



Computing Specializations: Perceptions of AI and Cybersecurity Among CS Students

Vidushi Ojha

Computer Science

University of Illinois at Urbana-Champaign
Urbana-Champaign, IL, USA
vojha3@illinois.edu

Brent Lagesse

Computing & Software Systems
University of Washington Bothell
Bothell, WA, USA
lagesse@uw.edu

Christopher Perdriau

Computer Science

University of Illinois at Urbana-Champaign
Urbana-Champaign, IL, USA
chp5@illinois.edu

Colleen M. Lewis

Computer Science

University of Illinois at Urbana-Champaign
Urbana-Champaign, IL, USA
colleenl@illinois.edu

ABSTRACT

Artificial intelligence (AI) and cybersecurity are in-demand skills, but little is known about what factors influence computer science (CS) undergraduate students' decisions on whether to specialize in AI or cybersecurity and how these factors may differ between populations. In this study, we interviewed undergraduate CS majors about their perceptions of AI and cybersecurity. Qualitative analyses of these interviews show that students have narrow beliefs about what kind of work AI and cybersecurity entail, the kinds of people who work in these fields, and the potential societal impact AI and cybersecurity may have. Specifically, students tended to believe that all work in AI requires math and training models, while cybersecurity consists of low-level programming; that innately smart people work in both fields; that working in AI comes with ethical concerns; and that cybersecurity skills are important in contemporary society. Some of these perceptions reinforce existing stereotypes about computing and may disproportionately affect the participation of students from groups historically underrepresented in computing. Our key contribution is identifying beliefs that students expressed about AI and cybersecurity that may affect their interest in pursuing the two fields and may, therefore, inform efforts to expand students' views of AI and cybersecurity. Expanding student perceptions of AI and cybersecurity may help correct misconceptions and challenge narrow definitions, which in turn can encourage participation in these fields from all students.

CCS CONCEPTS

- Social and professional topics → Computing education; Computer science education.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGCSE '23, March 15–18, 2023, Toronto, ON, Canada.

© 2023 Association for Computing Machinery.

ACM ISBN 978-1-4503-9431-4/23/03...\$15.00

<https://doi.org/10.1145/3545945.3569782>

KEYWORDS

Artificial intelligence; cybersecurity; computer science education; broadening participation in computing

ACM Reference Format:

Vidushi Ojha, Christopher Perdriau, Brent Lagesse, and Colleen M. Lewis. 2023. Computing Specializations: Perceptions of AI and Cybersecurity Among CS Students. In *Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 1 (SIGCSE 2023), March 15–18, 2023, Toronto, ON, Canada*. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3545945.3569782>

1 INTRODUCTION

Artificial intelligence (AI) and cybersecurity are growing fields: individuals with expertise in AI, cybersecurity, or both are increasingly in-demand for roles in industry and government [9, 24, 29]. Current research suggests that there are insufficient numbers of individuals with appropriate expertise to fulfill this increasing demand, particularly in cybersecurity [9] and in applying AI to cybersecurity [19]. A significant number of occupations in the technology industry require AI or cybersecurity skills, a number that is likely to increase in the near future [24]. Jobs in cybersecurity already exceed the number of qualified professionals to fill them: according to Cyberseek, a website providing data on the cybersecurity job market, “[o]n average, cybersecurity roles take 21% longer to fill than other IT [information technology] jobs” because employers cannot find workers with appropriate qualifications [9].

A factor further exacerbating the gap between qualified professionals and available positions is that the workforce in computing in general, and cybersecurity in particular, lacks proportional representation of people from different races, ethnicities, and gender identities [16, 17]. Specifically, in the U.S., people who identify as women, Black/African American, Hispanic, Latina/o/x, Native American, Native Alaskan, Native Hawaiian, Pacific Islanders, and/or disabled belong to historically underrepresented groups (HUGs) in computing [20]. These patterns of underrepresentation are also apparent in the educational system [7, 8, 25, 27], suggesting that in order to broaden participation, it is important to understand how students decide whether to pursue AI and cybersecurity, as these students go on to become qualified professionals in the workforce.

Previous work suggests that, in general, computer science (CS) students make decisions regarding specializations based on limited information [14], but little is known about specific factors that may influence their interest in AI or cybersecurity or discourage them from pursuing these fields entirely. A greater understanding of how students perceive these two fields may lead to interventions that correct misconceptions and highlight factors likely to draw students' interest, especially students from HUGs. To this end, we address the following research question:

RQ: How do undergraduate computing students perceive AI and cybersecurity?

