



Research article

Humans are still better than ChatGPT: Case of the IEEEExtreme competition

Anis Koubaa, Basit Qureshi, Adel Ammar, Zahid Khan, Wadii Boulila*, Lahouari Ghouti

Robotics and Internet-of-Things Laboratory, Prince Sultan University, Riyadh 12435, Saudi Arabia

ARTICLE INFO

Dataset link: <https://www.kaggle.com/datasets/riotulab/chatgpt-evaluation-on-ieeextreme-competitions>

Keywords:

ChatGPT
GPT-4
GPT-3.5
GPT performance
GPT limitations
OpenAI
NLP
Computer programming

ABSTRACT

Since the release of ChatGPT, numerous studies have highlighted the remarkable performance of ChatGPT, which often rivals or even surpasses human capabilities in various tasks and domains. However, this paper presents a contrasting perspective by demonstrating an instance where human performance excels in typical tasks suited for ChatGPT, specifically in the domain of computer programming. We utilize the IEEEExtreme Challenge competition as a benchmark—a prestigious, annual international programming contest encompassing a wide range of problems with different complexities. To conduct a thorough evaluation, we selected and executed a diverse set of 102 challenges, drawn from five distinct IEEEExtreme editions, using three major programming languages: Python, Java, and C++. Our empirical analysis provides evidence that contrary to popular belief, human programmers maintain a competitive edge over ChatGPT in certain aspects of problem-solving within the programming context. In fact, we found that the average score obtained by ChatGPT on the set of IEEEExtreme programming problems is 3.9 to 5.8 times lower than the average human score, depending on the programming language. This paper elaborates on these findings, offering critical insights into the limitations and potential areas of improvement for AI-based language models like ChatGPT.

1. Introduction

1.1. Background and motivation

Large Language Models (LLMs) [1] have emerged as a groundbreaking artificial intelligence technology, especially since the release of ChatGPT in late November 2022. LLMs can mimic human-level capabilities in various complex natural language processing and understanding tasks across multiple domains, such as virtual assistants, chatbots, language translation, sentiment analysis, and more. ChatGPT has been trained on an extensive corpus of data spanning various disciplines, enabling it to acquire a broad spectrum of knowledge. Its training data comprises diverse sources from multiple domains, including but not limited to science, literature, law, programming, finance, and many more. This various training data has given ChatGPT a global perspective, making it capable of understanding and generating responses across a wide range of subjects. The vast knowledge base of ChatGPT allows it to provide

* Corresponding author.

E-mail addresses: akoubaa@psu.edu.sa (A. Koubaa), qureshi@psu.edu.sa (B. Qureshi), aammar@psu.edu.sa (A. Ammar), zskhan@psu.edu.sa (Z. Khan), wboulila@psu.edu.sa (W. Boulila), lghouti@psu.edu.sa (L. Ghouti).

<https://doi.org/10.1016/j.heliyon.2023.e21624>

Received 27 May 2023; Received in revised form 6 October 2023; Accepted 25 October 2023

Available online 29 October 2023

2405-8440/© 2023 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

insights and solutions to complex problems that span different domains, making it an effective tool for various applications in natural language processing and understanding.

With the ChatGPT's unprecedented capabilities compared to other LLMs in competing with humans across various applications, there has been a significant increase in the number of studies investigating its performance in specialized and complex domains [2], such as healthcare [3], and finance [4]. However, despite the growing interest in evaluating ChatGPT's abilities in these areas, there has been a lack of research specifically focusing on its specific performance in problem-solving and programming assessment domains, which is the main focus of this paper. This research gap has motivated us to investigate and evaluate ChatGPT's abilities in these areas.

1.2. Objective

This paper aims to investigate the problem-solving capabilities of ChatGPT by evaluating its performance on programming problem benchmarks. Our objective is to assess how ChatGPT compares to human programmers and to extract valuable insights into its strengths and weaknesses in this domain-specific context.

To accomplish our research objective, we identified the IEEEExtreme Programming Challenge as the most reputable and prestigious competition that could serve as an appropriate benchmark for comparing the problem-solving abilities of ChatGPT and human programmers. The IEEEExtreme Programming Challenge is an annual international programming competition organized by the Institute of Electrical and Electronics Engineers (IEEE). This 24-hour competition attracts programming professionals from across the globe to compete in solving programming problems with varying degrees of complexity, which demand high-level problem-solving and programming skills.

In summary, the primary objective of this study is to evaluate the performance of ChatGPT in the programming and problem-solving-specific context. To this end, we will employ the IEEEExtreme Challenge competition as a benchmark, utilizing problems of varying complexities. Moreover, we aim to analyze the limitations of ChatGPT in solving specific problems and programming tasks and identify areas for improvement and optimization. By conducting this analysis, we aim to provide the community with insights into the effectiveness of ChatGPT in programming and problem-solving domains and provide recommendations for future developments in this area.

1.3. Methodology

For this study, we selected five IEEEExtreme programming competitions, each consisting of an average of 20 questions. To guide ChatGPT in designing solutions while meeting non-functional requirements such as memory usage and execution time, we designed well-crafted prompts. For each problem, we presented the prompt to ChatGPT and evaluated its corresponding solution using Hackerrank, which was used to generate scores. We evaluated solutions in three top programming languages: Python 3, C++ 11, and Java 7. In case of errors, we made up to seven attempts to guide ChatGPT toward the correct solution by providing the corresponding Hackerrank error message.

Furthermore, to ensure the consistency of the results, we conducted this process three times, using different ChatGPT chat windows for each programming problem of the five selected IEEEExtreme Challenges. The final results were analyzed, and we identified ChatGPT's limitations in solving specific problems and programming tasks. Additionally, we provided recommendations for areas of improvement and optimization, which could enhance ChatGPT's effectiveness in programming and problem-solving domains.

Research questions

In this study, we aim to respond to four research questions:

- 1) How does ChatGPT compare to human programmers in problem-solving and programming tasks, in the context of the IEEEExtreme Challenge competition?
- 2) In which specific programming tasks or problem types do humans outperform ChatGPT, and what are the underlying reasons for this disparity?
- 3) Is ChatGPT performance biased towards particular programming languages among the three selected languages, namely, Python, C++ 11, and Java?
- 4) What are the fundamental limitations of ChatGPT in programming and problem-solving, and how can these findings guide future research and development of ChatGPT and other domain-specific large language models?

