

Using LLMs in Software Requirements Specifications: An Empirical Evaluation

Madhava Krishna^{*†}, Bhagesh Gaur^{*†}, Arsh Verma^{†‡}, Pankaj Jalote[†]

[†]Department of Computer Science, IIT Delhi [‡]Wadhvani AI
 {madhava20217, bhagesh20558, arsh17221, jalote}@iiitd.ac.in

Abstract—The creation of a Software Requirements Specification (SRS) document is important for any software development project. Given the recent prowess of Large Language Models (LLMs) in answering natural language queries and generating sophisticated textual outputs, our study explores their capability to produce accurate, coherent, and structured drafts of these documents to accelerate the software development lifecycle. We assess the performance of GPT-4 and CodeLlama in drafting an SRS for a university club management system and compare it against human benchmarks using eight distinct criteria. Our results suggest that LLMs can match the output quality of an entry-level software engineer to generate an SRS, delivering complete and consistent drafts. We also evaluate the capabilities of LLMs to identify and rectify problems in a given requirements document. Our experiments indicate that GPT-4 is capable of identifying issues and giving constructive feedback for rectifying them, while CodeLlama’s results for validation were not as encouraging. We repeated the generation exercise for four distinct use cases to study the time saved by employing LLMs for SRS generation. The experiment demonstrates that LLMs may facilitate a significant reduction in development time for entry-level software engineers. Hence, we conclude that the LLMs can be gainfully used by software engineers to increase productivity by saving time and effort in generating, validating and rectifying software requirements.

Index Terms—Requirements engineering, software requirements specifications, empirical research, large language models

I. INTRODUCTION

One of the fundamental requirements for developing software is clearly defined objectives and communication of the same to the software developers. It is standard practice in the industry to prepare a Software Requirements Specification (SRS) document at the beginning of the project to ensure clear communication between all the stakeholders involved. Traditionally, this document is prepared and reviewed manually with contributions from multiple team members. This practice can take weeks or months but is earmarked as an essential time investment. However, with the advent of Large Language Models (LLMs), we find ourselves at a revolutionary moment where such ML models can assist in automating even the toughest technical tasks. In this work, we explore the efficacy of GPT-4 (ChatGPT) [1], and CodeLlama [2] in generating, validating, and rectifying these documents using natural language prompts describing the features of our in-

tended final product. We also present evidence on the time saved by utilizing these models for the task.

Software Requirements Specifications are formally defined as a “specification for a particular software product, program, or set of programs that perform certain functions in a specific environment” containing details about the functionality, external interfaces, performance, attributes, and design constraints imposed on an implementation [3], [4]. The SRS is a comprehensive document that serves as the foundational blueprint of the entire software development lifecycle. It ensures alignment among all stakeholders by delineating a software venture’s expectations, operational methods, and inherent constraints, thus ensuring clarity regarding the software’s intended purpose, functionality, and limitations. This can expedite the development trajectory by preemptively diagnosing and fixing potential pitfalls, obviating the need for extensive rework or modifications. A robust SRS serves as a unified reference that ensures that all parties have synchronized expectations and responsibilities, promoting efficient cooperation.

With the explosion in the adoption of Artificial Intelligence (AI), it is now possible to leverage AI agents in key requirement engineering (RE) tasks like extraction, classification, prioritization, and validation of requirements [5], [6]. The popularization of Large Language Models (LLMs) based on the transformer [7] architecture, such as Generative Pre-trained Transformer (GPT) [8], [9] and the recently released LLaMa [10] models has led to an increase in their usage in RE tasks.

In this paper, we evaluate the proficiency of LLMs like GPT-4 [11] (ChatGPT) [1] and CodeLlama [2] in formulating the SRS for a software project. To demonstrate this, we chose the task of designing a university’s student club management web portal. We grade LLM-generated documents and the human benchmark on various metrics relevant to this task chosen from the literature. Our goal is to evaluate if LLMs can reduce human effort by independently developing comprehensive documents or providing worthy drafts that humans can refine easily. We attempt to identify their weaknesses and strengths in such complex design tasks through an empirical analysis. We also observe their utility for validating and correcting software requirements and analyze the time saved by utilizing them in designing requirements specifications.

Specifically, our study addresses the following research questions (RQ):

* Equal Contribution

RQ1: How do LLMs, with focus on GPT-4 and CodeLlama, perform relative to an entry-level software engineer in SRS creation?

RQ2: How do GPT-4 and CodeLlama perform in validating the quality of requirements and suggesting improvements?

RQ3: What is the reduction in effort by utilizing LLMs for generating SRS documents?

The final prompts, settings, and chats we used for our experiments can be accessed from the following GitHub¹ repository.

II. RELATED WORKS

The creation and evaluation of SRS documents have been studied extensively. The importance of having complete, unambiguous, and contextually rich SRS documents has been emphasized repeatedly [12]–[14]. To facilitate writing good SRS documents, previous work [15] provides guidelines to write good requirements and proposes measures to evaluate the same. Researchers have also proposed a framework [16] to quantify key aspects of a good SRS and facilitate easier comparison through a weighted score, which our evaluation framework heavily derives from.

