Explore Public's Perspectives on Generative AI in Computer Science (CS) Education: A Social Media Data Analysis

Sunggyeol Oh
Department of Computer Science
Virginia Tech
Blacksburg, VA, USA
sunggyeol@vt.edu

Yi Cao
Department of Engineering Education
Virginia Tech
Blacksburg, VA, USA
caoyi@vt.edu

Andrew Katz
Department of Engineering Education
Virginia Tech
Blacksburg, VA, USA
akatz4@vt.edu

Jialu Zhao
Stanford Graduate School of Education
Stanford University
Stanford, CA, USA
jialuz@stanford.edu

***Abstract*—This research-to-practice full paper aims to analyze the public's comments on generative artificial intelligence (GAI) in computer science (CS) education, by the BERT-based model and Large Language Model (LLM) approaches to sentiment analysis and contextualize the results within broader educational and technological landscapes. Artificial intelligence (AI) has played a crucial role in advancing technical development throughout many areas. Evidence points toward the likelihood of major developmental breakthroughs unfolding soon in those sectors. Education is one such area. While there is certainly a possibility for hype and unfulfilled promises, the advent of available GAI platforms, such as ChatGPT, has caused a surge of scholarly interest in the impact of these technologies on CS education. Amid the growing debate, both the potential benefits and concerns of GAI in this sector are increasingly coming to the fore as people grapple with the tradeoffs associated with these technologies when applied in education settings. One can imagine the range of conversations around the topic, but that is difficult to use as input for policymakers and administrators without a more concrete understanding. To wit, there remain open questions about which benefits and concerns people tend to focus on when discussing GAI in education.**

This large-scale qualitative study addresses that gap by exploring the public's perspectives on GAI in CS education. We engage in this work by collecting and analyzing data from social media platforms, specifically Reddit comments. The social media dataset was analyzed using machine learning (ML) techniques to identify topics based on sentiment analysis. The study's objective was to document and characterize the public's perspectives concerning the general characteristics of GAI, its features related to learning, and its usability in educational settings. Through sentiment analysis using Large Language Models (LLM), the study revealed an overall positive public perception toward using generative AI in CS education, with over 57% of comments being favorable, while also identifying prominent topics of interest and concerns, such as the potential benefits of personalized learning support and automated grading, as well as issues like academic dishonesty, perpetuation of biases, over-reliance on AI hindering critical thinking, displacement of human instructors, and the need for updated curricula. The insights gleaned from the analysis will be instrumental in computing educators gaining a more profound comprehension of GAI's role in education and the subsequent development of GAI-enriched curricula.

***Keywords*—Generative AI, computer science education, social media dataset, sentiment analysis**

I. Introduction

The rapid advancement of AI technologies has disrupted numerous sectors, introducing both opportunities and challenges. One domain poised for potential transformation is education, particularly in CS. The recent emergence of available GAI platforms, such as ChatGPT, has sparked intense interest and debate regarding their impact on learning and teaching practices. These GAI systems, powered by large language models, possess the ability to generate human-like text, code, and even multimedia content based on natural language prompts. This capability holds significant promise for enhancing educational experiences by providing personalized learning support, automating grading and feedback processes, and fostering creative problem-solving skills among students [1-3].

However, the integration of GAI into CS education is not without concerns. Critics raise issues such as the potential for academic dishonesty, the perpetuation of biases inherent in training data, and the risk of over-reliance on AI systems, potentially hindering students' critical thinking and problem-solving abilities [4]. Additionally, there are concerns about the potential displacement of human instructors and the need for updated curricula and assessment methods to effectively harness the power of GAI while maintaining academic integrity.

GAI has been widely applied across various disciplines in higher education. As the discourse around GAI's role in CS education intensifies, it is crucial to understand the public's perspectives of GAI's role in CS education, which may include various stakeholders, like students, and educators. Social media platforms have emerged as vibrant forums where individuals express their views, share experiences, and engage in discussions on a wide range of topics, including emerging technologies and their implications for education [5]. Therefore, analyzing the comments from social media can provide us with many useful clues about how people conceptualize the GAI's role in CS education. Reddit, as one of the larger social media platforms in the US, is notable for its niche communities and discussion-based format, which makes it a fit data source to collect the public-specific perspective about GAI in CS education.

This study aims to explore the public's perspectives on GAI in CS education by leveraging the wealth of data available on social media platforms, specifically Reddit comments. By

2024 IEEE Frontiers in Education Conference (FIE) | 979-8-3503-5150-7/24/$31.00 ©2024 IEEE | DOI: 10.1109/FIE61694.2024.10893102

employing ML techniques, such as text embedding and clustering algorithms, combined with sentiment analysis, the researchers seek to identify and characterize the prominent topics, concerns, and perceived benefits associated with the integration of GAI into CS curricula.

II. Literature Review

In this era of relentless technological advancement that is reshaping human lifestyles, AI has emerged as a revolutionary force. AI has pervasively infiltrated every corner of our daily lives, silently yet profoundly altering our work methods, communication styles, and the way we control the world. In the research by Baidoo and Ansah [6], GAI is described as an unsupervised or semi-supervised ML framework that seamlessly generates artificial creations by analyzing existing digital content such as videos, images/graphics, texts, and audio. GAI has quickly become a focal point in academia, significantly increasing research activity [7-9].