We interviewed 17 students majoring or minoring in CS at a large, public university about their experiences and perceptions of AI¹ and cybersecurity. Using qualitative analysis methods, we identified patterns in students' beliefs regarding the nature of the work in AI and cybersecurity, the kind of people who work in the fields, and the societal impact such work may have. Our results suggest that participants have narrow perceptions of AI and cybersecurity. In particular, students tended to believe that all work in AI requires math and training models, while cybersecurity consists of low-level programming; that innately smart people work in both fields; that working in AI comes with ethical concerns; and that cybersecurity skills are important in contemporary society.

Our key contribution is identifying beliefs that students expressed about AI and cybersecurity that may affect their interest in pursuing the two fields and may, therefore, inform efforts to expand students' views of AI and cybersecurity. Many of our participants' beliefs mirror existing stereotypes about the work and the people in computing at large. Some of these stereotypes have been shown to impact participation of people from HUGs [10, 21], suggesting that they may play a similar role in AI and cybersecurity. In addition, other beliefs expressed by students are not supported by experts in the field. For example, participants suggested that all occupations in cybersecurity require low-level programming, but there exist a wide variety of cybersecurity occupations in management, analysis, policy, ethics, and software engineering [1]. These findings may provide partial explanations for why students do not choose to pursue AI or cybersecurity and, as such, can be used to create interventions with the goal of increasing student interest in them, as well as expanding student perceptions of what these fields are and who works in them.

2 PREVIOUS RESEARCH

Previous research has investigated how stereotypes about computing might affect student decisions to pursue it, how students specialize within computing, and efforts to recruit students into AI and cybersecurity courses.

2.1 Computing stereotypes

Stereotypes about AI and cybersecurity may affect student decisions to pursue the fields, especially since individuals working in these fields are often represented in popular media [13]. Although we are not aware of work on stereotypes specific to AI or cybersecurity, previous research has documented stereotypes that students tend

to have about the people and the work in science, technology, and mathematics (STEM) fields generally [5, 10, 15, 21] and in computing specifically [4, 11, 22, 23].

Like in other STEM fields, people in computing are stereotyped as requiring innate brilliance. For instance, in a study by Lewis et al., undergraduate computing students seemed to believe that computing was an ability you were "just born with" [23]. In a study by Leslie et al., respondents from STEM fields, including those in CS, were more likely to agree with statements such as "Being a top scholar of [discipline] requires a special aptitude that just can't be taught" [21]. Endorsement of such beliefs was inversely correlated with the percentage of PhDs in the field awarded to women [21], suggesting that the prevalence of this stereotype may be detrimental to broadening the participation of women in computing.

Some students also stereotype people in computing as not working with or for the benefit of others. In Diekman et al.'s survey of undergraduate students, respondents did not believe that STEM fields, including computing, would allow them to positively impact society [10]. The perceived asocial nature of CS was also a theme identified by a Lewis et al. study, where they found that many students viewed having to work alone as a requirement of the field [22]. Beliefs about the ability of computing to benefit society have been shown to affect the likelihood of an individual pursuing CS [4, 10].

Another stereotype that many students have about people in CS is that they are typically men. Lewis et al.'s list of factors that students used to assess their fit in computing included the perceived masculinity of the field [22]. This perception likely contributes to the relative lack of women pursuing computing. Indeed, Cheryan et al. suggest that the "masculine culture" of some STEM fields, such as CS, contributes to the gender gap in these fields [5].

Our focus on AI and cybersecurity allows us to investigate the open question of whether students have stereotypes particular to these fields and whether such beliefs impact whether they specialize in AI or cybersecurity.

2.2 How students pick CS specializations

Prior work has shown that students typically do not have clear goals when choosing a specialization within computing. To understand how CS majors make these decisions, Hewner interviewed computing students and advisors about how students choose courses [14]. Hewner's resulting model suggests that computing majors did not begin their academic careers with concrete goals and simply followed the curriculum until they encountered a course experience that was markedly more or less enjoyable than other courses [14]. This experience then helped the students create more specific goals and pursue a newly-defined interest [14]. This model suggests that students do not necessarily have clear preferences or goals of specialization within the major at the outset.