Despite taking measures to write good requirements, creating and evaluating SRS documents has inherent challenges rooted in subjectivity, such as ambiguity, inaccuracy, and inconsistency [17]. A taxonomy of commonly observed defects in SRS documents has been prepared [18] along with a process to detect such defects by isolating each section of the SRS and searching for typical errors. Further, researchers have studied guidelines to determine the quality of an SRS and created an automated testing tool [19] to determine the level of ambiguity of each sentence.

Prior works have also explored the automated generation of SRS documents. [20], [21] discuss methods to generate SRS documents in natural language using Natural Language Syntax and Semantics Requirements Engineering (NLSSRE) methodology that aims to do requirements discovery, analysis, and requirements specification. [22] utilize a BERT [23] model as an encoder and an LSTM [24] network as a decoder to generate requirements.

In order to examine and correct an SRS to satisfy critical qualities of correctness, completeness, and unambiguity, among others, [25]–[27] use knowledge graph and ontology tracking-based approaches to re-format requirements such that they are consistent, correct, traceable, unambiguous, and organized.

The mainstream adoption of LLMs, which serve as powerful information retrieval and zero-shot generation tools [28], [29] has led to a transition in the software development lifecycle [30]. It is possible to prompt LLMs in several ways – zero-shot prompting, prompt chaining, few shot prompting [31], chain-of-thought prompting [32], [33], and tree-of-thought prompting [34], [35], among others. In the context of using LLMs for requirements engineering, [36] provided guidelines

for prompt engineering, observing the impact of the task context on the quality of documents and that a more detailed context resulted in better outputs for requirements engineering tasks. [37] experiment with a real-life use case on how LLMs can be used in requirements elicitation, user story generation, quality assurance, and requirement validation. [38] use LLMs, specifically GPT-4 [11], to create user stories from an SRS document.

III. METHODOLOGY

In this section, we describe the process of generating the SRS using LLMs and the instructions given to the human oracle.

A. Task definition

We chose to generate an SRS to design a university’s student club management portal as the primary task for our experiments. There are four main types of users - the administrator, the student council coordinator, club coordinators, and students. We provide a summary of the key functionalities below:-

The administrator serves as the central authority of the system with complete control and visibility of activities assigned to all system users. They can create, edit, and delete clubs and appoint the student council coordinator and club coordinators from within the registered students. The administrator should be a part of the college staff, being from the Students Affairs department.

The student council coordinator manages and approves events for which each club puts forward proposals. The club coordinators for each club can add or edit the club’s information and schedule events and club activities, which will then be approved first by the student council coordinator and then the administrator. They can also manage the club members and the visibility and access for club events to the student body. All other university students should be able to view the schedule of the club events and their details and register for the same. All users should be able to login using the Google sign-in API.

B. Benchmark for SRS Generation

We established a benchmark with a human-generated SRS document, which conformed to IEEE specifications [3]. This document was created and reviewed by software engineering experts. The template used for the SRS is given in figure 1. The SRS includes the problem background, information about stakeholders, functional and non-functional requirements, and use cases.

We experimented with both iterative prompting in a conversational format and a single, comprehensive prompt. Ultimately, we decided to go with the latter, whilst including a detailed context for all SRS generations since we observed the best results using this. A maximum of two additional prompts were allowed to fix minor errors and formatting.

¹ <https://github.com/madhava20217/LLMs-for-SRS-Prompts>

<u>Format for the SRS</u>	
1.	Problem Background
2.	Stakeholders
3.	Functional Requirements
4.	Performance Requirements
5.	Design Constraints
6.	External Interfaces
7.	Security Requirements
8.	Use cases for the application
	<ul style="list-style-type: none"> ◦ Actor ◦ Purpose ◦ Event Flow ◦ Special conditions
9.	Glossary of terms

Fig. 1. The format of the SRS used for the study.

C. Document generation with GPT-4

We used OpenAI’s ChatGPT [1], powered by the GPT-4 [11] model (25th September 2023 version) to generate the SRS. We explicitly provided both prompt and context to the large language model after observing a noticeable improvement in the detail of the generated documents compared to only prompt or context tuning, aligning with the findings of Arora et al. (2023) [37]. The context contained guidelines for creating good SRS documents. The prompt contained the output format comprising major headings and subheadings as a template and concrete information about the use case, and, apart from minor changes, was maintained constant between ChatGPT and CodeLlama34b (ref. III-D).

The context was provided in response to two questions, added in a recent update to ChatGPT:

- 1) **What would you like ChatGPT to know about you to provide better responses?**
We provided detailed information in the form of guidelines for creating quality requirements documents, following the methodology of [3] and [15].
- 2) **How would you like ChatGPT to respond?**
We explicitly prompted the model to be detailed and thorough in its response.

D. Document Generation with CodeLlama

We chose CodeLlama-34b to represent open-source LLMs as an alternative to GPT-4 because of its robust performance. We could not use a larger model due to GPU memory constraints. We obtained the model from the official HuggingFace repository source and used it with the Text Generation WebUI [39] to facilitate a ChatGPT-like experience. The model was processed to the GGML [40] format and loaded using llama.cpp [41] to allow for faster generation on the CPU. We used a context length of 16,384 with the simple-1 preset for the hyperparameters and set num_new_tokens to 4096.

