

Computer Science Student Attitudes Towards AI Ethics and Policy: A Preliminary Investigation

James Weichert
Department of Computer Science
Virginia Tech
Blacksburg, VA, USA
jamesweichert@vt.edu

Hoda Eldardiry
Department of Computer Science
Virginia Tech
Blacksburg, VA, USA
hdardiry@vt.edu

Abstract—The explosive growth of artificial intelligence (AI) technologies in everyday settings in recent years has highlighted the need to develop comprehensive policies to promote the ethical use of AI. As the next generation of AI developers and policymakers receive training on technical AI foundations, it is also important to examine how discussions around AI ethics and policy are (or are not) woven into existing computer science (CS) curricula. Thus, the perceptions of current college students studying AI are valuable in two ways: (1) to assess how AI ethics is currently being taught at the university level; and (2) to understand the attitudes of this new generation of AI thinkers towards AI, both in general and with respect to AI ethics.

This paper summarizes the results of a preliminary survey of undergraduate CS students ($n = 41$) enrolled in a machine learning course at a large public university in the United States. We use our survey instrument to assess student attitudes towards AI, AI ethics, and AI policy. We find that while CS students are generally very positive about the benefit of AI and use AI tools regularly, they are nevertheless worried about the ethical impact of current and future AI technologies. Moreover, although nearly half of the students we surveyed would be interested in AI policy as a potential career path, only a third of respondents believed that their university courses were adequately preparing them to engage in discussions around AI policy and regulation. In this paper, we further evaluate these survey results and discuss implications for AI education at large.

Keywords—AI ethics, AI policy, CS education

I. INTRODUCTION

The emergence in recent years of ‘consumer-ready’ artificial intelligence (AI) technologies like ChatGPT or the “full self-driving” feature in some Tesla cars has justifiably raised concern over whether public and private sector AI governance is matured enough to effectively mitigate unethical applications of these technologies. Only recently are governments starting to ‘catch up’ by implementing regulatory frameworks for AI. In the European Union (EU), the EU AI Act was finally ratified in March 2024, and in China the central government has imposed restrictions on the sharing of private data and on generative AI [1] [2]. Even in the US, which has lagged behind both the EU and China on AI regulation [3], policy is beginning to take shape; in October 2023, President Biden signed an executive order on the *Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence* [4] and in his

March 2024 State of the Union address called for a ban on “AI voice impersonation and more” [5].

It is clear that policy and regulation represent a growing area of focus for AI, both in terms of research and application [6]. So it is reasonable to conclude that the demand for AI professionals with both technical and policy or governance skill sets will increase along with the advancement of AI technology itself [7]. Indeed, President Biden’s October Executive Order [4] calls for a government-wide “AI talent surge,” including AI policy positions. Yet it remains unclear whether current computer science (CS) curricula are adequately preparing students to fill these roles in the future [7], and whether these students are interested in policy careers in the first place.

To this end, this paper presents the results of a pilot study investigating college CS student attitudes towards AI, with a particular focus on AI ethics and policy. We design and field a survey to an undergraduate machine learning (ML) course at a large public U.S. university, and use the responses ($n = 41$) to evaluate student perceptions of AI, AI ethics and AI policy. Concretely, we focus on two principal lines of inquiry:

- 1) What are CS students’ general attitudes towards emerging AI technologies?
- 2) To what extent are students interested in and prepared for a job involving AI ethics, regulation and policymaking?

II. RELATED WORK

Recent literature on AI ethics has highlighted the importance of developing policies and governance structures for AI that facilitate ethical use. Kim et al. [6] document a “policy turn” away ethical *principles* towards the development of governance frameworks and policies to enable ethical *practice*. Given the nascency of ‘AI policy’ as a field of study, research has also focused on evaluating the efficacy of emerging regulatory schemes. Zhu et al. [7] conduct interviews with fifteen AI professionals in policy-related roles, who highlight the global politics of AI (e.g. the climate impact of AI training) and long-term evaluation of AI systems as two ethical concerns not widely covered in existing AI ethics literature. The authors conclude that the social and political contexts raised by the interviewees should not be omitted from AI ethics education. “Traditionally decontextualized [ethics] problems,” they argue,

should be substituted for exposure to “authentic, socially integrated problems in science and engineering programs.”