1.4. Overview of the paper structure

The paper is organized as follows. Section 2 provides a literature review on ChatGPT and its applications, human performance in programming tasks, and previous comparisons between AI and human performance. Section 3 describes the methodology, including selecting IEEEExtreme challenges, evaluation criteria and metrics, programming languages used, and the data collection and analysis approach. Section 4 presents the results of the study that compares ChatGPT and human performance in programming tasks of IEEEExtreme challenges and identifies the gap with human-level performance. The interpretation of the results and limitations of ChatGPT in programming and problem-solving tasks, as well as the reasons for disparities, are discussed in the same section. Finally,

Table 1
Related work on ChatGPT.

Paper	Approach	Main Finding	Limitations
Guo et al. [8]	NLP	ChatGPT is more helpful than human experts in finance and psychology questions.	Poor performance in the medical domain; may fabricate facts.
Qin et al. [29]	NLP	ChatGPT perform well on reasoning tasks but struggle with specific tasks.	Outperformed by models fine-tuned for specific tasks.
Kashefi and Mukerji [30]	Programming	ChatGPT can program numerical algorithms but face challenges generating long codes and using unknown libraries.	Issues with singular matrices and incompatible arrays.
Liu et al. [31]	Logical Reasoning	ChatGPT outperforms RoBERTa on most benchmarks but struggles with newly-released datasets.	Challenge in dealing with out-of-distribution and natural language inference datasets
Tian et al. [33]	Programming	Highlights the ability of ChatGPT to provide explanations and guidance to help users understand complex concepts and resolve technical issues.	No comparison with human programmers to evaluate ChatGPT's performance relative to human experts.
Hu et al. [34]	Refactory tool in Programming	ChatGPT achieves competitive results in semantic-based assignments repair.	Limited attention span affects performance; No comparison with human programmers.
Surameery and Shakor [35]	Programming Debugging	ChatGPT can enhance efficiency and effectiveness in bug identification and resolution.	Not a standalone solution; output should be verified and validated with traditional methods.
Piccolo et al. [36]	Bioinformatics Programming	ChatGPT can effectively complete a substantial percentage of basic to moderate-level bioinformatics programming tasks with the assistance of natural-language feedback.	The study does not address ChatGPT's performance on complex programming problems, leaving a gap in understanding its effectiveness in handling more challenging programming tasks.
Biwas et al. [37]	Programming	ChatGPT performs well in the code generation task but struggles to generalize to new and unseen problems.	The study is only limited to exploratory analysis without empirical results. No comparison with human performance
Chen et al. [38]	Programming assistance	GPTutor addresses a practical need in the programming community by offering a tool that can assist users in understanding code segments.	No quantitative performance evaluation.
Avila et al. [39]	Programming	The use of ChatGPT as a programming assistant to develop an online behavioral task using HTML, CSS, and JavaScript code.	There is no quantitative study to assess the performance of ChatGPT. It also lacks comparison with human performance.

Section 5 concludes the paper with a summary of findings, potential areas for improvement in ChatGPT, and suggestions for future research directions.

2. Related works

2.1. ChatGPT and its applications

As AI technologies continue to make significant strides across various domains, there is growing interest in how AI-human collaboration can be harnessed to augment and amplify human capabilities. This exploration spans a wide spectrum of applications, ranging from healthcare and education to industry and beyond. By investigating the interplay between AI and human contributions, researchers seek to uncover novel avenues for innovation, address limitations, and forge new frontiers that leverage the strengths of both entities.

More specifically, ChatGPT has made significant progress and it has been used in various applications. The detailed comparison of ChatGPT performance in various domains is shown in Table 1. In our previous study, ChatGPT's applications were classified into five main categories [5]:

- **NLP:** NLP stands for Natural Language Processing. It is a field of artificial intelligence and computational linguistics that focuses on the interaction between computers and human language [6]. In this type of application, ChatGPT generates human-like responses in natural language. Applications of ChatGPT in NLP [7] include building virtual assistants, chatbots, language translation systems, and text generation tasks such as summarization and question answering [8–10].
- **Healthcare:** ChatGPT has been used in various healthcare fields. It has been applied in healthcare decision support to provide relevant information and recommendations [11]. In addition, many recent research works have investigated the case of using ChatGPT in patient education, where ChatGPT provides patients with educational information about their health conditions, treatments, and medications [12,13]. Moreover, ChatGPT has been included in applications related to telemedicine to provide more efficient and accurate virtual diagnosis and treatment [14]. Furthermore, GPT-4's potential in surgery [15,16] and various branches of biomedical engineering [17] (including medical imaging, medical devices, bioinformatics, biomaterials, biomechanics, gene and cell engineering, tissue engineering, and neural engineering) revealed large opportunities for improving healthcare quality. However, these works came to the conclusion that ChatGPT/GPT-4 is not poised to supplant the role of physicians

and surgeons; rather, it is foreseeable that AI-powered tools like ChatGPT will advance and seamlessly integrate into medical procedures, complementing the proficiency of experienced physicians. By fostering collaboration between ChatGPT and human knowledge, the potential arises for enhanced outcomes and an elevated standard of healthcare delivery.

- **Ethics:** Many recent research works addressed the challenge of using ChatGPT for the benefit of society and how to maintain public safety [18]. Many authors explored using ChatGPT to generate student works and scientific publications [19]. Many other researchers focused on the ethical concerns, data biases, and safety issues related to ChatGPT [20].
- **Education:** ChatGPT has played an essential role in several applications in education [21–23]. It helped to improve the learning experience for students. It can provide personalized educational content. In addition, it can generate educational materials for students and tutors. ChatGPT is considered a promising tool for education, as it can provide insightful direction and feedback. Jeon and Lee [24] underscored the potential benefits of AI-powered chatbots like ChatGPT in education but emphasized that their success depends on a collaborative approach between teachers and AI. Teachers bring their expertise in pedagogy and ethical guidance, while AI can assist in various ways. The study's implications suggest that future implementations of AI-powered chatbots in education should prioritize this teacher-AI collaboration to maximize their educational impact. Ji et al. [25] conducted a systematic review that underscored the potential for AI-integrated language learning to benefit students and educators. While there are challenges to overcome, such as the need for more evidence of effective collaboration between AI and human teachers, the review provides guidance for future research and implementation in this field. Ultimately, the goal is to create a collaborative and effective learning environment that combines the strengths of AI and human educators.
- **Industry:** recently, various applications across many industries have been focused on using ChatGPT to improve efficiency, streamline processes, and enhance customer experiences [26–28]. Applications include the manufacturing industry, where it can monitor and control production processes. In addition, ChatGPT is used in the financial sector, where it can offer support to customers and company owners. Moreover, it can provide customer support to handle routine inquiries.