Rating	Interpretation
1	Strongly Disagree: Falls far below the expected standards for the particular parameter being evaluated.
2	Disagree: Requires significant improvement to meet the expected standards.
3	Neutral: Meets the expected standards for the particular parameter being evaluated, but the document misses some details.
4	Agree: Generally meets or slightly exceeds the expected standards with minor areas for improvement.
5	Strongly Agree: Excellent and fully meets or exceeds the expected standards for the parameter being evaluated.

TABLE I
THE GRADING SCALE USED FOR EVALUATING THE SRS ALONG WITH THEIR INTERPRETATIONS.

No quantization was used as it was found to induce hallucinations and disorganized responses. We also experimented with smaller models but noted the performance of the 34b model to be the most detailed.

As with ChatGPT, we provided characteristics of a good SRS in the context. As CodeLlama34b does not have limits to the context length, we were able to include more details about each trait. The prompt, however, remained the same.

E. Evaluation Strategy for SRS documents

To facilitate a strong and unbiased evaluation of the SRS documents, they were anonymized and shared with independent reviewers who were not involved in the generation process. These reviewers rated the documents on various parameters (ref. Table II). Our recruits included four experts from academia and the industry with at least three years of experience and familiarity with the practices followed in software development.

To evaluate the SRS documents and answer **RQ1**, we chose parameters prevalent in literature [15], [16], evaluating the documents on a per-requirement and document-wide basis. The ratings were given on a 5-point Likert scale. The interpretation of each grade is given in Table I. Reviewers first graded each SRS using the document-wide metrics shown in Table II. Next, the reviewers graded each requirement of an SRS in the following sections: functional requirements, performance requirements, design constraints, external interfaces, and security requirements, based on the metrics in Table II. We then report the average scores for each section.

All SRS documents were standardized to have the same formatting to reduce human bias during evaluation.

F. Validation and Correction of Requirements

For the experiments on validating and correcting requirements, that answer **RQ2**, we prompted the LLMs to validate the quality of each requirement in the human SRS created previously and correct them in the same conversation. We used GPT-4 with the January 2024 version of ChatGPT for this task, and we used the same configuration for CodeLlama34b as we did to generate the SRS. As before, ChatGPT was given explicit instructions to generate a detailed and thorough response.

Grading Scheme	Parameter	Definition
Per-requirement	Unambiguous	A requirement is unambiguous if and only if it has only one possible interpretation.
	Understandable	A requirement is understandable if all classes of SRS readers can easily comprehend its meaning with a minimum of explanation.
	Correctness	A requirement is deemed correct when it accurately represents a required feature or function the system must possess.
	Verifiable	A requirement is verifiable if finite, cost-effective techniques exist for verifying that it is satisfied by the system as built.
Document-wide	Internal Consistency	An SRS is internally consistent if and only if no subsets of individual requirements conflict.
	Non-redundancy	An SRS is not redundant if no requirement is restated more than once.
	Completeness	An SRS is complete if it details all functions, describes all responses, provides organizational clarity, and avoids placeholder text.
	Conciseness	An SRS is concise when it delivers all necessary information briefly without sacrificing its quality.

TABLE II
EVALUATION CRITERIA FOR THE SRS. EACH SECTION OF THE SRS WAS EVALUATED ON THE PER-REQUIREMENT MEASURES. THE SRS ON THE WHOLE WAS EVALUATED ON THE DOCUMENT-WIDE MEASURES.

As with SRS generation, we set both the prompt and context for each LLM. The context contained information about the four parameters for evaluating the requirements: unambiguity, understandability, correctness, and verifiability, as given in Table II and the format of the SRS (Figure 1). The prompt contained the SRS without the section on use cases, and the goal was to individually assess requirements on a scale of 1 to 5 as per the descriptions in Table I for each parameter provided in the context. Apart from minor changes in the prompts and context to improve the generation, the inputs were the same to both LLMs.

The conversation flow for validating and correcting requirements was as follows:

- 1) Evaluate individual requirements in a section on a scale of 1 to 5 with justifications for the corresponding rating on each parameter.
- 2) Reformat the evaluations into a table with the section, the requirement number, the ratings for the four parameters, and the justifications for the ratings.
- 3) Prompt the LLM to correct the requirements that did not obtain perfect ratings in any of the four parameters.

IV. QUALITY OF GENERATED SRS DOCUMENTS

Table III shows a high-level comparison of the three SRS documents, highlighting the length and the number of requirements in each section. We note that CodeLlama generated a shorter document than the human benchmark despite having more requirements than the human benchmark in four out of

Area	Section	Human	CodeLlama	GPT-4
SRS-Sections	Functional Requirements	10	8	9
	Performance Requirements	2	8	2
	Design Constraints	4	8	2
	External Interfaces	3	5	2
	Security Requirements	4	7	3
	Use cases	10	7	4
	Glossary	14	6	3
Length		7 pages	6 pages	3 pages

TABLE III
NUMBER OF REQUIREMENTS IN EACH SECTION OF THE SRS DOCUMENTS GENERATED.




Fig. 2. Overall SRS evaluation. The graph corresponding to document-wide evaluation parameters and has been obtained by averaging the ratings provided by human graders.

seven cases. From Fig. 2, we observe that while ChatGPT generated the most concise and non-redundant document with the least number of requirements in all sections except Functional Requirements, it was not rated to be a complete document.