Studies have explored GAI's potential in personalized learning, enhancing student engagement, predicting student performance, intelligent tutoring systems, and enhancing their educational experiences and resources. For instance, Dai et al. [10] explore the transformative role of ChatGPT and GAI in higher education, conceptualizing these technologies as student-driven innovations that enrich educational resources and experiences, while also highlighting the need for collaborative stakeholder efforts to address emerging challenges in training, curriculum development, and technology governance. Similarly, Hadi et al. [11] examined the development and impact of LLM, highlighting their transformative role in various fields including education, emphasizing GAI's capability to personalize learning by adapting and responding to individual linguistic needs. Rahman and Watanobe [4] also demonstrated GAI's broad prospects in educational fields, particularly in course planning, personalized learning support, rapid assessment, and addressing learners' queries. Additionally, research has shown GAI's (especially ChatGPT's) potential to revolutionize student engagement and interactive learning, although it has raised concerns about academic integrity and the need for strategic implementation to mitigate risks associated with its use in educational settings [2, 12].

GAI systems powered by large language models possess the remarkable ability to generate human-like text, code, and multimedia content based on natural language inputs [6]. This capability shows significant promise for enhancing educational experiences in CS. GAI could provide personalized learning support, automate grading and feedback, and foster students' creative problem-solving abilities related to programming and computational thinking [4, 5].

However, the integration of GAI into CS curricula raises important concerns as well. Critics point to issues like the potential for academic dishonesty if students leverage GAI for cheating [2, 12]. There are also concerns about perpetuating societal biases encoded in the training data of these models [3]. Some argue an over-reliance on GAI could hinder the development of critical thinking and problem-solving skills so vital for CS students [13]. The potential displacement of human instructors and the need to overhaul curricula and assessments to properly incorporate GAI while maintaining academic integrity has also been discussed [14].

As this debate intensifies, researchers have examined GAI's impacts specifically in the context of CS education. Studies show that GAI tools like ChatGPT can enrich the learning experience by generating code, explaining concepts, and creating programming exercises that foster deeper engagement [15, 16]. AI models have outperformed students on introductory coding problems while enhancing computational thinking [16]. Language models can provide multi-faceted explanations for code snippets to improve understanding of complex programming concepts [17].

At the same time, work has highlighted the importance of collaborative efforts to shape GAI's developmental trajectory in CS education responsibly [14]. There are open questions about cultivating problem-solving abilities, ethical considerations of AI systems, and the need for transparency and explainability as effective computing systems handle multimodal data [13, 18].

In industry, GAI is driving innovation and productivity across sectors like customer service, healthcare, and education technology [19, 20]. However, challenges like data privacy, bias, and workforce impacts require strategies to ensure responsible adoption [21]. GAI allows personalized learning experiences through intelligent tutoring, voice interaction, and customization based on student performance data [22, 23]. While GAI shows promise to revolutionize CS education, crucial gaps remain in understanding the specific benefits and concerns the public focuses on with these technologies in learning environments. This study aims to address that gap through a large-scale analysis of public perspectives shared on social media platforms about GAI's characteristics, features for learning, and usability in CS education settings.

The study highlights the transformative potential and inherent challenges of integrating GAI, particularly ChatGPT, into CS education. While these tools offer significant opportunities for personalized learning and efficiency, concerns about bias, academic integrity, and the erosion of critical thinking skills persist [1-3]. The ability of AI to generate tailored educational content and engage students in deeper learning processes presents a promising avenue to enhance educational outcomes [4, 5].

III. Methods

This process flow shows the general methods of data collection, preprocessing, and analysis in this study. Initially, data in the form of comments from Reddit were collected. This data underwent three main operations. Operation 1 involved preprocessing the raw data by removing duplicates and filtering irrelevant comments using regular expressions and keyword matching. Operation 2 utilized LLM to further classify each comment as related or unrelated to the topic of GAI in CS education. In Operation 3, the filtered and classified data underwent sentiment analysis. For sentiment analysis, a BERT-based model pre-trained on tweet data and an LLM with prompts were employed on the same datasets to interpret positive, negative, and neutral sentiments expressed in the comments. Decision points after each operation allowed for iterative refinements, such as re-classifying ambiguous comments or adjusting the filtering criteria. The final output of this process yielded insights into the overall public sentiment toward GAI in CS education. The overall process is illustrated in Fig. 1.

Fig. 1. Overview of the Data Analysis Process

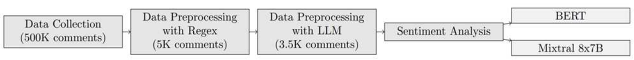

A. Data Collection

Since our study aimed to study the public perceptions of GAI in CS education, we collected data from the social media platform Reddit. We selected the 16 largest college Subreddits [24] and 42 CS-related Subreddits for this study. Using the Reddit Application Programming Interface (API), we gathered all comments containing keywords from these 58 Subreddits. 492,065 comments were collected from those Subreddits using these keywords. The list of subreddits and keywords used to collect data is in Appendix A.

B. Data Preprocessing

Before analyzing the dataset using ML techniques, we preprocessed the dataset by deleting duplicate data from 500K comments. The dataset was then filtered using regular expressions (Regex) in Python. Filtering using Regex consisted of 2 stages. Initially, we applied keywords related to GAI and CS education and filtered comments containing those keywords. Despite using filters for both CS education and GAI keywords, some comments about irrelevant topics such as job interviews and college admissions remained. These off-topic comments were filtered out using keywords that are not related to the research topic to refine the dataset, leaving 5,354 comments.

By randomly selecting and examining the 5,354 filtered comments, we were able to see that some datasets still contained topics that deviated from the research topic. Thus, before analyzing public perception through sentiment analysis, we used the Mixtral 8x7B model to classify each comment as either relevant or irrelevant to our topic of interest (i.e., GAI in CS education). These irrelevant comments were filtered out to clean the dataset, with 3,445 comments. The prompts used to classify datasets are listed in Appendix B.