In this paper, we are specifically interested in a novel and thought-provoking perspective that challenges the prevailing notion of AI's unparalleled performance. While ChatGPT has demonstrated remarkable aptitude across diverse applications, as discussed above, our focus narrows down to a domain where traditional human expertise appears to excel: computer programming. By delving into the IEEEExtreme Challenge competition, we empirically assess ChatGPT's performance against that of human programmers. This distinctive context not only enriches our understanding of ChatGPT's capabilities but also offers insights into the intricate interplay between AI and human problem-solving skills.

2.2. Previous comparisons between AI and human performance

In a recent study [8], Guo et al. compared the responses of ChatGPT and human experts to around 40 K questions in various domains, such as finance, psychology, medical, legal, and open-domain, in both English and Chinese languages. They analyzed ChatGPT's response characteristics, differences and gaps from human experts, and future directions for LLMs. The researchers discovered that ChatGPT's responses are generally more helpful than human experts' in over half of the questions, especially in finance and psychology, due to its ability to offer specific suggestions. However, ChatGPT performs poorly in the medical domain. The authors also found that ChatGPT writes in an organized manner, with clear logic, and tends to provide detailed answers with less bias and harmful information. However, it may fabricate facts. Notably, the study did not include programming tasks but only theoretical questions about computer science-related concepts taken from Wikipedia.

On another hand, Qin et al. [29] examined the zero-shot learning ability of ChatGPT. The evaluation was conducted on 20 commonly used natural language processing (NLP) datasets covering seven task categories, including natural language inference, question answering, dialogue, summarization, named entity recognition, and sentiment analysis. However, the study did not include any programming tasks. The authors performed extensive empirical studies to analyze the strengths and limitations of the current version of ChatGPT. They discovered that ChatGPT performed well on tasks that require reasoning abilities, such as arithmetic reasoning, but it struggled with specific tasks like sequence tagging. Additionally, ChatGPT was outperformed by previous models that had been fine-tuned for a specific task. The findings suggest that ChatGPT is still far from reaching perfection as a generalist model.

Kashefi and Mukerji [30] investigated the potential of ChatGPT in one specific aspect of programming, which is to produce numerical algorithms for solving mathematical problems. They explored generating code in different programming languages, debugging user-written code, completing unfinished code, rewriting code in different programming languages, and parallelizing serial code. Although the study outcomes demonstrated that ChatGPT is capable of programming numerical algorithms, certain limitations and challenges were encountered. These included issues such as generating singular matrices, producing incompatible arrays, and irregular interruption when generating long codes required for scientific simulations. Another challenge was the inclusion of unknown libraries. Despite these limitations, the study suggests that ChatGPT has the potential for further development and improvement in programming numerical algorithms.

Liu et al. [31] evaluated the performance of ChatGPT and GPT-4 on various logical reasoning tasks using multiple datasets, including both well-known benchmarks and newly-released ones. The experiments showed that ChatGPT performs better than the RoBERTa [32] fine-tuning method on most logical reasoning benchmarks. However, both ChatGPT and GPT-4 struggle with newly-released and out-of-distribution datasets. GPT-4 showed higher performance than ChatGPT on most logical reasoning datasets. Nevertheless, despite advancements in models like ChatGPT and GPT-4, the task of logical reasoning still poses significant challenges for these models, particularly when dealing with out-of-distribution and natural language inference datasets.

Tian et al. [33] presented an empirical study evaluating the potential of the ChatGPT generative large-scale language model as an assistant bot for programmers. The study assesses ChatGPT's performance on three code-related tasks: code generation, program repair, and code summarization. ChatGPT is found to perform well in the code generation task but struggles to generalize to new and unseen problems. The study also highlights the negative impact of long prompts on ChatGPT's inference capabilities. In the program repair task, ChatGPT achieves competitive results compared to Refactory [34], a state-of-the-art semantic-based assignments repair tool. However, prompts that are not related to bug information are found to make ChatGPT perform even worse due to its limited attention span. The study's limitation pertains to the absence of a comparison between ChatGPT's performance and that of human programmers, thus hindering the establishment of its proficiency in relation to human experts.

Similarly, Surameery and Shakor [35] suggested that ChatGPT can be a valuable addition to a programmer's toolkit for debugging. It can complement other debugging tools and techniques by leveraging its natural language processing abilities to provide human-readable insights into code issues. However, it should not be viewed as a standalone solution, and its output should be verified and validated using traditional debugging methods. The research underscores the potential benefits of integrating ChatGPT into a comprehensive debugging strategy to enhance the efficiency and effectiveness of bug identification and resolution in programming.

Piccolo et al. [36] conducted an evaluation of OpenAI's ChatGPT using 184 programming exercises from an introductory bioinformatics course. The study aimed to assess ChatGPT's ability to successfully complete basic to moderate-level programming tasks. The results of the evaluation showed that, on its first attempt, ChatGPT was able to solve 139 out of the 184 exercises, achieving a success rate of 75.5%. For the remaining exercises, the authors provided natural-language feedback to ChatGPT, guiding it to explore different approaches to solving the problems. Remarkably, within seven attempts or fewer, ChatGPT was able to successfully complete 179 out of the 184 exercises, achieving a high success rate of 97.3%. Nevertheless, contrary to the present study, these experiments did not include complex programming problems. We will show that ChatGPT's performance on this category of problems is markedly lower.

Biswas [37] explored existing language models and tools for computer programming. ChatGPT is introduced as a powerful and versatile tool that can perform a variety of programming-related tasks such as code completion, correction, optimization, and refactoring. The paper highlights the ability of ChatGPT to provide explanations and guidance to help users understand complex concepts and resolve technical issues. The use of ChatGPT is noted as a potential means to improve overall satisfaction with support services and build a reputation for expertise and reliability. In summary, the paper suggests that ChatGPT is a valuable resource for technical support and improving efficiency and accuracy in computer programming tasks. The author solved simple programs without comparing them with human performance. Additionally, the author's work is only limited to exploratory analysis without empirical results.

In a similar fashion, Chen et al. [38] introduced GPTutor, a programming tool powered by ChatGPT, which operates as a Visual Studio Code extension using the ChatGPT API to provide explanations for programming code. GPTutor can analyze provided code comprehensively, referencing relevant source code to explain selected code segments with pop-up messages. Initial evaluations suggest that GPTutor offers more concise and accurate explanations compared to vanilla ChatGPT and GitHub Copilot, and feedback from students and teachers indicates user-friendliness and satisfactory code explanations. The paper also discusses potential future research directions, including performance enhancement through prompt programming and evaluation with real users. However, no quantitative performance evaluation was provided.