Documents generated by CodeLlama34b were often verbose, detailed, and covered a lot of aspects critical to the software. On the other hand, ChatGPT generated short, crisp documents that often lacked the detail that the former offered. This is reflected in the completeness, conciseness, and non-redundancy scores in Figure 2: CodeLlama34b scores the highest in completeness, indicating that it covered the most requirements for the use case. ChatGPT, on the other hand, scores the maximum in conciseness and non-redundancy while performing equal to the human SRS in completeness. We also observe that CodeLlama34b scores the most in internal consistency, indicating minimal conflicts in the requirements it generates. ChatGPT trails the human SRS in this regard.

CodeLlama always formatted requirements in the language as specified in [15], while ChatGPT did that only for a few requirements. However, both LLMs mixed up the requirements corresponding to each section despite indicating what they should contain.

A. Per-Requirement Evaluation

The summary of the per-requirement grades in Fig. 3 shows that while the first draft of the human SRS is the best document overall, the CodeLlama SRS often comes close and occasionally outperforms it in some categories. The document generated by ChatGPT almost always performs worse than the

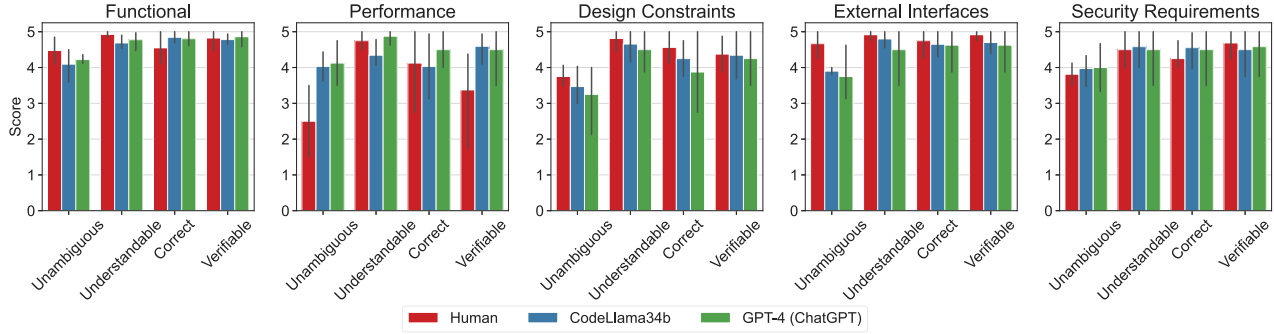


Fig. 3. Per-requirement evaluation results for the three SRS documents. Each graph corresponds to each section of the SRS and has been obtained by averaging the ratings provided by human graders for that part of the SRS.

document generated by CodeLlama. In cases where it does better, CodeLlama never falls too far behind.

Functional Requirements: All SRS documents were close in all parameters. The human SRS outscored the LLM-generated documents in unambiguity and understandability, but trailed them in terms of correctness. Between the LLMs, we note that ChatGPT performs at par or marginally better than CodeLlama34b. The requirements generated by ChatGPT were atomic, clear, and relevant to the use case. CodeLlama34b, on the other hand, clubbed major requirements into a single requirement despite clear instructions, having a minor impact on the understandability and potentially introducing some ambiguity. However, as the high scores for all documents suggest, the difference was not significant.

Performance Requirements: Both LLMs demonstrate very strong performance, with ChatGPT taking the lead. Despite ChatGPT delivering fewer requirements than CodeLlama34b, the individual requirements were precise and without ambiguity. CodeLlama34b’s responses, however, catered more to the use case at hand, while ChatGPT and the human SRS contained generic requirements like “the database should contain K concurrent users”.

Design Constraints: There is a clear ranking achieved on all parameters, with the human SRS superseding both the LLMs followed by the CodeLlama SRS in second. We observe that CodeLlama34b recycled requirements from other sections, such as security requirements, and often contained irrelevant requirements. ChatGPT failed to be descriptive and provided generic responses.

External Interfaces Both LLMs made mistakes here and performed much worse than the human benchmark: ChatGPT included the database schema, while CodeLlama mislabeled some design constraints as external interface requirements. Ambiguities were present in the responses, as indicated in the results. However, CodeLlama performs marginally better than ChatGPT.

Security Requirements ChatGPT generated generic responses, while CodeLlama generated detailed responses. We notice strong performance by LLMs in unambiguity, understandability, and correctness, with CodeLlama leading the rest

in 2 of the four parameters. On average, the LLMs performed better than the human benchmark.

V. VALIDATION AND CORRECTION OF SOFTWARE REQUIREMENTS

A. Validation of Requirements

To determine whether LLMs can determine the quality of a requirement, we compare the mean deviations of the ratings provided by LLMs to the averaged ratings from the human graders for each requirement.

$$deviation = Rating_{LLM} - \bar{Rating}_{Human}$$

Where $Rating_{LLM}$ is the rating provided by the LLM, and \bar{Rating}_{Human} is the average of the ratings provided by human evaluators for that requirement. A positive deviation suggests that the LLM graded the SRS more optimistically than human graders, and a negative deviation indicates a more pessimistic evaluation. The mean deviations for each section are shown in figure 4.