To this end, other studies have described the results of surveys of college student attitudes towards AI and AI ethics. Jang et al. [8] develop a survey instrument to assess respondent attitudes towards AI across five common ethical principles: fairness, transparency, non-maleficence, privacy, and responsibility. Their research revealed differences in attitudes between genders and levels of AI education among respondents enrolled in an online course at a Korean university. Ghotbi et al. [9] surveyed college students at a university in Japan regarding which of nine provided ethical issues associated with the future of AI (e.g. unemployment, increasing inequality, etc.) they deemed most significant. Hooper and Fletcher [10] develop but do not field a survey to measure the change in student attitudes towards AI ethics before and after taking an AI ethics course. Most relevant to this study, however, is the survey conducted by Kim et al. [11] of nearly 1000 students at a large public university in the U.S. Their study is useful to us as a baseline for comparing general student attitudes towards AI between CS and non-CS undergraduate students. The key limitation of the study is that it focuses on attitudes towards AI use in higher education, not AI ethics more broadly or AI regulation and policy. To our knowledge, there has been no study of college-level CS students focused on AI ethics and policymaking *competencies* and *attitudes* towards AI regulation, only studies focused on measuring general ethical values.

Thus, we set our survey apart from prior work in two key ways. First, we expand the focus of the survey beyond ethics to include both general attitudes towards AI *and* AI policy attitudes and competencies. In this way, we can compare the perceptions of students towards AI from a decontextualized perspective (*general* attitudes towards AI) with student attitudes within the context of government AI regulation. Furthermore, the inclusion of survey questions focused on AI policy attitudes and competencies allows us to begin to assess the extent to which a more explicit incorporation of AI *policy* education is needed in CS curricula, and whether current AI students have the desire and skills necessary to become AI policy professionals. Second, we use as a starting point the Kim et al. [11] survey of students from across all disciplines, allowing us to compare responses to similar questions between CS and non-CS students.

III. METHODOLOGY

A. Study Population

The study population for this research consists of undergraduate and graduate students at a large public university in the United States enrolled in advanced-level computer science courses focusing on AI or ML. We refer to students in this population as “computer science students,” without regard for their specific major or research focus.

For this preliminary phase of the study, the survey was distributed to students enrolled in an introductory undergraduate machine learning course taught in Spring 2024. This course

is taught by the computer science department and consists primarily of third- and fourth-year undergraduate students, most of whom are computer science majors. The Spring 2024 offering of the course had 50 students enrolled.

B. Survey Instrument

The survey instrument designed for this study builds on the work of Kim et al. [11], who surveyed student and faculty attitudes towards generative AI technology across STEM and non-STEM disciplines. A key limitation of their survey is the inability to identify key differences in attitudes among CS students—who presumably learn about and use AI tools most often—and non-CS students. We view this distinction as more decisive in the context of AI use and attitudes than the STEM vs. non-STEM divide. To extend on their work, our survey includes a section on *General Attitudes Towards AI* containing two Likert scale statements taken from the Kim et al. survey (with slight modifications): (1) “I believe that generative AI tools are helpful in my university studies.” and (2) “Using generative AI tools increases my productivity.” In both cases, the word “generative” was omitted in our survey, which asks instead about AI tools in general. The inclusion of these questions will enable a comparison of attitudes between CS students and undergraduate students at large.

Our survey also includes two additional sections on *AI Ethics* and *AI Regulation and Policy*, both with Likert scale questions assessing students’ concerns over the ethical development of AI technology, AI ethics education, and AI governance. The latter is the key novel contribution of this survey, which seeks to identify student attitudes towards AI regulation and to assess the extent to which their CS education has prepared them for a possible career in AI policy.

Table I contains all survey questions. All Likert scale questions were preceded with the instruction “Please rate your agreement with the following statements on a scale of 1 (strongly disagree) to 5 (strongly agree).”

C. Data Collection

Data were collected using an anonymous online survey accessible to students using their university email. After obtaining approval from our university’s Institutional Review Board (IRB), the survey was advertised to students enrolled in the machine learning course with the permission of the instructor. An announcement about the survey was made in-person during class time and through the course’s learning management system (LMS). As a participation incentive, two survey respondents were randomly selected to each receive a \$10 gift card. Additionally, each student who completed the survey was awarded extra credit in the course amounting to 0.5% of their overall grade. The survey remained open for responses for a three-week period in April and May 2024.