In reference [39], the authors discussed the challenges faced by behavior analysts in automating and systematizing experimental tasks. With the development of online platforms, OpenAI ChatGPT has emerged as a chatbot that can generate text responses similar to humans in a conversational context. One of its key functions is the ability to generate programming code blocks in various programming languages. The article presents the use of ChatGPT as a programming assistant to develop an online behavioral task using HTML, CSS, and JavaScript code. While ChatGPT cannot replace programmers entirely, it can provide detailed programming solutions and reduce the time associated with programming. The authors assess the performance of the ChatGPT with random problems in diverse directions. There is no quantitative study to assess the performance of the ChatGPT. It also lacks comparison with human performance.

The previous studies on ChatGPT mainly explored its performance in various contexts, but most of them did not follow a quantitative approach. In contrast, our study quantitatively evaluates ChatGPT's performance in solving IEEE Xtreme problems and compares it to average human performance in three different programming languages.

3. Methodology

3.1. The IEEEExtreme competition

There are several global programming competitions including, IEEE Extreme Programming Competition (IEEEExtreme) [40], ACM International Collegiate Programming Contest (ICPC) [41], Google Code Jam [42], Facebook / Meta Hacker Cup [43], and International Olympiad in Informatics (IOI) [44], to name a few. The IEEEExtreme is a global programming competition organized by the Institute of Electrical and Electronics Engineers (IEEE). The IEEEExtreme programming competition has been running annually since 2006. The number of participants in the competition has been increasing each year. In the early years, the competition had around 500 teams participating. In recent years, the number of participating teams has grown to around 10,000 or more, with participants from over 100 countries. The competition has become a major event in the global programming community, attracting top talent worldwide. The competition provides a platform for students to showcase their technical skills and talent. Winning the competition is a significant achievement that can help participants stand out to potential employers or graduate schools.

Table 2
Classification of challenges in IEEEXtreme Competitions.

Competition Version	Total	Available	Easy	Medium	Hard	Advanced
Extreme 8	23	4	1	3	0	0
Extreme 9	29	29	5	8	8	8
Extreme 10	24	17	0	0	17	0
Extreme 15	26	26	8	10	8	0
Extreme 16	26	26	9	10	7	0

It is a 24-hour coding marathon where teams of up to three students worldwide compete to solve a series of challenging programming problems. The problems are usually related to topics in computer science, mathematics, and engineering. The competition is designed to encourage and develop programming skills in students, as well as to promote teamwork and creativity. Participants must rely on their knowledge of algorithms, data structures, and programming languages to solve problems within the time limit. The competition is judged based on the number of problems solved, with ties broken based on the time taken to solve them. Each problem has a set number of test cases defined in the evaluation platform. Students need to provide a programming solution to the given problem by passing all the test cases to earn scores.

3.2. Selecting IEEEXtreme challenges

IEEEXtreme competition editions 11 and beyond were held on the CSacademy platform [45], while the earlier editions were conducted on the Hackerrank platform [46]. These platforms are accessible worldwide and open to all, and both may rely on Amazon Web Services (AWS) for their infrastructure. The entire problem set for IEEEXtreme competitions versions 15 and 16, which were hosted in 2021 and 2022 respectively, are available on the CSacademy platform. The earlier versions are unfortunately not available. Selected problems from versions 8, 9, and 10 are available on the Hackerrank platform, however, these can be accessed using the practice community website only. In this work, we rely on the problem sets presented in IEEEXtreme versions 8, 9, 10, 15, and 16. Table 2 shows a list of problems presented in each of these competitions. Each problem is classified based on difficulty level defined by the organizers as easy, medium, hard, and advanced. While all problems are available for IEEEXtreme versions 9, 15, and 16, only a select few are available for versions 8 and 10, on the hackerrank [47] practice community website.

3.3. Solving and scoring a problem

As mentioned in the previous section, each competition comprises a certain number of problems. A problem statement generally includes a brief description of the problem, its input and output format, and any constraints or limitations that apply to the solution. The problem may also include sample test cases that provide examples of the expected input and output. Participants are expected to write a computer program using a programming language of their choice, that solves the problem statement and produces the expected output. The solution is then submitted to the competition's online system, which tests the program against a variety of test cases and assigns a score based on its accuracy and efficiency. In recent versions of the competition, participants are expected to submit correct solutions that satisfy the minimum execution time and memory limitations. This usually requires participants to post optimized solutions based on the correct choice of programming constructs, efficient data structures, and algorithms.

To obtain scores for each problem, participants can submit solutions multiple times to pass most, if not all, of the hidden test cases. However, multiple submissions to the same problem typically result in point deductions, known as penalties, which are factored into the total score in case of tiebreakers between teams scoring the same number of points. A Programmer may solve any of these problems using a programming language of their choice including but not limited to C/C++, java, python, etc. At the end of the competition, the website displays relevant information for each problem, including the average score earned per team and the percentage of teams that attempted to solve the problem. This important information factors into Human performance in this research work and is compared with AI performance, as explained in the next section.

3.4. ChatGPT code generation and data collection

As the IEEEXtreme programming competition is open to participants worldwide, it is essential that each participant has access to the necessary resources to solve the problems. With this goal in mind, we have developed a method for participants in this study to simulate the experience of competing in IEEEXtreme by using ChatGPT to solve the problems. We evaluated the problem-solving performance of ChatGPT and human programmers using three primary programming languages: Python 3, Java 7, and C++. These languages were selected based on their popularity in the programming community and their suitability for solving the programming problems provided by the IEEEXtreme Challenge competition. We ensured that all participants were familiar with these languages and had equivalent levels of proficiency.

As a problem appears, the participants simply copy and paste the problem statement from the competition website into the prompt space. ChatGPT then generates detailed results that include an explanation of the algorithm, complete executable code in the selected programming language, sample test cases, and other relevant information needed to solve the problem. If for any of the following reasons, the provided code fails to execute, the participant will repeat the process for a maximum number of 7 trials until

Table 3
Average ChatGPT and Human Scores on the IEEEExtreme problems per Programming Language.