C. Data Analysis

To analyze the social media datasets, we used sentiment analysis approaches to analyze this study. The sentiment analysis was done with two models. A BERT-based model pre-trained with tweet data and an LLM with prompts were used to interpret sentiments.

The first step of our data analysis was sentiment analysis. Using the Pysentimiento model, each comment was labeled with the sentiment. The base model of the Pysentimiento model is RoBERTa. It is an open-source Python library that was trained with 5K tweets [25]. We labeled our dataset with 3.5K sentiments classified as POS, NEG, and NEU.

The next step was using LLM to classify the same datasets. We used the Mixtral 8x7B model to classify the same comments into Positive, Negative, and Neutral.

The prompt we used for this step starts as:

"As a world-class sentiment analyst, your expertise lies in deeply analyzing texts to detect expressed sentiments, pinpointing whether they are positive, negative, or neutral. When reading comments, you excel in distinguishing nuanced sentiments, clearly identifying what aspects evoke specific feelings."

The full prompt used to classify datasets is listed in Appendix C.

IV. Results

A. Sentiment Analysis: BERT-Based Analysis vs. LLM-Based Approach

The results present sentiment analysis results performed on Reddit datasets using two different methods, BERT-based and LLM-based models.

Both models categorized the sentiments of comments into three categories: positive, neutral, and negative. The distribution of these emotions is shown in Figures 2 and 3. Specifically, the BERT-based model categorized 29.32% of the comments as positive, 19.85% as negative, and 50.83% as neutral.

Fig. 2. Sentiment Analysis on Reddit Dataset using BERT-based Model

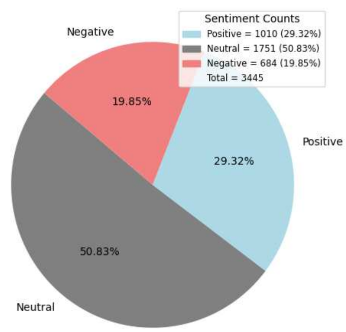

In the sentiment analysis conducted using LLM, the Mixtral 8x7b model showed a different distribution. In the LLM-based model, those with more comments (57.16%) were classified as positive, 20.35% as neutral, and 22.50% as nega

tive. These results are shown in Fig. 3.

Fig.3. Sentiment Analysis on Reddit Dataset Using Mixtral 8x7B

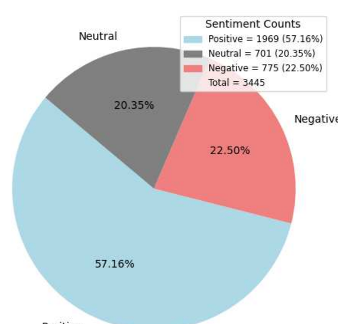

B. Overall Results of Sentiment Analysis

The comparison of results between BERT-based and LLM-based models highlights that when classifying comments, each model has a different approach to interpreting comments. This demonstrates the difference between sentiment analysis obtained through BERT-based models and LLM-based approaches in interpreting the semantic context of textual data. BERT-based models are conservative and classify many comments as neutral, while LLM-based analysis with our prompt is more assertive in classifying comments as positive.

C. Edge Cases in Two Different Sentiment Analysis

Comments that BERT classified neutral but positively by LLM were paraphrased by GPT-4 and its sentiment analysis results are compared with each other. The following edge cases show the difference in emotional classification between the two models:

1) *Paraphrased Comment 1*: I ended up learning Python because I needed to correct some ChatGPT code. Experienced programmers will find it more beneficial than beginners, as beginners lack the fundamental coding knowledge necessary to comprehend unfamiliar code. If you have some knowledge of any programming language, there are hardly any drawbacks to using GPT. It's incredibly useful for figuring out which functions are available for specific tasks.

a) BERT Analysis: Neutral, probas={NEU: 0.572, POS: 0.314, NEG: 0.115}

b) LLM Analysis: Positive, The user expresses a positive sentiment towards Python and chat-GPT. They mention that they ended up learning Python as a consequence of needing to fix chat-GPT code, implying a positive experience with it. They also describe chat-GPT as amazing and helpful for understanding functions in code. However, they do mention that beginners might have difficulty with it, which introduces a slight element of negativity, but overall, the sentiment is positive due to the user's appreciation for Python and chat-GPT.

2) *Paraphrased Comment 2:* I began programming just two months ago, so my experience is quite limited. However, by asking a series of questions, I was able to guide ChatGPT into writing a complete Python program for a ping-pong game.

a) BERT Analysis: Neutral, probas={NEU: 0.773, POS: 0.221, NEG: 0.006}

b) LLM Analysis: Positive, Despite acknowledging being a novice at programming, the author expresses satisfaction and accomplishment in making ChatGPT create a ping-pong game program in Python. This positive sentiment is derived from the successful outcome of their interaction with the AI, which led to the creation of a full program.

3) *Paraphrased Comment 3:* I turn to ChatGPT for help when I'm stuck with an error or unsure about how to tackle something new. I find the explanations it provides to be much clearer than the official .Net documentation, which is helpful since I'm still quite new to .Net.

a) BERT Analysis: Neutral, probas={NEU: 0.492, POS: 0.490, NEG: 0.018}

b) LLM Analysis: Positive, The user has a positive sentiment towards ChatGPT, finding the explanations easier to understand than the official documentation. The user also mentions being new to .NET, which might indicate a slight challenge or unfamiliarity, but this does not negatively impact the overall positive sentiment expressed in the comment.