However, many students have likely encountered, for example, "hackers" or AI robots in popular culture, which prior research has shown can lead to preconceived notions and stereotypes about the kinds of people who work in cybersecurity or AI [13, 31, 32]. An open question for our work to address is whether Hewner's model is reflected in students' experiences with AI or cybersecurity, where they may have already decided their level of interest in the field based on prior cultural exposure.

¹As most of our participants equated AI and machine learning (ML), we combine the two and refer only to AI for the remainder of this paper, excepting participant quotes.

2.3 Recruitment into AI and cybersecurity

Prior research on recruitment into AI has focused on the need to broaden the participation of people from HUGs in computing. A survey of undergraduates at by Barretto et al. showed that students from HUGs were less likely to take an AI course, and the explanation proffered by many participants was that they lacked an interest in the field, especially in the technical content [3]. Another barrier to student recruitment into AI is that students do not always have a clear understanding of what the content of the field really is. For example, Ottenbreit-Leftwich et al. interviewed 4th grade students about what they thought AI was, and common answers referred to “robotic vacuums” and “search engines” [30]. Kreinsen and Schulz found similar themes in their interviews with students from grades 7 through 10, where many interviewees centered their view of AI as “the brain of the robots” [18]. These studies suggest that, at least prior to any exposure at the college level, students are unlikely to have accurate or specific ideas of what AI work looks like.

Prior work has shown that a lack of sufficient cybersecurity course offerings may be contributing to the challenge of recruiting students to the field [6]. A study found that none of the 10 highest-ranked computer science undergraduate programs require cybersecurity and that most CS programs ranked in the top 50 offered fewer than five cybersecurity electives [6]. This dearth of available courses makes it harder for students to develop an interest in the field or learn what kind of work cybersecurity entails. Increasing knowledge of cybersecurity may be a particularly important factor in attracting students to the field, as suggested by the National Initiative for Cybersecurity Education (NICE) Workforce Framework for Cybersecurity [28]. This framework was created to more precisely describe cybersecurity work in support of the education of students and the recruitment of employees. For example, the 52 “Work Roles” listed by the NICE Framework describe areas of cybersecurity that an individual may be responsible for, including occupations such as “Cyber Legal Advisor” and “Technical Support Specialist” [1]. Students may not be aware that such roles are examples of cybersecurity work, making it harder to attract those with cybersecurity skills to these positions.

3 METHODS

3.1 Participant Recruitment

We conducted a qualitative study using semi-structured interviews with undergraduate students majoring or minoring in CS. We aimed to recruit students further in the degree program because these students were more likely to be taking advanced electives and to have thought about which specialization they want to pursue. Following IRB approval, participants were recruited by asking leaders of CS-specific student clubs to share our recruitment materials with their constituents. In particular, we recruited from the institution’s chapter of the Association for Computing Machinery and a club for women in CS. Students were incentivized to participate with a \$20 Amazon gift card.

In total, we conducted 17 interviews; these were conducted over Zoom, lasted an hour, and were recorded. Of our participants, 16 were majoring and one was minoring in CS; 11 identified as women, five as men, and one did not share their gender identity; 12 identified as Asian or Asian American, two as Caucasian or White,

and three did not share their racial/ethnic identity. We deliberately oversampled women due to our interest in broadening the participation of women in AI and cybersecurity. Each participant was assigned a pseudonym by the research team. Individual participant identities are not shared due to confidentiality concerns as well as because identifying differences between groups was not the focus of our work. Instead, we hope to have captured student perceptions that can later be used to identify such differences with the goal of broadening participation in AI and cybersecurity.

3.2 Interview Protocol

Consistent with recommended practice [26], interview questions were updated throughout data collection. Our initial protocol asked students about their CS background, their experiences with internships, and their perceptions about AI and cybersecurity, such as what they believe day-to-day work in these fields looks like. Based upon ongoing analysis, the interview protocol was refined three times to add questions where further detail was needed. We added questions about stereotypes students may have about AI and cybersecurity; whether they felt AI and cybersecurity should be required for all CS students; and beliefs about whether certain computing specializations were more difficult, rigorous, or prestigious than others, and their impact on society.