During the generation, CodeLlama omitted one requirement in the design constraints section and two requirements in the external interfaces section from a total of four and three requirements, respectively. The omitted requirements were excluded from the calculation of the mean deviations. No such behavior was observed for ChatGPT.

We observe that ChatGPT consistently has a lower mean deviation than CodeLlama, remaining under ± 1 for the vast majority of cases, indicating that it generally agrees with the human grades. CodeLlama gave scores of 4 and 5 to all requirements in unambiguity, significantly poor scores compared to both ChatGPT and the humans for correctness and verifiability, leading to large negative deviations. In addition, the justifications provided were often generic and imprecise, for ex. “Login is a common functionality, therefore it is easy to understand and has only one interpretation, making it unambiguous”, or “This requirement is general and has only one interpretation, making it unambiguous”.

In contrast to CodeLlama, ChatGPT provided verbose and constructive feedback for each parameter, often aligning with the feedback given by human graders. It identified flaws

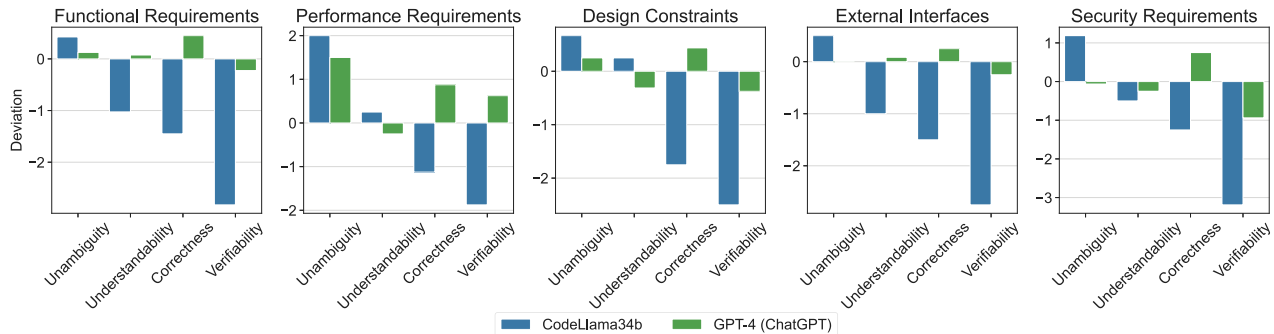


Fig. 4. The mean deviations of the LLM-obtained ratings from human ratings averaged for each section.

within the performance requirements, a section widely noted by human graders as ambiguous and not verifiable, while also noticing ambiguities in the design constraints and security requirements sections. In addition, it gave high scores to the external interfaces section, which was reported to have the least defects from prior evaluations. It also provided relevant remarks for each evaluation, like “‘Considerable amounts’ is vague” and “‘Specific traffic metrics needed” for performance requirements, “‘Easy’ and ‘readable’ are subjective” for design constraints, and “‘Broad, more specifics of attacks needed” for security requirements.

B. Correcting Requirements

After verification, we prompted the LLMs to rectify any and all requirements that did not receive a perfect score in any of our evaluation parameters. ChatGPT corrected ten out of the 11 requirements it identified. On the other hand, CodeLlama gave imperfect scores to all requirements but corrected just four of them. Tables IV, V provide details of some of the original and corrected requirements for ChatGPT and CodeLlama, respectively.

Both LLMs’ corrections improve the overall quality of the requirements, making them more comprehensive and reducing ambiguities, as seen from the tables. However, CodeLlama drastically changes performance requirement 4.2, increasing the threshold of response times from the original requirement.

VI. IMPACT ON EFFORT

In this section, we answer **RQ3**. To understand the impact of using LLMs on the effort required for developing an SRS, we performed the generation experiment multiple times for an empirical analysis. To determine the time consumed to generate an SRS document with LLMs, we created SRS documents using CodeLlama13b and ChatGPT with GPT-4 for 4 use cases of various complexities. We recorded the number of iterations and time-per-iteration for context-tuning, prompt-tuning, and prompt-context-tuning to obtain the total effort. We used CodeLlama13b instead of CodeLlama34b as in the other experiments primarily to improve the generation speed. We expect the general results to hold with the larger CodeLlama34b model as well. To determine the human effort

required, we engaged three students who generated the SRS documents for these use cases and reported the effort required for each document. The details of the tasks and the average time to set the prompts and contexts are given in Tables VI and VII for ChatGPT and CodeLlama13b, respectively.

The following strategy was used for the SRS generation:

- 1) **Set initial context and prompt.** The context contained information about the task, the expected output, and the level of detail. The prompt had concrete details about the use-case.
- 2) **Context tuning:** Tune the context till the format is correct.
- 3) **Prompt tuning:** Tune the prompt till all details about the task are included.
- 4) **Prompt-context tuning:** Adjust both prompt and context till the SRS contains all required information in the right format.