Three edge examples illustrate the differences in sentiment analysis between the BERT-based model and the LLM-based model. For example, in the first comment, a user discusses a positive experience of unintentionally learning Python while modifying the code generated by ChatGPT. The BERT-based model classified it as neutral with a 57.2% probability, while LLM classified the user's experience as overall positive based on the nuances of the comments. Similarly, in the second comment, users share positive experiences of making ping-pong games with Python through ChatGPT as programming beginners. BERT-based models were classified as neutral, while LLM classified them as positive based on user satisfaction and achievement. Finally, the third comment talks about the positive aspects of ChatGPT rather than the official .NET documentation. The BERT-based model classified it as neutral, with the margin between neutral (0.492) and positive (0.490) being very close. In contrast, the LLM-based model classified it as positive based on the overall sentiment expressed in the comment. These examples show that in comments containing mixed sentiments, LLM tends to detect user attitudes and find overall tones to classify more positives, while BERT-based models classify sentiments more conservatively.

D. Analysis of Public Perceptions Through LLM-Based Theme Identification

Sentiment analysis using two models, BERT and LLM, found different results on sentiment toward GAI in CS education. The BERT-based approach classified a small percentage of comments as positive compared to the LLM-based approach.

To identify how each model classified the themes, we separated the sentiment summaries by BERT and LLM and investigated the topics among the different methodologies. First, the classified comments from each model were merged into six separate documents according to the sentiment classification. Then, each document containing all the classified comments was divided into 10 chunks so that GPT-4o-mini could handle them at a time. For each chuck, GPT-4o-mini extracted three main themes, generating a total of 30 themes per sentiment category for each model. GPT-4o then analyzed these themes, extracting two key themes for each sentiment classification and providing a brief explanation for each theme.

Of the 30 themes for each emotion identified in the analysis using the BERT-based model, two key themes were extracted for each sentiment. Table 1 shows that the positive sentiments of the BERT-based model are mainly related to improving the learning experience, productivity, and efficiency. The neutral sentiments were about the role of GAI in education and the importance of critical thinking. The negative sentiments focus on the frequent generation of inaccurate responses from GAI and the reduction in the value of human technology due to its dependence.

TABLE I. BERT-BASED ANALYSIS THEMES

Sentiment	Themes
Positive	**Enhanced Learning Experience:** *Generative AI tools provide instant, personalized guidance, simplifying complex topics and improving understanding.*
	Increased Productivity and Efficiency: *AI tools automate repetitive coding tasks and debugging, significantly boosting developers' productivity.*
Neutral	**Generative AI in Education:** *Generative AI is valuable for personalized learning and feedback but should not replace fundamental learning.*
	Importance of Fundamentals and Critical Thinking: *AI can assist in coding, but solid programming fundamentals and critical thinking are crucial to avoid over-reliance.*
Negative	**Lack of Genuine Understanding in AI:** *Generative AI tools often produce incorrect or nonsensical outputs due to their inability to comprehend context and nuances, leading to mistrust in their reliability.*
	Diminished Value of Human Skills and Overreliance on AI: *Relying on AI tools can devalue human skills, promote superficial competence, and hinder genuine learning and critical thinking.*

LLM-based sentiment analysis also revealed key themes. Table 2 shows LLM-based models, like BERT-based models, rated themes of improved learning and understanding, increased productivity, and efficiency as the main topics of positive evaluation. Although LLM-based models rated comments more positively (57%) than BERT-based models (29%), weight differences are likely due to the different detailed nuances of the comments evaluated as the main themes of the comments were similar. Neutral topics included restrictions on improving education and critical engagement. Negative topics highlighted the excessive reliance on AI and the issue of AI-generated outcomes. A full list of topics identified in each sentiment is listed in Appendix D.

TABLE II. LLM-BASED SENTIMENT THEMES

Sentiment	Themes
Positive	**Enhanced Learning and Understanding:** *Generative AI tools significantly aid in learning and comprehending complex programming concepts, acting as personal tutors that provide immediate, tailored explanations and examples.*
	Increased Productivity and Efficiency: *These AI tools streamline coding tasks, automate repetitive processes, and boost productivity, allowing users to focus on higher-level problem-solving and creative aspects of development.*
Neutral	**Educational Enhancement and Productivity:** *Generative AI aids learning and boosts productivity in computer science by helping students understand complex concepts and generate code snippets efficiently*.
	Limitations and Critical Engagement: *AI tools require critical thinking and verification, as reliance on them without understanding programming fundamentals can result in errors and missed learning opportunities.*
Negative	**Overreliance on Generative AI:** *Commenters worry that reliance on generative AI tools without understanding their limitations can hinder the development of critical thinking and problem-solving skills.*
	Quality and Reliability Issues: *There is concern about the frequent production of subpar or erroneous AI-generated outputs, leading to mistrust in these tools' effectiveness for coding and education.*

V. DISCUSSION

This section reflects on the findings from both the BERT-based and Large Language Model (LLM) approaches to sentiment analysis, contextualizing the results within broader educational and technological landscapes.

The contrasting outputs from the BERT-based model and the LLM highlight different methodological perspectives on sentiment classification. The BERT-based model, with a conservative approach, labeled a significant portion of the comments as neutral, suggesting a cautious stance in sentiment interpretation. This conservatism is likely due to BERT's design to focus on neutrality when mixed sentiments are present. On the other hand, the LLM displayed a propensity to classify comments more assertively as positive, indicating its capacity to capture nuanced sentiments and user appreciation within the text.