3.3 Data Analysis

Interview audio recordings were transcribed by a third-party service, Rev.com, and were anonymized and corrected by the authors. Data analysis was conducted in Saturate, an application for text-based qualitative analysis [35], using inductive approaches with no codes determined a priori. After six interviews had been independently read by two team members, we created an initial list of codes (words or phrases capturing some aspect of the text) based upon patterns apparent in these interviews. Our codes used the principles of concept and descriptive coding, i.e., were based on high-level ideas and descriptions in the text [33]. The two team members used this initial list to code three of the six interviews each and then met to refine the application of the codes. This formed the basis of the codebook used to code all subsequent interviews. To increase the confirmability of the findings [34], the team met weekly while coding the rest of the interviews. During these meetings, we reviewed codes by category in order to ensure consistency in how codes were applied. We added and removed codes as necessary, as well as resolved any conflicts or confusion regarding the meaning of the codes. Finally, upon the completion of coding, the data for related codes were reviewed to identify themes and variations in participants’ perceptions of the fields of AI and cybersecurity.

4 RESULTS

Our findings are organized into four categories relating to student perceptions of AI and cybersecurity that emerged from our analysis: their beliefs about (1) what skills are required in the field, (2) what the day-to-day work looks like, (3) what kind of people are in the field, and (4) what impact the field has on society. In all participant quotes, an ellipsis (“...”) indicates omitted words, a dash (“—”) pauses, square brackets added words for clarity, and *sic* an error in the original quote.

4.1 Skills required

4.1.1 *AI*. Many participants expressed the belief that AI requires mathematical skill and, perhaps for this reason, viewed the field as challenging and difficult. For example, Quincy believed that AI is heavily mathematical and stated, when discussing the importance of math in computing, that *“I mean, I think all of AI/ML is essentially just math.”* Other participants shared similar views, such as Gail’s perception that *“people who are good at statistics and math and, you know, will be encouraged to go into the field [of AI].”* Among those who endorsed this view, participants also seemed to agree that AI is difficult. They described AI as *“intimidating”* (Diane) and *“rigorous”* (Parth) and highlighted the difficulty of the AI course they had taken, which Mahir described as being *“considered one of the most challenging [courses], probably.”* Mahir further noted that it *“involved a lot of math and I think that’s what contributed to its perception as being difficult.”* It seems likely, therefore, that the perception of AI as being difficult is tied to its perception as mathematical, and that some students are choosing not to take AI courses at all because they think it would be too hard. For instance, Kasey expressed a reluctance to take an AI course because of its perceived difficulty: *“I have [considered it], I’m just afraid that it will be too difficult.”*

4.1.2 *Cybersecurity*. Much like AI, participants seemed to view cybersecurity as a difficult field, in this case due to the belief that it requires systems programming. When asked about what parts of computing were rigorous, Gail said, *“Just kind of with all the low-level, like assembly stuff or like security, right now I’m taking security, I feel like it’s really rigorous to be honest.”* Other participants shared the perception that cybersecurity incorporates low-level programming and is difficult. When Nora was asked what stereotypes she had heard regarding work in cybersecurity, she expressed that *“system[s] programming are (sic) always really time consuming and tedious.”* The reputation of cybersecurity as low-level also appeared to deter Idris from exploring it further, as they note that this has made them less interested in it: *“I guess like, a lot of security is lower level than I like to work.”*

Beyond its use of systems programming, participants seemed to suggest cybersecurity has a reputation as unfriendly to newcomers. For instance, when asked about what kind of people do cybersecurity, Gail recounted an experience attending a cybersecurity student club meeting and feeling *“out of place”* because *“everyone else kind of know (sic) what they’re doing,”* while she felt that *“I don’t even understand what’s happening.”* This experience left Gail feeling *“not good enough,”* highlighting how discouraging the experience was for her. Similarly, Diane, sharing her overall perspective of the field, explained that the cybersecurity course at the institution is *“notorious,”* and that she has *“felt that sense of intimidation and just never really wanted to explore it because it sounded like a heavy and difficult course.”* This sense of difficulty and intimidation may discourage students from exploring cybersecurity.

4.2 Day-to-day work

4.2.1 *AI*. When asked about what the day-to-day work of AI looks like, many participants described their understanding of AI work

as training models. Parth described his experience during an internship, talking to coworkers who were engaged in AI work, as follows:

“They’re developing some model or something or the other with machine learning. So they take a bunch of derivatives and then they use that to figure out, how should the model look like? And then they write that and then they train it or—yeah, they train it and then they try to vary the parameters of that model to best fit it. And then they call it a day.”