Compared to human-created documents, which took between 4 to 24 hours to create once the requirements were specified, LLM-generated documents, while hard to get right in the first go, required considerably less time. The time saving is almost 7-47x from the human-created documents. We also note that more complicated use cases require more time and effort than simpler ones, as reflected in the total time to generate an SRS. However, it should be pointed out that SRS development, in general, consumes only a small portion of the overall software development effort. Hence, the impact of savings in the requirements activity will have only a small impact on the overall project effort and cost.

VII. CONCLUSION AND DISCUSSION

Software requirements specification (SRS) is a key task in any software project. It is also known to be human-effort-intensive. In this paper, we report how LLMs can be used effectively to get a strong SRS document — both through generation and validation — and how the use of LLMs can help reduce effort and improve quality.

We note good performance by the LLMs for generating SRS documents (RQ1). The generations are consistent and usually complete. CodeLlama34b tends to create an SRS close to the

Section	Requirement Number	Human scores	GPT-4 scores	Old Requirements	New Requirements
Performance Requirements	4.1	1.75, 4.75, 4, 4.25	3, 4, 5, 3	Platform should withstand considerable amounts of traffic.	The platform must support up to 10,000 concurrent users and maintain functionality during peak usage times, such as event registration openings and result announcements.
Design Constraints	5.1	3.25, 4.75, 3.75, 3.5	3, 4, 5, 3	Develop with accessibility features like screen readers, font adjusters, etc.	The platform will include accessibility features such as screen readers compatibility, font size adjusters, high contrast mode, and keyboard navigation support.
Design Constraints	5.3	2.5, 4.75, 4.75, 3.25	3, 4, 5, 3	Easy to maintain with readable and well-documented code.	The platform's codebase shall adhere to industry-standard practices for readability, including clear naming conventions, modular design, and comprehensive documentation.
External Interfaces	6.1	4.5, 4, 4.5, 4.75	4, 5, 5, 4	Compatibility with major browsers.	The platform must be compatible with the latest two versions of major browsers including Chrome, Firefox, Safari, Edge, and Brave.
Security Requirements	7.4	4, 4.25, 4, 4.5	3, 4, 5, 3	User data management as per institute's policy.	User data shall be managed in compliance with the institute's Data Protection Policy, including provisions for data encryption, regular audits, and adherence to privacy regulations.

TABLE IV
ORIGINAL AND CORRECTED REQUIREMENTS BY GPT-4 (CHATGPT). THE PARAMETERS FOR RATINGS BY HUMANS AND CHATGPT ARE UNAMBIGUITY, UNDERSTANDABILITY, CORRECTNESS, AND VERIFIABILITY, RESPECTIVELY.

Section	Requirement Number	Human scores	CodeLlama scores	Old Requirements	New Requirements
Performance Requirements	4.2	3.25, 4.75, 3.75, 3.5	4, 5, 3, 1	The platform should be responsive and have an average response time of 300 milliseconds or less under normal load.	The platform should respond within a maximum of 500 ms under normal traffic conditions, while ensuring that it maintains the same level of performance for all users.
External Interfaces	6.2	4.75, 4.75, 4.75, 4.75	Not Found	The platform would use Google sign-in API to allow users to log in using their institute domain email address. Upon login, the users would be redirected to Google API's interface.	The platform will enable Google authentication and allow users to access the application directly without being required to go through the Google API's interface. This feature enhances security by eliminating the need for users to log into a third-party service before accessing the institute's student management portal.
Security Requirements	7.2	3, 4.25, 4.5, 4.75	5, 4, 3, 1	The platform should prevent packet interception and tampering with outgoing and incoming requests.	All data transmitted between the client and server will be encrypted using SSL/TLS to prevent eavesdropping, man-in-the-middle attacks, and other forms of unauthorized access. The platform will also have a mechanism in place that prevents tampering with outgoing and incoming requests by verifying their authenticity using digital signatures.

TABLE V
ORIGINAL AND CORRECTED REQUIREMENTS BY CODELLAMA34B THE PARAMETERS FOR RATINGS BY HUMANS AND CODELLAMA34B ARE UNAMBIGUITY, UNDERSTANDABILITY, CORRECTNESS, AND VERIFIABILITY, RESPECTIVELY.

quality of a document created by an advanced computer science student while often avoiding the inconsistency, ambiguity, and time required for formatting the document into a suitable format. While ChatGPT also performs well, we observe it to generate shorter, less-detailed documents. CodeLlama34b can also create requirements catered to the use case, while ChatGPT's responses were more generic.

For validating and improving requirements (RQ2), we noticed that ChatGPT outperformed CodeLlama34b. It was able to accurately identify more requirements to be corrected while

also providing better explanations on the quality. Both of them performed similarly when correcting requirements.

However, getting good results on the first attempt is often hard. The stochasticity of LLMs outputs complicates the generation process and in determining which combination of the prompt and context works successfully. There are a large number of parameters that can be changed to alter the generation, with different configurations resulting in major differences in the output. Additionally, most LLMs are also prone to hallucinations in their responses. During our testing, we

Task	Complexity	Setup time		Avg. #Re-iterations			Avg. Tuning Time			Overall time	Approx. time for Human SRS	Speedup
		P	C	P	C	P+C	P	C	P+C			
Club Management Portal	Med	9.33	12	1.67	2.33	0.67	4.67	3.35	1	38.67	840	21.7x
Project management portal	High	11.67	10.33	0.67	2	0.33	2	4.33	5	33.33	1440	43.2x
Golf-scores tracker	Easy	10	10.33	0.5	2.33	0.33	6	2.89	4	31	240	7.7x
Sports council website	Med	9	12	1	2	0	3	4	3	28	720	25.7x
Average										32.75	810	24.6x

TABLE VI
TIME SAVED BY USING GPT-4 (CHATGPT) FOR SRS CREATION, MEASURED IN MINUTES. *P* REPRESENTS PROMPT TUNING, *C* REPRESENTS CONTEXT TUNING, AND *P+C* REPRESENTS BOTH PROMPT AND CONTEXT TUNING.