41 students (82% of the class) responded to at least one of the questions (each of which was optional), with the number of responses for all questions ranging between 31 and 41. The survey responses exhibited some attrition, with the average number of respondents per question dropping from 40.8 in

TABLE I
SURVEY QUESTIONS

Question	Response Type
General Attitudes Towards AI	
I am familiar with modern artificial intelligence (AI) technologies like neural networks, large language models (LLMs), and autonomous vehicles (“self-driving cars”). I can explain to someone who is not familiar with AI, in general terms, how most AI technologies work. I believe that current AI tools provide more benefits than drawbacks. I believe that future AI tools will provide more benefits than drawbacks. I believe AI tools are helpful to me in my day-to-day life. I believe AI tools are helpful in my university studies. Using AI tools increases my productivity.	Likert ¹
How often do you use AI tools (e.g. ChatGPT, Google Gemini, Tesla Autopilot, etc.) in your day-to-day life? How often do you use AI tools in your university studies? What, if any, AI technologies do you use?	Multiple Choice
AI Ethics	
In general, I think existing AI tools are ethical. I believe that most developers of AI tools design their AI systems with ethics in mind. I worry about the ethical impact of current AI technology. I worry about the ethical impact of future AI technology. I can explain to someone not familiar with AI how AI can make biased or harmful decisions. My artificial intelligence and/or machine learning courses at university integrate discussion about the ethics of AI into the curriculum. My computer science education focuses sufficient attention on AI ethics. My courses at university have prepared me to discuss and mitigate potential harms that AI can cause.	Likert
AI Regulation and Policy	
I believe AI technologies are currently adequately regulated by the government. The US government is doing a good job of balancing promoting innovation and protecting users with respect to AI technology. The US government should do more to protect users from the potential harms of AI technology. We talk about the regulation of AI by the government in my artificial intelligence and/or machine learning courses at university. My courses at university are adequately preparing me to engage in conversations about AI policy and regulation. I am interested in AI policy and regulation as a potential career path.	Likert

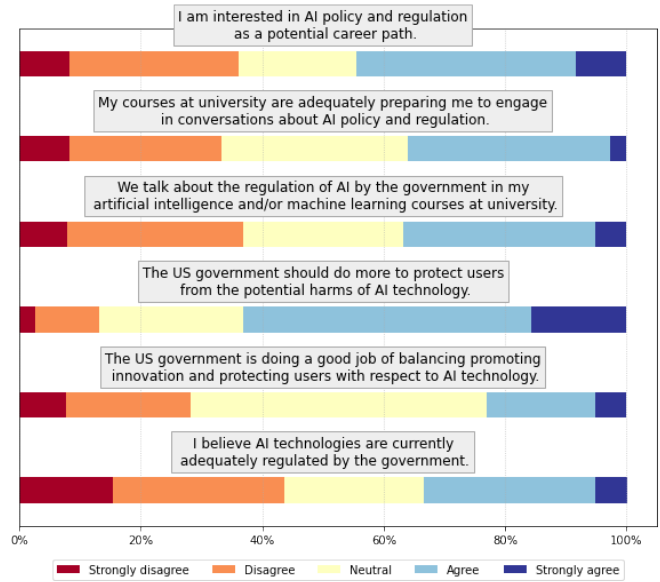


Fig. 1. Undergraduate CS student attitudes towards AI policy and regulation.

the first section of the survey to 36.3 and 37.7 in the second and third sections, respectively.

IV. FINDINGS

A. Attitudes Towards AI, Ethics and Policy

Overall, the responses to the Likert scale questions exhibit a high degree of uniformity, with ‘Agree’ (4 on a 5-point scale) constituting the majority or plurality response across all 15 questions related to general attitudes and attitudes towards AI ethics (sections 1 and 2 of the survey). For example, nearly 93% of respondents agreed or strongly agreed with the statement “I can explain to someone who is not familiar with AI, in general terms, how most AI technologies work,” while 78% of respondents agreed or strongly agreed that “future AI tools will provide more benefits than drawbacks.” There was less agreement for ethics-related questions, such as “I believe that most developers of AI tools design their AI systems with ethics in mind,” (42% agree or strongly agree) and “courses at university have prepared me to discuss and mitigate potential harms that AI can cause,” (52% agree or strongly agree).

The responses in the *AI Regulation and Policy* section of the survey were altogether less homogeneous than in the previous sections. Likewise, agreement with the statements was noticeably lower. For questions in this section, the average Likert rating was 3.06 (close to 3 - ‘Neutral’) compared to averages closer to 4 (‘Agree’) in sections 1 ($\mu = 3.99$) and 2 ($\mu = 3.58$). The average standard deviation for responses in this section was also 30% higher than section 1 and 14% higher than section 2, hinting at a divergence in attitudes for policy-related questions, including U.S. AI regulation and discussion of AI policy-related topics in CS courses. Figure 1 shows the full Likert scale distribution for questions from this section.