The substantial difference in sentiment distribution between the two models emphasizes the importance of model selection in sentiment analysis. The BERT-based model's higher neutral classification (50.84%) contrasts sharply with the LLM's higher positive classification (57.16%). This divergence suggests that LLMs may be more sensitive to positive expressions in user comments, whereas BERT may miss subtle positive cues, leading to a more balanced but less enthusiastic sentiment distribution.

Examining the edge cases further illuminates the nuanced capabilities of the LLM in sentiment analysis. In comments where mixed emotions are present, the LLM successfully identifies underlying positive sentiments, which the BERT-based model often classifies as neutral. This distinction is crucial in understanding user perceptions, especially in educational contexts where positive reinforcement can significantly impact engagement and learning outcomes.

For instance, in the case of comments reflecting unexpected positive experiences with GAI, such as learning Python through ChatGPT, the LLM's positive classification aligns with the user's overall satisfaction and achievement. These insights demonstrate the LLM's effectiveness in capturing the holistic sentiment expressed by users, providing a more accurate reflection of public perception.

The overall positive sentiment toward GAI in CS education, as highlighted by the LLM analysis, suggests a general acceptance and appreciation of its role in enhancing learning experiences. Users frequently mention GAI's ability to boost productivity, provide immediate assistance, and facilitate deeper understanding through iterative interactions. These benefits underline the transformative potential of GAI in making CS education more accessible and efficient.

However, a significant portion of neutral sentiments, especially as identified by the BERT-based model, indicates cautious optimism among users. This neutrality points to the need for further development and discussion around the integration of GAI in education. Users' concerns about over-reliance on AI, potential hindrances to developing critical thinking skills, and ethical considerations must be addressed to ensure balanced and effective use of GAI in educational settings.

While the positive perceptions are encouraging, the challenges associated with GAI integration cannot be overlooked. Issues such as academic dishonesty, biases in AI training data, and the displacement of human instructors highlight the complexities of implementing GAI in education. These concerns necessitate comprehensive strategies to mitigate negative impacts while maximizing the benefits of GAI.

VI. Limitations

Theoretically, we admit that it would be difficult to extract direct conclusions from the results of sentiment analysis, with one method (BERT) about 29% exhibit a positive reaction; with the second method (Mixtral) 57% exhibit a positive reaction. It would be a limitation for our data analysis to use the sentiment analysis method but mixed research methods to analyze this type of data. However, we should provide more analysis approaches to triangulate our conclusions in our next steps.

VII. Conclusions

The data analysis yielded valuable insights into public perceptions of GAI in CS education settings. Sentiment analysis using both BERT-based and large language models revealed an overall positive sentiment, with over 57% of comments expressing a favorable outlook toward the use of GAI in educational contexts. However, a notable portion, around 20%, maintained a neutral stance, highlighting the need for continued discussion and development to better define GAI's role in education.

Conversely, key concerns raised pertained to issues of academic dishonesty, the perpetuation of societal biases encoded in training data, an over-reliance on AI potentially hindering critical thinking development, the displacement of human instructors, and the necessity of updated curricula and assessment methods to effectively incorporate GAI while maintaining academic integrity.

While the study highlights the transformative potential of GAI to enhance educational experiences and outcomes, it also underscores the inherent challenges and concerns that must be addressed. Future research should focus on longitudinal studies to assess the long-term impact of continuous AI interaction on students' learning behaviors, critical thinking, and problem-solving capabilities. Additionally, comprehensive investigations into ethical, privacy, and accessibility concerns are crucial to ensure responsible AI integration that bridges educational disparities rather than exacerbates them.

Regarding the preliminary findings from sentiment analysis, for the next step, we would focus on more longitudinal studies to evaluate the long-term effects of GAI on students' learning behaviors, critical thinking, and problem-solving skills. Additionally, addressing ethical, privacy, and accessibility issues is crucial to ensure that GAI serves as an inclusive tool that bridges educational gaps rather than widening them.

Appendixes

Appendix A. Keywords for Data Collection and Filtering

A. List of 56 Subreddits used in data collection

Universities: UIUC, Berkeley, aggies, gatech, UTAustin, OSU, ucf, UCSD, rutgers, Purdue, rit, UMD, uofm, ucla, ASU, VirginiaTech

Computer Science and Programming: AskComputerScience, ChatGPT, cscareerquestions, CSEducation, learnprogramming, OpenAI, programming, compsci, cpp, csharp, devops, dotnet, emberjs, Frontend, gamedev, golang, homelab, java, javascript, Kotlin, linux, MachineLearning, netsec, nextjs, node, opensource, ProgrammerHumor, ProgrammerTIL, Python, rails, reactjs, ReverseEngineering, ruby, rust, SideProject, sysadmin, technology, userexperience, vuejs, web_design, webdev

B. Keywords used to collect data from Subreddits

Keywords: ai, bard, chatgpt, computer science, copilot, gemini, generative ai, gpt, artificial intelligence, cs

C. Keywords used to filter datasets

Keywords: ai, bard, chatgpt, computer science, copilot, gemini, generative ai, gpt, artificial intelligence, cs

GAI Keywords: GPT, Generative AI, Copilot, GPT-4, GPT-3, GPT-3.5, GPT3, Bard, Gemini, Hallucination, Prompt Engineering, ChatGPT, GPT-like, LLMs, LLM, neural network, transformer, superintelligent