This general view, that AI work involves training models with data, was repeated by other participants, but very few of them had concrete explanations of what this process looked like. For instance, when asked what she thought the day-to-day work of AI looks like, Lillian suggested *“I guess it’d be like reading research papers or like just training models.”*

4.2.2 *Cybersecurity*. Participants’ conception of cybersecurity was that the work often entails *“hacking”* into systems or defending against attacks. For example, when Henry was asked what he thinks the daily work of cybersecurity looks like, he explained that he did not have first-hand experience, but added his perception that *“I guess there’s like white-hat hacking, that’s a thing, right. Where hackers like try to find vulnerabilities for companies, not to actually hack into them, but so that they can fix them.”* On the defensive side, Gail recounted having heard her cybersecurity professor talk about *“the adversarial mindset”* and *“think[ing] from the perspective of [an] attacker.”* She went on to say that the specific skills depend on what kind of cybersecurity work one wants to do, but that *“in general, it’s the ability to be able to think from the perspective of my attacker... com[ing] up with a way to defend against the attacker for [a] system.”* These perspectives suggest that, among our participants, cybersecurity had a narrow, specific definition related to adversarial hacking, a definition they are likely using in deciding whether it is something they want to pursue.

4.3 Perception of people in the field

4.3.1 *AI*. Participants expressed the belief that the people pursuing AI are seen as *“cool”* and highly intelligent. For example, while discussing which specializations within CS are considered prestigious, Nora noted that *“AI/ML is really trending. So if you do that, then people think you’re really cool.”* The idea of AI being *“cool”* is also seen in Quincy’s response to the same question about prestigious specializations, where he listed the current *“buzzwords”*: *“Today’s buzzwords are AI, ML, VR, future, crypto, blockchain or whatnot. Yeah. Um, so obviously that’s gonna attract a lot of buzz towards the AI/ML courses.”* Nora and Quincy’s comments suggest that, for some students, feeling that people who do AI are *“cool”* and prestigious may contribute to an interest in pursuing AI. People studying or working in AI were also viewed as particularly knowledgeable and capable, likely related to the belief of it being a highly challenging field. Kasey, when asked about any stereotypes of those who do AI, responded simply, *“That they’re really smart.”* Other participants expressed similar sentiments, such as Diane describing the field as *“intimidating.”* In response to a question regarding what kind of person is encouraged to do AI, Idris identified that *“they’ve been coding for ages, and they know all the software things,”* and Nora felt that

“strong-minded” people are encouraged to go into AI. These quotes suggest the belief that, to pursue AI, one must be very intelligent and capable, a suggestion that may deter less confident students.

4.3.2 Cybersecurity. Among our participants, the primary perception of the people who work in cybersecurity was that they are also very intelligent and are usually men. For instance, Farah noted that “cybersecurity seems like one of those fields where you have to be super duper smart to succeed.” Students seemed to believe that you must be particularly intelligent to do cybersecurity because mastering the content of the field is inherently impressive. Gail suggests this when she reported hearing the stereotype that “security people are like super smart because you can do security.” A second stereotype Gail mentioned encountering is that cybersecurity is “very male-dominated.” Other participants shared this impression, such as Esha, who explained that her mother had told her that cybersecurity is “in terms of gender, [an] imbalanced field.” Mahir also acknowledged having heard a similar stereotype when asked what kind of person is encouraged to go into cybersecurity, pointing out that cybersecurity has the same “bias towards men” that he described as existing in computing generally:

Well, besides, you know, the bias towards men in general in computer science, I don't think any other biases exist [in cybersecurity] that I'm aware of... Yeah, I think there [is] this new connection between like, [a] hacker is a man usually, which is, I guess, unnecessary, cause it's really quite accessible or it should be accessible to anyone.

In pointing out that cybersecurity “should be accessible to anyone,” Mahir appears to be highlighting that this is not the current perception of the field.

Notably, cybersecurity was not as popular among our participants as AI and lacked the perception of being “cool.”

4.4 Impact on society

4.4.1 AI. Some of our participants cited the belief that they could use AI to solve a large variety of problems, indicating a large potential for impact on society. For Idris, this was a specific reason for their decision to take courses in AI, because the problems they wanted to work on “can be solved with AI... I think AI has the ability to solve a lot of really interesting problems.” Although they do not specify which problems these are, it is clear that they see AI as a field that can be applied to many different contexts. Similarly, Parth explains that “[AI and ML] is like a very applicable (sic) and like, every business is trying to see how they can leverage machine learning to make their lives easier and produce more profit.” In Parth’s view, not only can AI be used in many different contexts, but this utility is the reason it is an in-demand skill among employers.

