Fig2		<p>Table listing various requirements with three columns: Category, Sub-Category, and Examples. Categories include Functional Requirement for Mobile App and Commercial Platform, Non-Functional Requirement, and Social Media Integration Requirement. Sub-Categories cover User Management, Product Catalog &amp; Search, Shopping Cart &amp; Checkout, Inventory &amp; Order Management, Social Media Integration, Security, Performance, Scalability, Usability, Data Analytics, Marketing Automation, and User Behavior Insights. Examples provided include Registration, Login, Profile Editing, Browsing, Filtering, Product Details, Cart Operations, Payment Methods, Order Confirmation, Stock Management, Order Tracking, Sales Dashboard, Social Login, Behavior Data Collection, Campaign Tracking, GDPR, Encryption, Secure Payments, Fast Response Time, API Performance, High Availability, Auto-scaling, Cloud Infrastructure, Responsive Design, Accessibility, Intuitive Navigation, Sentiment Analysis, Influencer Detection, Campaign Scheduling, Performance Analytics, and Trends Tracking, Personalized Recommendations.</p>

Page no	Fig/Photo	Thumbnail	Alt-text Description																																	
			<p>Infrastructure, Responsive Design, Accessibility, Intuitive Navigation, Sentiment Analysis, Influencer Detection, Campaign Scheduling, Performance Analytics, Trends Tracking, and Personalized Recommendations.</p>																																	
7	Fig3	 <table border="1"> <thead> <tr> <th>Layer</th> <th>Technologies / Tools</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>Presentation Layer (Web)</td> <td>Next.js (React Framework)</td> <td>API-first web front-end with SSR and SSG support, optimized for scalability and performance</td> </tr> <tr> <td>Presentation Layer (Mobile)</td> <td>Flutter</td> <td>Cross-platform development framework for mobile applications</td> </tr> <tr> <td>Application Layer</td> <td>Node.js with NestJS</td> <td>Modular backend with dependency injection, promoting clean architecture and scalability</td> </tr> <tr> <td>AI Layer</td> <td>OpenAI API, LangChain</td> <td>Dedicated layer for prompt management, code generation, automated testing and documentation</td> </tr> <tr> <td>Database Layer</td> <td>PostgreSQL (Cloud RDS)</td> <td>Robust relational database for structured data storage</td> </tr> <tr> <td>Infrastructure</td> <td>AWS Lambda, API Gateway</td> <td>Serverless architecture to ensure scalability, cost-efficiency, and event-driven execution</td> </tr> <tr> <td>DevOps &amp; Quality</td> <td>GitHub Actions, SonarQube</td> <td>Automated CI/CD pipeline with integrated code quality analysis</td> </tr> <tr> <td>Monitoring &amp; Observability</td> <td>Sentry, Prometheus</td> <td>Real-time error tracking, system monitoring, and performance metrics</td> </tr> <tr> <td>API Management</td> <td>AWS API Gateway or Kong</td> <td>API lifecycle management, security, and analytics</td> </tr> <tr> <td>LLM Assistants</td> <td>GitHub Copilot, ChatGPT API</td> <td>AI-assisted code completion, documentation generation, testing support</td> </tr> </tbody> </table>	Layer	Technologies / Tools	Description	Presentation Layer (Web)	Next.js (React Framework)	API-first web front-end with SSR and SSG support, optimized for scalability and performance	Presentation Layer (Mobile)	Flutter	Cross-platform development framework for mobile applications	Application Layer	Node.js with NestJS	Modular backend with dependency injection, promoting clean architecture and scalability	AI Layer	OpenAI API, LangChain	Dedicated layer for prompt management, code generation, automated testing and documentation	Database Layer	PostgreSQL (Cloud RDS)	Robust relational database for structured data storage	Infrastructure	AWS Lambda, API Gateway	Serverless architecture to ensure scalability, cost-efficiency, and event-driven execution	DevOps & Quality	GitHub Actions, SonarQube	Automated CI/CD pipeline with integrated code quality analysis	Monitoring & Observability	Sentry, Prometheus	Real-time error tracking, system monitoring, and performance metrics	API Management	AWS API Gateway or Kong	API lifecycle management, security, and analytics	LLM Assistants	GitHub Copilot, ChatGPT API	AI-assisted code completion, documentation generation, testing support	<p>A table listing various layers, their corresponding tools or terminologies, and descriptions. The layers include "Process Edition Layer (Node)," "Process Edition Layer," "API Layer," "Data/Service Layer," "Infrastructure," "UNICA &amp; OLAP/BI," "Marketing &amp; Channel Mgmt," "Portfolio Management," and "LUV Analysis." Each layer is associated with specific tools or terminologies like "Node.js," "Hadoop," "OpenAPI 3.0," "PostgreSQL," and others. Descriptions provide details about the functionality and applications of each layer, such as API management, data storage, and marketing analysis.</p>
Layer	Technologies / Tools	Description																																		
Presentation Layer (Web)	Next.js (React Framework)	API-first web front-end with SSR and SSG support, optimized for scalability and performance																																		
Presentation Layer (Mobile)	Flutter	Cross-platform development framework for mobile applications																																		
Application Layer	Node.js with NestJS	Modular backend with dependency injection, promoting clean architecture and scalability																																		
AI Layer	OpenAI API, LangChain	Dedicated layer for prompt management, code generation, automated testing and documentation																																		
Database Layer	PostgreSQL (Cloud RDS)	Robust relational database for structured data storage																																		
Infrastructure	AWS Lambda, API Gateway	Serverless architecture to ensure scalability, cost-efficiency, and event-driven execution																																		
DevOps & Quality	GitHub Actions, SonarQube	Automated CI/CD pipeline with integrated code quality analysis																																		
Monitoring & Observability	Sentry, Prometheus	Real-time error tracking, system monitoring, and performance metrics																																		
API Management	AWS API Gateway or Kong	API lifecycle management, security, and analytics																																		
LLM Assistants	GitHub Copilot, ChatGPT API	AI-assisted code completion, documentation generation, testing support																																		
8	Fig4	 <pre> graph TD     Start(( )) --&gt; C1[Customer places the order]     C1 --&gt; S1[System checks stock]     S1 --&gt; D1{Items available?}     D1 -- No --&gt; E1[Error: Item unavailable]     D1 -- Yes --&gt; C2[Customer selects payment method]     C2 --&gt; S2[System validates payment]     S2 --&gt; D2{Payment valid?}     D2 -- No --&gt; E2[Payment error notified]     D2 -- Yes --&gt; C3[Order Preparation]     C3 --&gt; C4[Order Dispatched]     C4 --&gt; C5[Customer receives the order]     C5 --&gt; C6[Order status Completed]     E1 --&gt; End(( ))     E2 --&gt; End     C6 --&gt; End   </pre>	<p>Flowchart illustrating an order processing system. The process begins with the customer placing an order, followed by the system checking stock availability. If items are available, the customer selects a payment method, and the system validates the payment. If the payment is valid, the order is confirmed, prepared, dispatched, and received by the customer, completing the order. If the payment is invalid, a payment error is notified. If items are unavailable, an error message is displayed. Key steps include stock check, payment validation, and order completion.</p>																																	