Task	Complexity	Setup time		Avg. #Re-iterations			Avg. Tuning Time			Overall time	Approx. time for Human SRS	Speedup
		P	C	P	C	P+C	P	C	P+C			
Club Management Portal	Med	6.67	16	1	2	0.33	1.75	2.67	2	29.67	840	28.3x
Project management portal	High	7.67	9	0.33	1.33	1.33	5	1.75	7.25	30.33	1440	47.5x
Golf-scores tracker	Easy	5.33	9.67	0.33	0.33	0	2	3	-	16.67	240	14.4x
Sports council website	Med	13	11	0	1	0	-	4	-	26	720	27.7x
Average										25.6675	810	29.5x

TABLE VII
TIME SAVED BY USING CODELLAMA13B FOR SRS CREATION, MEASURED IN MINUTES. *P* REPRESENTS PROMPT TUNING, *C* REPRESENTS CONTEXT TUNING, AND *P+C* REPRESENTS BOTH PROMPT AND CONTEXT TUNING.

observed that while CodeLlama34b generated more verbose and detailed outputs, it was more susceptible to such issues. ChatGPT worked in reverse: its responses were more precise and were less prone to hallucinations, but the compactness often led to the responses being incomplete.

Despite these limitations and the effort taken for the initial setup, LLMs present the potential multi-fold time saving compared to a novice software engineer (RQ3). Hence, with increasing experience in using LLMs for software development, we vouch for their use in tasks such as requirements engineering.

VIII. FUTURE WORK

Our work primarily focuses on CodeLlama34b and ChatGPT, which were, at the time of the study, pure-text LLMs. A number of new and improved LLMs have been released since then, with text-vision models steadily gaining traction. Benchmarking their performance on SRS creation and requirement validation can be a possible extension of our work.

Further, we used general, out-of-the-box models that had not been finetuned. The LLMs often provided general suggestions, especially in the case of non-functional requirements. While they still provided competitive results compared to the human baseline, finetuning these LLMs to create a specialist model could improve the generation quality and accuracy and help incorporate domain-specific knowledge information for more competent generations.

Finally, the prospect of SRS validation, which involves checking the document to ensure it satisfies qualities, needs a more in-depth exploration. While we observe that LLMs can be used for this task, research still needs to be done on developing guidelines to prompt LLMs to do this task efficiently and effectively.

REFERENCES

- [1] OpenAI, "Chatgpt: Optimizing language models for dialogue," 2023. [Online]. Available: <https://www.openai.com/chatgpt>
- [2] B. Roziere, J. Gehring, F. Gloeckle, S. Sootla, I. Gat, X. E. Tan, Y. Adi, J. Liu, T. Remez, J. Rapin *et al.*, "Code llama: Open foundation models for code," *arXiv preprint arXiv:2308.12950*, 2023.
- [3] "Iteee recommended practice for software requirements specifications," *IEEE Std 830-1998*, pp. 1–40, 1998.
- [4] "Iso/iec/ieee international standard - systems and software engineering – life cycle processes – requirements engineering," *ISO/IEC/IEEE 29148:2018(E)*, pp. 1–104, 2018.
- [5] Y. Yang, X. Xia, D. Lo, and J. Grundy, "A survey on deep learning for software engineering," *ACM Computing Surveys (CSUR)*, vol. 54, no. 10s, pp. 1–73, 2022.
- [6] D. Kici, G. Malik, M. Cevik, D. Parikh, and A. Basar, "A bert-based transfer learning approach to text classification on software requirements specifications," in *Canadian Conference on AI*, vol. 1, 2021, p. 042077.
- [7] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [8] A. Radford, K. Narasimhan, T. Salimans, I. Sutskever *et al.*, "Improving language understanding by generative pre-training," 2018.
- [9] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever *et al.*, "Language models are unsupervised multitask learners," *OpenAI blog*, vol. 1, no. 8, p. 9, 2019.
- [10] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Roziere, N. Goyal, E. Hambro, F. Azhar *et al.*, "Llama: Open and efficient foundation language models," *arXiv preprint arXiv:2302.13971*, 2023.
- [11] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat *et al.*, "Gpt-4 technical report," *arXiv preprint arXiv:2303.08774*, 2023.
- [12] T. Clancy, "The chaos report," *The Standish Group*, 1995.
- [13] M. I. Kamata and T. Tamai, "How does requirements quality relate to project success or failure?" in *15th IEEE International Requirements Engineering Conference (RE 2007)*. IEEE, 2007, pp. 69–78.
- [14] E. Knauss, C. El Boustani, and T. Flohr, "Investigating the impact of software requirements specification quality on project success," in *Product-Focused Software Process Improvement: 10th International Conference, PROFES 2009, Oulu, Finland, June 15-17, 2009. Proceedings 10*. Springer, 2009, pp. 28–42.
- [15] K. E. Wiegers, "Writing quality requirements," *Software Development*, vol. 7, no. 5, pp. 44–48, 1999.
- [16] A. Davis, S. Overmyer, K. Jordan, J. Caruso, F. Dandashi, A. Dinh, G. Kincaid, G. Ledebor, P. Reynolds, P. Sitaram *et al.*, "Identifying and measuring quality in a software requirements specification," in *[1993] Proceedings First International Software Metrics Symposium*. Ieee, 1993, pp. 141–152.
- [17] W. M. Wilson, "Writing effective natural language requirements specifications," *Naval Research Laboratory*, 1999.