Language	Avg ChatGPT Score	Avg Human Performance
Python 3	9.06	44.5
Java 7	11.46	44.5
C++	7.67	44.5

either the code works perfectly, or the number of trials has exhausted. To this end, several prompts would be used to improve the quality of results generated by chatGPT. The reasons include:

- Incomplete code produced
- Compile Errors
- Runtime Errors
- Memory Limit Exceeded
- The execution time limit Exceeded
- Failing test cases

After a series of preliminary tests, we opted for the following unified prompts to be used by all participants to improve the quality of the response generated:

- “Provide a complete code for this problem using [language]”: this prompt is used in the first trial when no time/memory constraints were imposed in the IEEEExtreme problem statement.
- “Provide an optimized code using [language] that runs the program in a minimum time of x minutes and memory limitations of y Megabytes”: this prompt is used in the first trial when time/memory constraints were imposed.
- “Following up on this code, improve it to solve this test case: [provided test case with the expected output generated for certain input]”: this prompt is used in the following trials when the first trial did not pass all the test cases.

The participant runs the code generated by ChatGPT on the platform and records the success rate, which includes the number of passed test cases and the maximum score earned for each problem. This procedure is replicated for all problems, utilizing all three programming languages. The results are recorded by all participants in a shared data repository. The data collected include:

- Competition edition
- Problem title identifier
- Difficulty level of the problem
- Language used (C++, Java or Python)
- Number of trials/iterations
- Scores earned by chatGPT generated code
- Average Human performance for the problem

Data is gathered from all five editions of the IEEEExtreme competition, encompassing a total of 102 problems.¹

For every problem set, code is generated using each of the three programming languages, with at least three iterations performed. Participants conducted three to seven trials per iteration to complete the execution of the generated code. Since ChatGPT’s behavior has changed over time [48], it is important to note that our experiments were conducted from April 5 to April 30, 2023. The data is analyzed to generate conclusions for this study. These are presented in the next sections.

4. Discussion

4.1. Interpretation of the results

Table 3 shows the average score that ChatGPT achieved on the set of programming problems for Python 3, Java 7, and C++ 11, as well as the average human performance on the same set of problems. It appears that ChatGPT’s average score is significantly (from 3.9 to 5.8 times) lower than the average human performance for all three languages, which shows that there is still a large room for improvement in ChatGPT’s programming abilities. It is noteworthy that ChatGPT’s performance varies among the three programming languages, with Java 7 showing the highest average score, followed by Python 3 and then C++. This observation may suggest that the size and quality of available learning materials for each language in ChatGPT’s dataset are not equal. Indeed, Java has been the most widely used language for many years and has extensive documentation, which could have contributed to ChatGPT’s better performance on problems written in Java 7 compared to Python 3 and C++.

¹ <https://www.kaggle.com/datasets/riotulab/chatgpt-evaluation-on-ieeeextreme-competitions>.




Fig. 1. Comparison of average ChatGPT and human scores by problem complexity, in the IEEEExtreme competition.

Table 4
Average ChatGPT and Human Scores on the IEEEExtreme problems per programming language and complexity category.

Language	Complexity	Avg ChatGPT Score	Avg Human Score
Python 3	Easy	22.17	51.96
	Hard	8.69	46.70
	Medium	1.53	38.99
	Advanced	3.19	21.89
Java 7	Easy	30.06	51.96
	Hard	9.21	46.70
	Medium	4.24	38.99
	Advanced	0.00	21.89
C++	Easy	22.02	51.96
	Hard	4.35	46.70
	Medium	4.64	38.99
	Advanced	0.00	21.89

Fig. 1 shows the average ChatGPT and human performances for programming problems categorized by their complexity level. The complexity levels are Easy, Hard, Medium, and Advanced. As expected, when the complexity level of the problems increases, both ChatGPT’s average score and human performance significantly decrease. However, the decrease is much sharper for ChatGPT. Its score is 23 times lower for the Advanced category compared to the Easy category, while this decrease is only 2.4 times for human performance. It should be noted that the categories ‘Hard’ and ‘Medium’ used by IEEEExtreme competition may not be accurate indicators of problem difficulty, as both humans and ChatGPT demonstrate significantly better performance in the ‘Hard’ category compared to the ‘Medium’ category. This also highlights the subjective character of this categorization. On another hand, the correlation coefficient between ChatGPT’s and human scores is low (0.21), which indicates that the easiest programming problems for human programmers are not necessarily the easiest for ChatGPT to solve, and vice versa. This lack of correlation between ChatGPT’s and human scores could be due to various factors, such as differences in problem-solving approaches and strategies used by ChatGPT and human programmers, variations in the level of programming expertise, and the type and complexity of the problems presented to them.

Table 4 breaks down further these results per complexity levels. We notice again that ChatGPT performs better on Java for all complexity levels except the ‘Advanced’ category. For this category, all the tests on all problems completely failed except one test on the “Finite Domain Constraints” problem from Xtreme9 that gave partial success (12.74%) only on Python. Therefore, we cannot draw a general conclusion from this single partial success. On another hand, the average human score presented in this table is the same for all three languages because it was provided by the IEEEExtreme website as an overall average over all programming languages, and no average scores per programming language were available.

Fig. 2 shows the complete score distributions of ChatGPT and average human programmers. This figure demonstrates the marked superiority of average human programmers over ChatGPT, with ChatGPT obtaining a null score in the large majority of cases (72%), while only 10.0% of cases correspond to an average human performance less than 10%. Fig. 3 compares the sunburst charts of ChatGPT and human scores per programming language and complexity level. The color of inner sectors (representing programming languages and complexity categories) corresponds to the average colors of outer sectors belonging to them. The darker the color,

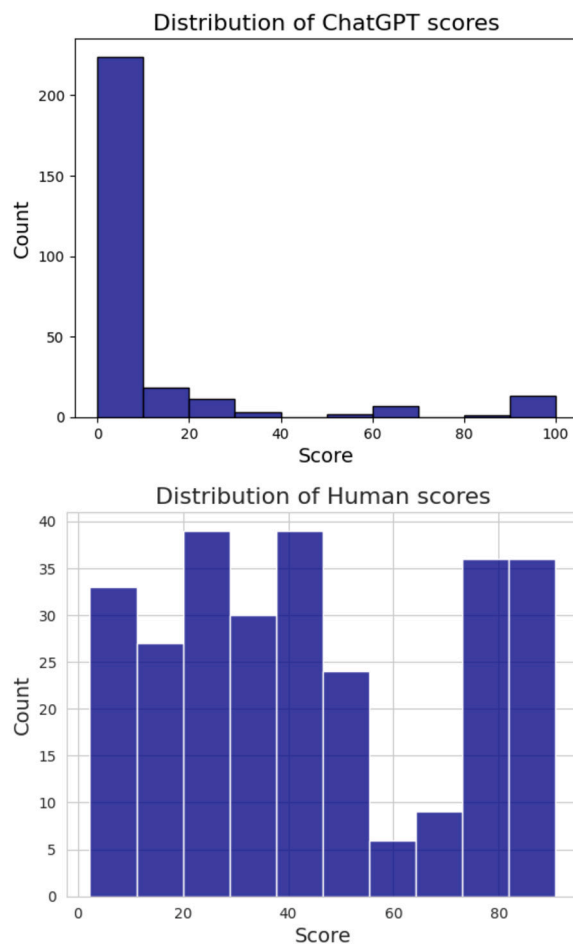


Fig. 2. Histogram of the Distribution of ChatGPT scores (top) and human programmers' average scores (bottom).

the better the results. This figure provides additional evidence of the dominance of average human programmers over ChatGPT in almost all tested cases.