CS Keywords: Computer Science, CS, Coding, Programming, Algorithm, Data Structures, Machine Learning, Artificial Intelligence, Web Development, Software Engineering, Database, Network, Security, Operating System, Mobile App Development, Cloud Computing, Internet of Things, Big Data, Data Science, Natural Language Processing, Computer Vision, Robotics, Game Development, Virtual Reality, Augmented Reality, Blockchain, Cybersecurity, UI/UX Design, Front-end Development, Back-end Development, Full-stack Development, Data Analysis, Computer Graphics, Parallel Computing, Embedded Systems, programmer, software engineering, developer, development

Non-Related Keywords: Alumni, Career, Degree, Admission, rejected, Accepted, GRE, TOEFL, IELTS, Scholarship, Fellowship, Resume, Cover Letter, Interview, Offer, Salary, Negotiation, Promotion, Layoff, Resign, Quit

Appendix B. Prompt for Data Classification

Act as if you are an expert in computer science education and Generative AI. You specialize in analyzing comments and identifying whether the comments are related to Generative AI in computer science education. The comments provided to you are extracted from either computer science or AI-related subreddits. That means although even the comments directly mention AI or computer science education, it is likely to be talking about the same topic.

Your task is to classify each comment as either related to AI in computer science education or not. Follow these steps to complete your analysis:

1) Read the comment given to you in the <text> tag.

2) Determine whether the comment is related to AI in CS education based on its content.

Provide your classification of the comment as either related to Generative AI in CS education or not related to Generative AI in CS education. Your response should always start with "My classification: ". Do not provide explanations for your classification, only the classification itself.

Appendix C. Prompt for Sentiment Analysis

Act as if you are the world's best sentiment analyst. You specialize in analyzing texts and identifying sentiments expressed in those texts. You are especially skilled in identifying nuances in those sentiments by identifying the aspects that people are expressing positive, neutral, or negative sentiments about. I need your help. I have a collection of comments collected from Reddit. In the comments, there may be a combination of positive and negative sentiments expressed. Your task is to identify the sentiments expressed in the text and what those sentiments are being expressed about by following two steps. First, read the comments given to you in the <text> tag. Second, create an enumerated list in which you summarize each sentiment expressed and what it is being expressed about. For each item in your list, you should put the sentiment followed by the aspect about which that sentiment was expressed. If there are two instances of positive sentiments expressed about two separate things, then they should each have an entry in your list. For example, for the following instance from a movie review "the actors did a great job in the action scenes. I also love the soundtrack." Your response should look like:

1) *Positive:* actors did great job in action scenes

2) *Positive:* loved the soundtrack.

You can look for positive, negative, or neutral sentiments. If nothing in the text expresses a sentiment, you should just say "no sentiments identified". Your response should always start with "My expert analysis:". It is essential that you identify each instance of a sentiment expressed but that you also not make up things that are not in the text.

Appendix D. Comprehensive List of Identified Themes

A. *Positive Themes from BERT-based Sentiment Analysis*

Enhanced Learning and Problem Solving, Increased Productivity and Creativity, Future Workforce Integration and Computational Thinking, Fostering Collaboration and Innovation, Enhanced Productivity and Efficiency in Learning and Coding, AI as a Personalized Learning Assistant and Mentor, Fostering Creativity and Innovative Problem Solving, AI as an Enhancer of Learning and Skill Acquisition, Increased Productivity through Automation, Fostering Creativity and Collaborative Development, Increased Productivity in Coding Tasks, Valuable Companion for Problem-Solving, Enhanced Learning and Understanding of Coding Concepts, Productivity and Efficiency in Academic and Professional Work, Creative Collaboration and Support in Assignment Generation, Increased Productivity and Efficiency, Support for Collaboration and Problem Solving, Increased Efficiency and Productivity, Increased Productivity in Programming, Facilitating Creativity and Innovation, Enhancing Creativity and Problem-Solving in Game Development, Improving Debugging and Coding Efficiency, Democratizing Access to Programming Knowledge, Personalized Learning Experience, Enhancing Creative Outputs, Streamlined Educational Processes, Enhanced Learning Experiences, Increased Productivity for Developers, Accessibility and Collaboration, Real-World Application and Problem Solving

B. *Neutral Themes from BERT-based Sentiment Analysis*

Perceived Limitations of Generative AI, Evolving Roles of Educators and Students, New Skill Sets and Job Functions in AI-Enhanced Education, Aiding in Programming Efficiency, The Role of Generative AI as a Tool for Developers, The Importance of Human Oversight in Generative AI Outputs, Evolving Expectations in Software Development, Tool for Learning and Problem-Solving, Limited Utility in Complex Development Tasks, Encouraging Creativity and Innovation, Enhancement of Learning and Productivity, Efficiency and Productivity, Critical Thinking and Skill Development, Tool for Efficiency and Learning Enhancement, Collaboration Between AI and Human Skill, Varied Expectations and Use Cases, Complementing Human Creativity and Coding Skills, Practical Applications of Generative AI in Education, Challenges with Reliability and Trustworthiness of AI, The Role of AI as a Learning Aid, User Satisfaction and Perceived Benefits, Generational Divide in Adoption of Generative AI, Concerns About Impact and Misuse, Generational Differences in AI Usage, Skepticism and Concerns about AI Misuse, Support for Learning and Problem-Solving, Efficiency and Productivity Boost, Skepticism Towards Reliance on AI, Ethical Considerations in Academic Integrity, Tools for Collaborative Projects