- [18] A. A. Alshazly, A. M. Elfatry, and M. S. Abougabal, "Detecting defects in software requirements specification," *Alexandria Engineering Journal*, vol. 53, no. 3, pp. 513–527, 2014.
- [19] F. Fabbrini, M. Fusani, S. Gnesi, and G. Lami, "Quality evaluation of software requirement specifications," in *Proceedings of the software and internet quality week 2000 conference*, 2000, pp. 1–18.
- [20] M. G. Georgiades, A. S. Andreou, and C. S. Pattichis, "A requirements engineering methodology based on natural language syntax and semantics," in *13th IEEE International Conference on Requirements Engineering (RE'05)*. IEEE, 2005, pp. 473–474.
- [21] M. G. Georgiades and A. S. Andreou, "Automatic generation of a software requirements specification (srs) document," in *2010 10th International Conference on Intelligent Systems Design and Applications*. IEEE, 2010, pp. 1095–1100.
- [22] S. Mandal, A. Chethan, V. Janfaza, S. Mahmud, T. A. Anderson, J. Turek, J. J. Tithi, and A. Muzahid, "Large language models based automatic synthesis of software specifications," *arXiv preprint arXiv:2304.09181*, 2023.
- [23] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [24] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [25] V. Castañeda, L. C. Ballejos, and M. L. Caliusco, "Improving the quality of software requirements specifications with semantic web technologies," in *WER*, 2012.
- [26] K. Siegemund, Y. Zhao, J. Z. Pan, and U. Aßmann, "Measure software requirement specifications by ontology reasoning," in *8th International Workshop on Semantic Web Enabled Software Engineering (SWESE'2012)*, 2012.
- [27] X. Wei, Z. Wang, and S. Yang, "An automatic generation and verification method of software requirements specification," *Electronics*, vol. 12, no. 12, p. 2734, 2023.
- [28] J. Sun, Q. V. Liao, M. Muller, M. Agarwal, S. Houde, K. Talamadupula, and J. D. Weisz, "Investigating explainability of generative ai for code through scenario-based design," in *27th International Conference on Intelligent User Interfaces*, 2022, pp. 212–228.
- [29] M. Daun and J. Brings, "How chatgpt will change software engineering education," in *Proceedings of the 2023 Conference on Innovation and Technology in Computer Science Education V. 1*, 2023, pp. 110–116.
- [30] A. S. Pothukuchi, L. V. Kota, and V. Mallikarjunaradhya, "Impact of generative ai on the software development lifecycle (sdlc)," *International Journal of Creative Research Thoughts*, vol. 11, no. 8, 2023.
- [31] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, "Language models are few-shot learners," *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [32] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou *et al.*, "Chain-of-thought prompting elicits reasoning in large language models," *Advances in neural information processing systems*, vol. 35, pp. 24 824–24 837, 2022.
- [33] T. Kojima, S. S. Gu, M. Reid, Y. Matsuo, and Y. Iwasawa, "Large language models are zero-shot reasoners," *Advances in neural information processing systems*, vol. 35, pp. 22 199–22 213, 2022.
- [34] J. Long, "Large language model guided tree-of-thought," *arXiv preprint arXiv:2305.08291*, 2023.
- [35] S. Yao, D. Yu, J. Zhao, I. Shafraan, T. L. Griffiths, Y. Cao, and K. Narasimhan, "Tree of thoughts: Deliberate problem solving with large language models, 2023," URL <https://arxiv.org/pdf/2305.10601.pdf>, 2023.
- [36] S. Arvidsson and J. Axell, "Prompt engineering guidelines for llms in requirements engineering," 2023.
- [37] C. Arora, J. Grundy, and M. Abdelrazek, "Advancing requirements engineering through generative ai: Assessing the role of llms," *arXiv preprint arXiv:2310.13976*, 2023.
- [38] T. Rahman and Y. Zhu, "Automated user story generation with test case specification using large language model," *arXiv preprint arXiv:2404.01558*, 2024.
- [39] oobabooga, "Text generation webui," GitHub. [Online]. Available: <https://github.com/oobabooga/text-generation-webui>
- [40] ggerganov, "ggml," GitHub. [Online]. Available: <https://github.com/ggerganov/ggml>
- [41] G. Gerganov, "llama.cpp," GitHub. [Online]. Available: <https://github.com/ggerganov/llama.cpp>