To get an idea about the progress achieved in GPT-4 compared to GPT-3.5, in terms of programming capabilities, we tested their performance on a representative set of 6 problems using the Python 3 language. The results are presented in Table 5. GPT-4 showed a slight improvement in one problem (“Counting Molecules”) with an average score increasing from 65% to 70% (but still lower than average human score), and a clear improvement in another problem (“Painter’s Dilemma”) where it went from complete failure to complete success. However, for the remaining 4 problems, both GPT-3.5 and GPT-4 obtained a score of zero. For the “Painter’s Dilemma”, which is an optimization problem, we also prompted ChatGPT to generate C++ and Java solutions. Both tests yielded a 100% success in GPT-4, compared to 4.76% and 0%, respectively, in GPT-3.5. These results indicate that the improvement in GPT-4 programming abilities, compared to GPT-3.5 is limited to specific types of problems.

4.2. Limitations of ChatGPT in programming tasks

Based on the results of our experiments, we can draw several general conclusions about the limitations of ChatGPT in programming tasks. First, ChatGPT’s performance on programming tasks is significantly lower than that of an average human programmer, indicating that there is still a way to go before ChatGPT can fully match human intelligence in programming. This is especially true for more complex problems, where the performance gap between ChatGPT and humans is even more significant. This suggests that ChatGPT still has limitations in understanding and solving complex programming problems that require high-level reasoning and expertise. Second, the lack of correlation between ChatGPT’s and human scores indicates that the easiest programming problems for human programmers are not necessarily the easiest for ChatGPT to solve, and vice versa. This suggests that ChatGPT may have limitations in problem-solving approaches and strategies that differ from those of human programmers. Finally, although there have been some improvements in the GPT-4 compared to GPT-3.5 in terms of programming capabilities, there is still a significant performance gap between ChatGPT and human programmers, especially for more complex problems. This suggests that there are still limitations in the current state-of-the-art language models for programming tasks and that further research and development are needed to bridge the gap between ChatGPT’s performance and that of human programmers.

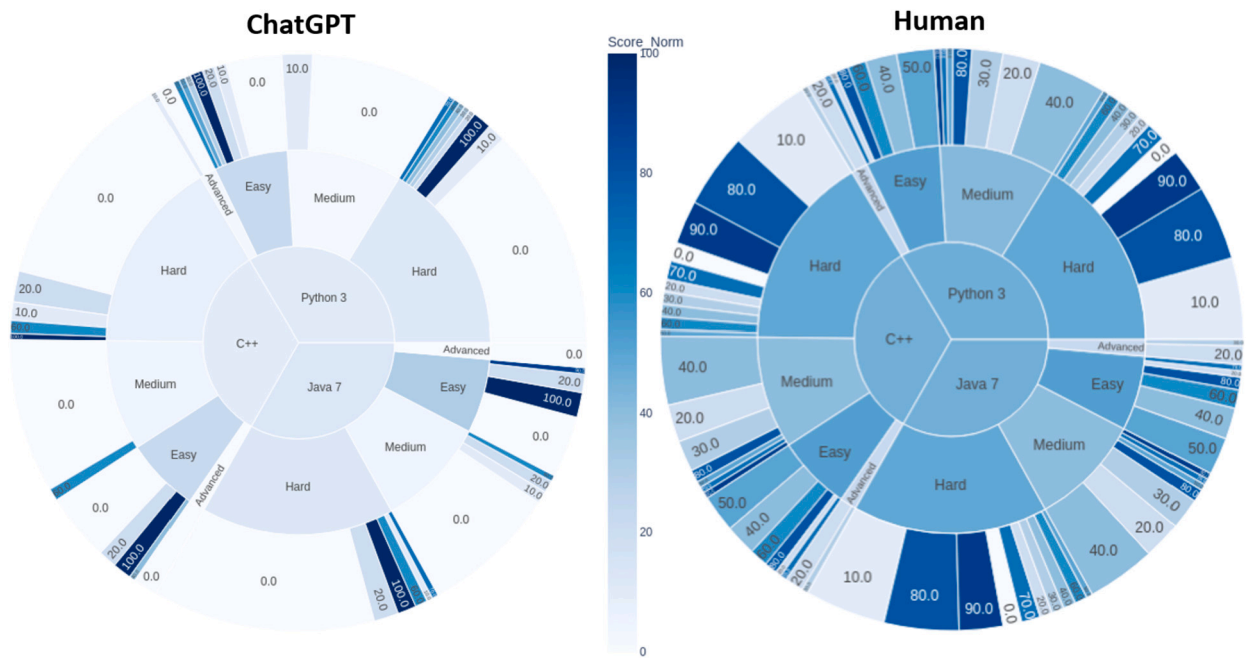


Fig. 3. Sunburst charts comparing programming language proficiency scores between ChatGPT (left) and human programmers (right) across different complexity levels and programming languages. The scores in the outer sectors have been rounded to the nearest 10.

Table 5
Comparison of GPT-3.5 and GPT-4 average scores on a set of selected problems, using Python 3 language.

Problem Title	Extreme Edition	Complexity	GPT-3.5 Score	GPT-4 Score	Human score
Back to Square 1	Xtreme8	Medium	0.0	0.0	77.00
Prediction Games	Xtreme9	Medium	0.0	0.0	15.38
Counting Molecules	Xtreme10	Hard	65.0	70.0	84.70
Painter's Dilemma	Xtreme10	Hard	0.0	100.0	80.39
Travel Service	Xtreme16	Medium	0.0	0.0	38.89
My Treat	Xtreme16	Easy	0.0	0.0	58.49

In summary, while ChatGPT represents a significant breakthrough in language modeling, its limitations in programming tasks suggest that there is still much room for improvement. Further research and development are needed to improve ChatGPT's performance on programming tasks, especially for more complex problems, and to bridge the gap between ChatGPT's performance and that of human programmers.