However, several participants also expressed concern about the societal impact AI could have. Henry described these as “moral gray areas” in his discussion of what stereotypes he had heard about AI:

I don't know if it's really a stereotype, but there's a lot of negative effects that AI can have, whether it's reinforcing prejudice on like Google and stuff or, well, Tesla self-driving cars running into people, but yeah, there's a lot of moral gray areas when it comes to AI.

In this quote, Henry pointed to bad societal outcomes that could happen as a result of working in AI, which are likely to factor into his decision-making when determining whether he wants to work in the field himself. Concern about AI’s societal impact is also suggested by Farah’s apparent frustration when sharing her perception of the ethical concerns with using AI and the people who work in it:

I feel like people are, like, yeah, we want to eliminate bias from AI and all this stuff. But like, I think the other thing that they have to acknowledge is that humans are biased, and, if you're training an algorithm or model or whatever it is, it's also going to be biased.

Other participants shared similar concerns about the use of AI, such as when Parth said that AI “can and has, intentionally and inadvertently, been harmful.” Although our participants did not cite these ethical concerns as a reason to avoid pursuing AI, it is likely that many of them are weighing the impact of AI on society in their own decision of whether to pursue the field.

4.4.2 Cybersecurity. Multiple participants explained that cybersecurity is a useful skill because of the importance of secure data in the world today, highlighting cybersecurity’s potential for societal impact. Lillian, when asked whether cybersecurity should be a required course in the undergraduate CS curriculum, responded that “I do see its utility. In the future—I mean, already data privacy is so important,” suggesting that cybersecurity’s utility is tied to the value it can bring people or companies, by keeping their data private. This value is especially important as technology becomes more prevalent in daily life, as Gail suggested when asked about why she felt that cybersecurity is an important field. She noted that “it's very important to make all the system[s] secure because a lot of things are moving online, like credit card[s], online transaction[s], e-commerce.” Similarly, Henry believed that cybersecurity is a “fairly well-regarded area because it's protecting the security of the internet.” His suggestion that cybersecurity is valued “because” of “the security of the internet” further reinforces the idea that cybersecurity skills are valued because they can significantly impact society.

We note participants did not express ethical concerns about cybersecurity work, despite multiple references to “hacking.”

5 DISCUSSION

Our findings extend prior work and indicate several ways in which student perceptions of AI and cybersecurity may hinder efforts to broaden the participation of students from HUGs.

Our results suggest that students have similar preconceptions about AI and cybersecurity as about computing in general. In particular, participants’ beliefs reflected stereotypes about the innate brilliance and masculine culture needed to participate in AI and cybersecurity. These beliefs have been shown to deter students from HUGs from computing [5, 21, 23]. In addition, the perception that AI requires mathematical skill may have a particularly detrimental effect on the participation of women because prior research has documented that women have lower math self-efficacy than men and that this difference affects their interest in pursuing math [12]. Participants’ concerns about AI’s unintentional negative societal impact may align with the belief that work in computing does not satisfy the goal of having a positive impact on society, which is a

goal more often endorsed by students from HUGs [10]. This suggests that students with goals to help society may be less likely to pursue AI due to their belief that working in AI could have a more negative societal impact. Taken together, our work extends our collective knowledge about field-specific perceptions and suggests a particular need to highlight the ways in which work in AI and cybersecurity extends beyond stereotypes of requiring innate brilliance, involving math, and having a negative impact on society.

Our findings align with prior work that students at the undergraduate level understandably do not have complete or accurate knowledge of specializations within computing, as some of their beliefs about AI and cybersecurity are not supported by experts. For example, participants strongly associated cybersecurity with systems programming, but the NICE Framework lists a number of cybersecurity “work roles” whose responsibilities include computing skills that are *not* systems programming, such as the *Warnings Analyst*, who “[c]ollects, processes, analyzes, and disseminates cyber warning assessments” [1]. Similarly, although participants suggested that the main work of AI is training models, AI hiring managers report that this is a relatively small subset of the work [2]. Expanding student perceptions of the skills required by and the work of AI and cybersecurity may be crucial to broadening participation in these fields, because systems programming and training models may seem disconnected from societal impact, a factor shown to be important to students from HUGs when selecting a field [10].