C. *Negative Themes from BERT-based Sentiment Analysis*

Diminished Value of Human Skills, Potential Negative Impact on Education and Learning, Limitations of Generative AI in Understanding Context, Risk of Reduced Learning and Dependency, Quality of Output and Misinformation, Declining Utility of AI in Learning, Ethical and Security Concerns, Overreliance on AI and Loss of Critical Thinking Skills, Impediments to Learning and Understanding, Inaccuracies and Unreliability of AI Outputs, Increased Dependence on Human Expertise, Detrimental Effects on Learning and Understanding, Misleading Competence and Overreliance on AI, Technical Limitations and Inaccuracy Issues, Inaccuracy and Hallucination of AI Content, Ethical Concerns and the Erosion of Trust, Concerns Over Academic Integrity and Learning Loss, Limitations and Inaccuracies of Generative AI in Practical Applications, Ethical and Regulatory Concerns Surrounding AI Technology, Ineffectiveness of Generative AI in Teaching Programming, Risks of Inaccurate Code Formation and Reliability, AI as a Crutch Leading to Skill Degradation, Overreliance on Generative AI Leads to Lack of Critical Understanding, Accuracy and Quality Concerns around AI-Generated Code, Replacement of Traditional Learning and Development Practices, Concerns Over Dependence on AI Tools, Quality and Reliability of AI-Generated Code, Potential Erosion of Academic Integrity and Traditional Learning, Limitations in Understanding and Creativity, Ethical and Academic Integrity Issues

D. *Positive Themes from LLM-based Sentiment Analysis*

Enhanced Learning and Personal Development, Practical Applications and Automation, Bridging the Skills Gap, Empowerment through Learning and Skill Development, Increased Productivity and Efficiency, Creative Collaboration and Idea Generation, Enhancing Learning and Problem Solving, Increased Productivity and Efficiency in Coding, Encouragement of Creative Exploration, Learning Enhancement and Support for Beginners, Efficiency in Coding Tasks, Creative Collaboration in Game Development, Enhanced Learning and Understanding Through AI Assistance, Encouragement of Creative Problem-Solving, Enhanced Learning and Understanding, Increased Productivity and Efficiency, Practical Application in Real-world Projects, Enhancing Learning and Problem Solving, Facilitating Collaboration and Ideation, Enhancing Learning and Proficiency in Coding, Automating and Streamlining Development Processes, Adapting Curriculum and Teaching Approaches in Education, Enhancing Learning and Understanding, Accelerating Development and Project Completion, Problem Solving and Support in Coding, Increased Efficiency in Learning Process, Bridging the Generational Divide, Coding Assistance and Productivity, Boosting Productivity and Efficiency, Democratizing Access to Expertise and Resources.

E. Neutral Themes from LLM-based Sentiment Analysis

The Role of Human Expertise and Collaboration, Evolving Educational Needs and Career Skills, Educational Support and Enhancements, Limitations and Caution in Use, AI's Role in Learning Processes, Utilization of Generative AI as a Learning Tool, The Role of AI in Developing Proficiency Among Developers, The Limitations of Generative AI and the Importance of Critical Thinking, Adaptation to New Tools, Quality Control and Verification, Future of AI Integration in Education, Importance of Fundamentals in Programming, Generative AI as a Tool for Learning and Problem-Solving, Awareness of Limitations and Role of Human Oversight, Variability in Effectiveness of Generative AI Tools in Coding, The Need for Critical Engagement with AI Outputs, The Role of Generative AI in Educational Settings, Generative AI as a Tool for Assistance, Limitations and Challenges of Generative AI, The Evolving Role of Developers, Ethical Considerations and Integration in Education, Learning Enhancement and Resource Optimization, The Mixed Role of AI in Practical Development, Certification and Skill Validation in AI's Era, AI as a Tool for Enhancing Education and Productivity, Ethical Considerations and Future Implications of Generative AI in Education, Tools and Adaptation in Programming Workflows, Educational Enhancement, Resource Availability, Ethical and Critical Thinking Development

F. Negative Themes from LLM-based Sentiment Analysis

Lack of Understanding and Overreliance on Generative AI, Concerns About Quality and Reliability, Potential Negative Impacts on Learning and Critical Skills, Ineffective Learning and Comprehension, Quality and Reliability of AI-Assisted Outputs, Ethical Implications and Academic Integrity, Accuracy and Reliability Concerns, Impacts on Learning and Skill Development, Lack of Understanding and Dependency on AI Tools, Devaluation of Software Development Skills, Concerns About AI's Reliability and Long-Term Viability, Limitations of Generative AI in Technical Contexts, Devaluation of Skills and Creativity, Ethical and Copyright Concerns, Impeding Genuine Learning and Problem-Solving, The Perception of Generative AI as a Crutch, Threat to Individual Expression and Craftsmanship, Dependence on AI Diminishing Critical Thinking Skills, Quality Deterioration of Coding Standards, Encouragement of Plagiarism and Academic Dishonesty, Overreliance on Generative AI Leading to Skills Degradation, Lack of Understanding and Critical Skills Development, Concerns Over Quality and Accuracy in Outputs, Ethical and Professional Implications, Lack of Intellectual Value in Learning, Dependency on Errors and Limitations of AI, Concerns Over Learning Erosion, Ethical and Security Implications, Erosion of Fundamental Skills, Lack of Critical Thinking Development

References

[1] T. Adıgüzel, M. H. Kaya, and F. K. Cansu, “Revolutionizing education with AI: exploring the transformative potential of ChatGPT,” Contemporary Educational Technology, 2023.