To substantiate the hypothesis that attitudes across sections differ significantly, we conducted a one-way ANOVA test on




Fig. 2. Frequency of AI tool use among undergraduate CS students.

the means of each Likert scale question, grouped by survey section. Based on the high F statistic (20.13) returned by the test, we ran a Tukey’s Honestly Significant Difference (HSD) test [12]. The HSD p-values not only confirmed that the lower mean for the *AI Regulation and Policy* section is statistically significant (at $p < 0.01$), but so too is the difference in average agreement between sections 1 and 2 (at $p < 0.05$).

B. Frequency of AI Use

The *General Attitudes Towards AI* survey section also included two questions asking about the frequency of use of “AI tools”. In “day-to-day life”, 78% of respondents use AI technology at least a few times per week. For “university studies”, the frequency is lower but still high, with 66% of participants reporting using AI tools at least a few times per week. Figure 2 shows the full distributions for day-to-day and university-related AI use. A similar question about generative AI use for university studies was posed in the Kim et al. [11] study. While their study focuses on the more narrow term “generative AI”, we believe that the questions are similar enough to allow comparison. In contrast to the CS students in our study, the general undergraduate student population in the Kim et al. survey was found to use AI tools infrequently. Less than 30% of student respondents reported using AI at least once per week.

V. DISCUSSION

A. General Attitudes Towards AI

The outlook of the CS students we surveyed is, in general, very positive. Nearly 80% of respondents agreed with the statement that “current AI tools provide more benefits than drawbacks.” The percentage was the same when asking about future AI tools. Furthermore, 78% of surveyed students find AI tools useful, both in everyday life and at university.

Our survey also provides an indicator as to which AI technologies are used by college-aged individuals. Nearly all respondents reported using Large Language Model (LLM) chatbots (e.g. ChatGPT), while 78% of students use AI in the form of social media recommender systems (e.g. TikTok’s “For Your Page”). Approximately 40% of students use AI writing or coding assistants (e.g. Grammarly AI or GitHub Copilot). Perhaps unsurprisingly, only 2 of 41 survey respondents listed autonomous driving as an AI technology they use. Survey data from the wider undergraduate population will

likely show lower AI tool use across all categories, but we posit that the relative popularity of both LLMs and social media recommendation algorithms will persist. For future surveys, it would be beneficial to disentangle the use of LLMs for *question answering* tasks (e.g. prompting “What are the main literary themes in ‘The Great Gatsby’?”) from *composition* tasks (e.g. prompting “Write a college English essay about symbolism in ‘The Great Gatsby’”). Doing so would elucidate whether the AI tasks performed by CS students differ substantially from those performed by non-CS students. In other words, do CS students tend to use LLMs less for *composition* because their courses tend to rely less on essays?

B. AI Ethics

Although average agreement with the Likert scale statements remained higher for *AI Ethics* questions compared to *AI Regulation and Policy* questions (averages of 3.58 vs. 3.06 on a 5-point scale), the Likert responses in section 2 indicate a level of concern among CS students for the ethical impacts of AI. For example, nearly a third of respondents disagreed with the statement that “most developers of AI tools design their AI systems with ethics in mind,” (42% agreed or strongly agreed). Likewise, 58% of respondents indicated that they “worry about the ethical impact of current AI technology,” and 76% indicated that they “worry about the ethical impact of future AI technology.”

For AI ethics education, 80% of students expressed agreement that “discussion about the ethics of AI” is integrated into their AI and ML courses. The extent and efficacy of this integration, however, is less clear. While a majority of respondents indicated that there is “sufficient attention” paid toward AI ethics in their CS education, 32% disagreed or strongly disagreed. Nevertheless, only 52% of students thought that their courses had prepared them to “discuss and mitigate potential harms that AI can cause.”. These results suggest that AI ethics topics are present in the curriculum, but (1) these topics may not be taught in a manner that is engaging for all students; and (2) the existing ethics curricula may not instill the competencies needed for students to act as responsible and ethical AI practitioners after graduating.

C. AI Regulation and Policy

The survey results support our view that *AI policy* education merits more attention in CS curricula and CS education in general. Although 44% of respondents expressed interest in AI policy as a potential career path, only 36% of students felt that their courses are “adequately preparing me to engage in conversations about AI policy and regulation.” We expect that AI developers will in the future be called upon to operationalize complex and perhaps contradictory policy requirements. Thus, even if a majority of students are not interested in AI policy as a standalone career, their jobs in the AI industry may nevertheless require a more nuanced understanding of AI policy than their CS education currently provides.