4.3. Implications for AI development and applications

The implications of these results for AI development and applications in the programming field are significant. While ChatGPT and other language models have shown promise in natural language processing and generation, their limitations in complex programming tasks indicate that they may not be suitable for fully automated programming, at least not yet. However, they can still be useful for tasks such as the generation of simple programs, code completion, code summarization, and documentation generation.

To fully harness the potential of language models in programming, further research is necessary to develop models capable of comprehending and rationalizing code in a manner akin to human programmers. This will require a better understanding of the cognitive processes involved in programming and the ability to incorporate this knowledge into AI models. Additionally, more comprehensive and diverse datasets need to be developed that better capture the variety of programming tasks and languages used in real-world programming.

Overall, the limitations of ChatGPT in programming tasks highlight the need for continued research and development in AI and programming and the importance of understanding the strengths and limitations of AI models in different contexts.

5. Conclusion

Numerous studies have demonstrated the impressive performance of ChatGPT, which often rivals or even surpasses human capabilities in various tasks and domains. However, this paper presented an alternative perspective by showing a situation where human performance excels over ChatGPT in typical tasks that suit it, specifically in relatively complex computer programming. To evaluate this claim quantitatively, we used the IEEEExtreme Challenge competition as a benchmark, which offers a range of programming problems with varying levels of difficulty. We executed a diverse set of 102 challenges drawn from five IEEEExtreme editions, using three major programming languages: Python, Java, and C++. We then compared ChatGPT's score to the average score achieved by human competitors.

Our empirical analysis demonstrated that human programmers maintain a significant advantage over ChatGPT in certain aspects of problem-solving within the programming context. This paper offers critical insights into the potential areas of improvement for ChatGPT and other AI-based language models. The present experiments were conducted in April 2023, and it is worth acknowledging that potential future iterations or releases of ChatGPT may yield differing outcomes.

Future research could investigate the factors that enable humans to outperform ChatGPT in programming tasks, measure the precise influence of prompt formulation on the quality of code provided by ChatGPT, and explore ways to address the limitations of AI-based language models in this area, such as improving their understanding of programming languages and their ability to work with complex code structures.

CRedit authorship contribution statement

Anis Koubaa: Conceptualization, Data curation, Formal analysis, Methodology, Project administration, Resources, Validation, Visualization, Writing – original draft, Writing – review & editing. **Basit Qureshi:** Conceptualization, Data curation, Formal analysis, Methodology, Project administration, Resources, Validation, Visualization, Writing – original draft, Writing – review & editing. **Adel Ammar:** Methodology, Resources, Validation, Visualization, Writing – original draft, Writing – review & editing. **Zahid Khan:** Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. **Wadii Boulila:** Resources, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. **Lahouari Ghouti:** Resources, Software, Validation, Visualization, Writing – original draft, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The data is made publicly available at <https://www.kaggle.com/datasets/riotulab/chatgpt-evaluation-on-ieeeextreme-competitions>.

Acknowledgement

The authors would like to acknowledge the support of Prince Sultan University for funding the Article Processing Charges (APC) of this publication.

References

- [1] W.X. Zhao, K. Zhou, J. Li, T. Tang, X. Wang, Y. Hou, Y. Min, et al., A survey of large language models, arXiv preprint, arXiv:2303.18223, 2023.
- [2] H. Alotaibi, H.S. Al-Khalifa, D. AlSaeed, Teaching programming to students with vision impairment: impact of tactile teaching strategies on student's achievements and perceptions, *Sustainability* 12 (13) (2020) 5320.
- [3] J. Qiu, L. Li, J. Sun, J. Peng, P. Shi, R. Zhang, Y. Dong, et al., Large AI models in health informatics: applications, challenges, and the future, *IEEE J. Biomed. Health Inform.* (2023).
- [4] S. Wu, O. Irsoy, S. Lu, V. Dabrovolski, M. Dredze, S. Gehrmann, P. Kambadur, D. Rosenberg, G. Mann, BloombergGPT: a large language model for finance, arXiv preprint, arXiv:2303.17564, 2023.
- [5] A. Koubaa, W. Boulila, L. Ghouti, A. Alzahem, S. Latif, Exploring ChatGPT capabilities and limitations: a critical review of the NLP game changer, Preprints, <https://doi.org/10.20944/preprints202303.0438.v1>, 2023.
- [6] R. Fatima, N. Samad Shaikh, A. Riaz, S. Ahmad, M.A. El-Affendi, K.A. Alyamani, M. Nabeel, J. Ali Khan, A. Yasin, R.M.A. Latif, A natural language processing (nlp) evaluation on Covid-19 rumour dataset using deep learning techniques, *Comput. Intell. Neurosci.* (2022).
- [7] S.L. Marie-Sainte, N. Alalayani, S. Alotaibi, S. Ghouzali, I. Abunadi, Arabic natural language processing and machine learning-based systems, *IEEE Access* 7 (2018) 7011–7020.
- [8] B. Guo, X. Zhang, Z. Wang, M. Jiang, J. Nie, Y. Ding, J. Yue, Y. Wu, How close is ChatGPT to human experts? Comparison corpus, evaluation, and detection, arXiv preprint, arXiv:2301.07597, 2023.
- [9] Y. Bang, S. Cahyawijaya, N. Lee, W. Dai, D. Su, B. Wilie, H. Lovenia, Z. Ji, T. Yu, W. Chung, et al., A multitask, multilingual, multimodal evaluation of chatgpt on reasoning, hallucination, and interactivity, arXiv preprint, arXiv:2302.04023, 2023.
- [10] M. Ortega-Martín, Ó. García-Sierra, A. Ardoiz, J. Álvarez, J.C. Armenteros, A. Alonso, Linguistic ambiguity analysis in chatgpt, arXiv preprint, arXiv:2302.06426, 2023.
- [11] S. Liu, A.P. Wright, B.L. Patterson, J.P. Wanderer, R.W. Turer, S.D. Nelson, A.B. McCoy, D.F. Sittig, A. Wright, Assessing the value of chatgpt for clinical decision support optimization, medRxiv, 2023, pp. 2023–02.