[2] A. Haidar, “ChatGPT and generative AI in educational ecosystems: transforming student engagement and ensuring digital safety,” in Preparing Students for the Future Educational Paradigm, pp. 70-100, IGI Global, 2024.

[3] I. D. Raji, M. K. Scheuerman, and R. Amironesei, “You can't sit with us: exclusionary pedagogy in AI ethics education,” in Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, pp. 515-525, March 2021.

[4] M. M. Rahman and Y. Watanobe, “ChatGPT for education and research: opportunities, threats, and strategies,” Applied Sciences, vol. 13, no. 9, p. 5783, 2023.

[5] R. Yilmaz and F. G. K. Yilmaz, “The effect of generative artificial intelligence (AI)-based tool use on students' computational thinking skills, programming self-efficacy and motivation,” Computers and Education: Artificial Intelligence, vol. 4, p. 100147, 2023.

[6] D. Baidoo-Anu and L. O. Ansah, "Education in the era of generative artificial intelligence (AI): Understanding the potential benefits of ChatGPT in promoting teaching and learning," Journal of AI, vol. 7, no. 1, pp. 52-62, 2023.

[7] Z. Bahroun, C. Anane, V. Ahmed, and A. Zacca, "Transforming education: A comprehensive review of generative artificial intelligence in educational settings through bibliometric and content analysis," Sustainability, vol. 15, no. 17, p. 12983, 2023.

[8] O. Hazzan and Y. Erez, "Generative AI in Computer Science Education," in Proc. 55th ACM Technical Symposium on Computer Science Education, vol. 2, pp. 1899-1899, March 2024.

[9] A. Ghimire, "Generative AI in Education From the Perspective of Students, Educators, and Administrators," 2024.

[10] Y. Dai, A. Liu, and C. P. Lim, "Reconceptualizing ChatGPT and generative AI as a student-driven innovation in higher education," Procedia CIRP, vol. 119, pp. 84-90, 2023.

[11] M. U. Hadi, R. Qureshi, A. Shah, M. Irfan, A. Zafar, M. B. Shaikh, ... S. Mirjalili, "A survey on large language models: Applications, challenges, limitations, and practical usage," Authorea Preprints, 2023.

[12] A. Al Yakin, M. Muthmainnah, E. Apriani, A. J. Obaid, and A. A. Elngar, "Transforming Online Learning Management: Generative Models on ChatGPT to Enhance Online Student Engagement Scale (OLE)," Idarah (Jurnal Pendidikan dan Kependidikan), vol. 7, no. 2, pp. 135-148, 2023.

[13] L. Jacques, "Teaching CS-101 at the Dawn of ChatGPT," ACM Inroads, vol. 14, no. 2, pp. 40-46, 2023.

[14] B. A. Becker, P. Denny, J. Finnie-Ansley, A. Luxton-Reilly, J. Prather, and E. A. Santos, "Programming is hard—or at least it used to be: Educational opportunities and challenges of AI code generation," in Proc. 54th ACM Technical Symposium on Computer Science Education, vol. 1, pp. 500-506, March 2023.

[15] C. Zastudil, M. Rogalska, C. Kapp, J. Vaughn, and S. MacNeil, "Generative AI in computing education: Perspectives of students and instructors," in 2023 IEEE Frontiers in Education Conference (FIE), pp. 1-9, October 2023, IEEE.

[16] J. Finnie-Ansley, P. Denny, B. A. Becker, A. Luxton-Reilly, and J. Prather, "The robots are coming: Exploring the implications of OpenAI Codex on introductory programming," in Proc. 24th Australasian Computing Education Conference, pp. 10-19, February 2022.

[17] P. Denny, V. Kumar, and N. Giacaman, "Conversing with copilot: Exploring prompt engineering for solving CS1 problems using natural language," in Proc. 54th ACM Technical Symposium on Computer Science Education, vol. 1, pp. 1136-1142, March 2023.

[18] K. Cortiñas-Lorenzo and G. Lacey, "Toward explainable affective computing: A review," IEEE Transactions on Neural Networks and Learning Systems, 2023.

[19] Y. Chen and P. Esmaeilzadeh, "Generative AI in medical practice: in-depth exploration of privacy and security challenges," Journal of Medical Internet Research, vol. 26, e53008, 2024.

[20] W. M. Lim, S. Kumar, S. Verma, and R. Chaturvedi, "Alexa, what do we know about conversational commerce? Insights from a systematic literature review," Psychology & Marketing, vol. 39, no. 6, pp. 1129-1155, 2022.

[21] N. Rane, "Role and challenges of ChatGPT and similar generative artificial intelligence in business management," Available at SSRN 4603227, 2023.

[22] H. Khosravi, S. B. Shum, G. Chen, C. Conati, Y. S. Tsai, J. Kay, ... D. Gašević, "Explainable artificial intelligence in education," Computers and Education: Artificial Intelligence, vol. 3, p. 100074, 2022.

[23] M. Castelli and L. Manzoni, "Generative models in artificial intelligence and their applications," Applied Sciences, vol. 12, no. 9, p. 4127, 2022.

[24] K. Saha and M. De Choudhury, "Assessing the mental health of college students by leveraging social media data," XRDS: Crossroads, The ACM Magazine for Students, vol. 28, no. 1, pp. 54–58, 2021. [Online]. Available: https://doi.org/10.1145/3481834

[25] J. M. Pérez, M. Rajngewerc, J. C. Giudici, D. A. Furman, F. Luque, L. A. Alemany, and M. V. Martínez, "pysentimiento: a python toolkit for opinion mining and social NLP tasks," arXiv preprint arXiv:2106.09462, 2021.