VI. CONCLUSION

In this paper, we present the findings of a preliminary, small-scale ($n = 41$) survey of undergraduate computer science students enrolled in a machine learning course. The survey reveals that while these students have a positive outlook towards AI in general and use AI tools very often compared to their non-CS peers, they are nevertheless concerned about the ethical impacts of both current and future AI technologies.

A. Implications for CS Education

Student responses regarding the incorporation of AI ethics and policy discussions in existing CS curricula were mixed. While a majority of respondents agreed that “sufficient attention” is paid to AI ethics in their CS education (53% agree or strongly agree vs. 32% disagree or strongly disagree), students split on similar questions related to classroom discussion of AI regulation (37% vs. 37%) and preparing students to engage in policy discussions (36% vs. 33%). Most significantly, our survey results identify AI policy as a career path that a significant portion of CS students *are* interested in exploring. These data provide educators and education researchers with evidence that curricular development in the area of AI policy is worth exploring. Furthermore, this survey, which took students on average only 2 minutes to complete, could be fielded as-is at the beginning and end of an AI ethics course to explore changes in student perceptions towards AI ethics and policy.

B. Future Work

We intend to substantiate these conclusions by expanding the study population to all undergraduate and graduate CS students enrolled in AI, ML and CS ethics courses at our university, fielding a larger survey in Fall 2024. These new results will help to validate the findings in this paper, and to compare differences in attitudes between undergraduate and graduates. This study could also be extended to multiple universities, both in the U.S. and abroad. Coupling the survey responses with an analysis of each university’s AI ethics and policy curricula could aid in the identification of relationships between teaching and student attitudes towards AI ethics. Likewise, survey responses could inform a review of existing curricula to improve student ethics and policy competencies.

REFERENCES

- [1] H. Roberts, J. Cows, E. Hine, J. Morley, V. Wang, M. Taddeo, and L. Floridi, “Governing artificial intelligence in China and the European Union: Comparing aims and promoting ethical outcomes,” *The Information Society*, vol. 39, no. 2, pp. 79–97, Sep. 2022.
- [2] E. Hine and L. Floridi, “New deepfake regulations in China are a tool for social stability, but at what cost?” *Nature Machine Intelligence*, vol. 4, pp. 608–610, 2022.
- [3] S. Parinandi, J. Crosson, K. Peterson, and S. Nadarevic, “Investigating the politics and content of US State artificial intelligence legislation,” *Business and Politics*, pp. 1–23, Mar. 2024.
- [4] Exec. Order No. 14110, 88 C.F.R. 75191 (2023).
- [5] J. Biden, “Remarks of President Joe Biden — State of the Union Address As Prepared for Delivery,” U.S. Capitol, Mar. 2024.
- [6] D. Kim, Q. Zhu, and H. Eldardiry, “Exploring Approaches to Artificial Intelligence Governance: From Ethics to Policy,” in *Proceedings of the 2023 IEEE International Symposium on Ethics in Engineering, Science, and Technology (ETHICS)*, West Lafayette, IN, USA, 2023, pp. 1–5.
- [7] Q. Zhu, D. Kim, H. Eldardiry, and M. C. Ausman, “A preliminary investigation of the ethics policy concerns of artificial intelligence: insights from AI professionals working in policy-related roles,” in *Proceedings of the 2023 IEEE Frontiers in Education Conference (FIE)*, College Station, TX, USA, 2023, pp. 1–4.
- [8] Y. Jang, S. Choi, and H. Kim, “Development and validation of an instrument to measure undergraduate students’ attitudes toward the ethics of artificial intelligence (AT-EAI) and analysis of its difference by gender and experience of AI education,” *Education and Information Technologies*, vol. 27, pp. 11 635–11 667, May 2022.
- [9] N. Ghotbi and M. T. Ho, “Moral awareness of college students regarding artificial intelligence,” *Asian Bioethics Review*, vol. 13, no. 4, pp. 421–433, Sep. 2021.
- [10] K. Hooper and T. Fletcher, “An analysis of engineering and computing students’ attitudes to AI and ethics,” in *Proceedings of the 2022 ASEE Annual Conference & Exposition*, Minneapolis, MN, Aug. 2022.
- [11] J. Kim, M. Klopfer, J. R. Grohs, H. Eldardiry, J. Weichert, L. A. Cox II, and D. Pike, “An examination of faculty and students’ perceptions of generative AI use for university classes,” unpublished, Apr. 2024.
- [12] H. Abdi and L. J. Williams, “Tukey’s Honestly Significant Difference (HSD) Test,” in *Encyclopedia of Research Design*, N. Salkind, Ed. Thousand Oaks, CA: Sage, 2010.