- [12] M. Sallam, Chatgpt utility in healthcare education, research, and practice: systematic review on the promising perspectives and valid concerns, *Healthcare* 11 (6) (2023) 887, MDPI.
- [13] A. Alzahem, S. Latif, W. Boulila, A. Koubaa, Unlocking the potential of medical imaging with chatgpt's intelligent diagnostics, arXiv preprint, arXiv:2305.07429, 2023.
- [14] L. Iftikhar, et al., Docgpt: impact of chatgpt-3 on health services as a virtual doctor, *EC Paediatr.* 12 (1) (2023) 45–55.
- [15] K. Cheng, Z. Li, Q. Guo, Z. Sun, H. Wu, C. Li, Emergency surgery in the era of artificial intelligence: chatgpt could be the doctor's right-hand man, *Int. J. Surg.* 109 (6) (2023) 1816–1818.
- [16] K. Cheng, Z. Sun, Y. He, S. Gu, H. Wu, The potential impact of chatgpt/gpt-4 on surgery: will it topple the profession of surgeons?, *Int. J. Surg.* (2023) 10–1097.
- [17] K. Cheng, Q. Guo, Y. He, Y. Lu, S. Gu, H. Wu, Exploring the potential of gpt-4 in biomedical engineering: the dawn of a new era, *Ann. Biomed. Eng.* (2023) 1–9.
- [18] P. Hacker, A. Engel, M. Mauer, Regulating chatgpt and other large generative ai models, arXiv preprint, arXiv:2302.02337, 2023.
- [19] S. Fergus, M. Botha, M. Ostovar, Evaluating academic answers generated using chatgpt, *J. Chem. Educ.* (2023).
- [20] D.R. Cotton, P.A. Cotton, J.R. Shipway, Chatting and cheating: ensuring academic integrity in the era of chatgpt, in: *Innovations in Education and Teaching International*, 2023, pp. 1–12.
- [21] S. Frieder, L. Pinchetti, R.-R. Griffiths, T. Salvatori, T. Lukasiewicz, P.C. Petersen, A. Chevalier, J. Berner, Mathematical capabilities of chatgpt, arXiv preprint, arXiv:2301.13867, 2023.
- [22] A. Tlili, B. Shehata, M.A. Adarkwah, A. Bozkurt, D.T. Hickey, R. Huang, B. Agyemang, What if the devil is my guardian angel: chatgpt as a case study of using chatbots in education, *Smart Learn. Environ.* 10 (1) (2023) 15.
- [23] B. Qureshi, Exploring the use of chatgpt as a tool for learning and assessment in undergraduate computer science curriculum: opportunities and challenges, arXiv preprint, arXiv:2304.11214, 2023.
- [24] J. Jeon, S. Lee, Large language models in education: a focus on the complementary relationship between human teachers and chatgpt, in: *Education and Information Technologies*, 2023, pp. 1–20.
- [25] H. Ji, I. Han, Y. Ko, A systematic review of conversational ai in language education: focusing on the collaboration with human teachers, *J. Res. Technol. Educ.* 55 (1) (2023) 48–63, <https://doi.org/10.1080/15391523.2022.2142873>, Online.
- [26] S.A. Prieto, E.T. Mengiste, B. García de Soto, Investigating the use of chatgpt for the scheduling of construction projects, *Buildings* 13 (4) (2023) 857.
- [27] M. Dowling, B. Lucey, Chatgpt for (finance) research: the bananarama conjecture, *Finance Res. Lett.* 53 (2023) 103662.
- [28] A.B. Kocaballi, Conversational ai-powered design: chatgpt as designer, user, and product, arXiv preprint, arXiv:2302.07406, 2023.
- [29] C. Qin, A. Zhang, Z. Zhang, J. Chen, M. Yasunaga, D. Yang, Is chatgpt a general-purpose natural language processing task solver?, arXiv preprint, arXiv:2302.06476, 2023.
- [30] A. Kashefi, T. Mukerji, Chatgpt for programming numerical methods, arXiv preprint, arXiv:2303.12093, 2023.
- [31] H. Liu, R. Ning, Z. Teng, J. Liu, Q. Zhou, Y. Zhang, Evaluating the logical reasoning ability of chatgpt and gpt-4, arXiv preprint, arXiv:2304.03439, 2023.
- [32] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, V. Stoyanov, Roberta: a robustly optimized bert pretraining approach, arXiv preprint, arXiv:1907.11692, 2019.
- [33] H. Tian, W. Lu, T.O. Li, X. Tang, S.-C. Cheung, J. Klein, T.F. Bissyandé, Is chatgpt the ultimate programming assistant—how far is it?, arXiv preprint, arXiv:2304.11938, 2023.
- [34] Y. Hu, U.Z. Ahmed, S. Mechtaev, B. Leong, A. Roychoudhury, Re-factoring based program repair applied to programming assignments, in: *2019 34th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, IEEE, 2019, pp. 388–398.
- [35] N.M.S. Surameery, M.Y. Shakor, Use chat gpt to solve programming bugs, *Int. J. Inf. Technol. Comput. Eng. (IJITC)* (ISSN 2455-5290) 3 (01) (2023) 17–22.
- [36] S.R. Piccolo, P. Denny, A. Luxton-Reilly, S. Payne, P.G. Ridge, Many bioinformatics programming tasks can be automated with chatgpt, arXiv preprint, arXiv:2303.13528, 2023.
- [37] S. Biswas, Role of chatgpt in computer programming: chatgpt in computer programming, *Mesop. J. Comput. Sci.* 2023 (2023) 8–16.
- [38] E. Chen, R. Huang, H.-S. Chen, Y.-H. Tseng, L.-Y. Li, Gptutor: a chatgpt-powered programming tool for code explanation, arXiv preprint, arXiv:2305.01863, 2023.
- [39] L. Avila-Chauvet, D. Mejía, C.O. Acosta Quiroz, Chatgpt as a support tool for online behavioral task programming, Available at SSRN 4329020, 2023.
- [40] Ieeextreme competition website, [Online]. Available: <https://ieeextreme.org/>, 2023.
- [41] The international collegiate programming contest, [Online]. Available: <https://icpc.global/>, 2023.
- [42] Google code jam, [Online]. Available: <https://codingcompetitions.withgoogle.com/codejam>, 2023.
- [43] Meta hackercup, [Online]. Available: <https://www.facebook.com/codingcompetitions/hacker-cup>, 2023.
- [44] International olympiad in informatics (ioi), [Online]. Available: <https://ioinformatics.org/>, 2023.
- [45] Cs academy platform, [Online]. Available: <https://csacademy.com/ieeextreme>, 2023.
- [46] Hackerrank platform, [Online]. Available: <https://www.hackerrank.com/ieeextreme7>, 2023.
- [47] Hackerrank practice community website for ieeextreme competition, [Online]. Available: <https://www.hackerrank.com/contests/ieeextreme-challenges/challenges>, 2023.
- [48] L. Chen, M. Zaharia, J. Zou, How is chatgpt's behavior changing over time?, arXiv preprint, arXiv:2307.09009, 2023.