Participants’ perceptions of AI and cybersecurity are inconsistent with student experiences reported by Hewner, wherein students did not seem to have preconceived notions about a field prior to taking a course in it [14]; rather, many of our participants had narrow views of what AI and cybersecurity were like even before taking a course in them. A possible explanation is that prior exposure to these fields in media may be shaping student perceptions, as documented in prior work [13, 31, 32]. Our results also suggest that expanding student views of cybersecurity and AI cannot only take place in courses covering those subjects, as not all students will enroll in those courses.

6 IMPLICATIONS FOR PRACTICE

The results of our work suggest that many of participants’ beliefs about AI and cybersecurity are likely to differentially impact students from HUGs, as these views mirror existing stereotypes that have been shown to deter students from HUGs in computing. Thus, in order to broaden participation in these fields, there is a need to expand students’ views of what AI and cybersecurity entail. Our work suggests that interventions to expand these views cannot be restricted to AI or cybersecurity classrooms and indicates a need to explore venues such as student clubs, conferences, introductory CS courses, and other places where students who do not have a pre-existing interest in AI and cybersecurity may be found. Strategies to enhance students’ exposure to the breadth of career options in these two fields could include inviting speakers whose work does not fit the stereotype of AI or cybersecurity, showcasing projects that highlight how these fields can benefit society, and sharing the NICE Framework in courses. In addition, future work in research and teaching may aim to challenge the existing stereotypes about these fields and computing generally, such as an intervention in

a CS1 course designed to encourage a growth mindset, the belief that intellectual ability is not fixed and can grow [36], which could dispel the belief that these fields require innate brilliance.

7 LIMITATIONS AND FUTURE WORK

We note that our work is limited due to the influence of the institution’s specific course offerings, which are likely to shape students’ perceptions. The cybersecurity course at this institution requires a systems programming course as a prerequisite, which may explain students’ beliefs that cybersecurity work is closely tied to systems programming; this belief may not be shared by students at other institutions. The lack of sufficient cybersecurity course offerings and its impact on student perceptions is a previously documented concern [6], as most of these courses focus primarily on building secure systems rather than the full depth of cybersecurity such as risk analysis & management, policy & law, penetration testing, and secure software engineering [1].

An additional limitation of our work is that participant responses to our questions may have been influenced by their unwillingness to share a perspective that may not appear to be socially desirable, such as a stereotype based in racial/ethnic identities. In order to mitigate this, our interview questions allowed for the conflation of participants’ own views with beliefs they had heard expressed by others; as a consequence, our findings may reflect participants’ impression of popular beliefs, rather than their own perceptions of AI and cybersecurity.

Our work’s relatively small sample size limits our ability to conduct comparisons between groups, such as comparing views held by men versus women. However, we note that the existence of these perceptions and stereotypes impacts broadening participation efforts regardless of who expresses these views.

Future work in this area may investigate whether the beliefs held by our participants are shared by students at different institutions, such as those attending institutions that offer a greater number and variety of cybersecurity courses.

8 CONCLUSION

In this study, we interviewed computing undergraduate students in order to investigate their perceptions of AI and cybersecurity, particularly regarding the work, people, and societal impact of these fields. We found that student views did not always reflect the reality of working in AI or cybersecurity as described by experts and often reinforced stereotypes about computing. Our work suggests that there is a need to expand students’ views of what working in AI and cybersecurity could be like, such as by demonstrating that they can be used for societal good. This need is particularly important in any effort to broaden the participation in these fields of students from HUGs, as many of the beliefs students held have been shown to exacerbate the existing inequitable representation of gender and racial/ethnic groups in computing.

9 ACKNOWLEDGEMENTS

This work is partially funded by the National Science Foundation, Grant Nos. 2113954 and 2113955.

REFERENCES

[1] 2021. The Workforce Framework for Cybersecurity (NICE framework) Work Roles. <https://niccs.cisa.gov/about-niccs/workforce-framework-cybersecurity-nice-framework-work-roles>

[2] Author(s). In preparation. Title anonymized for review. (In preparation).

[3] Daphne Barretto, Julienne LaChance, Emanuelle Burton, and Soohyun Nam Liao. 2021. Exploring Why Underrepresented Students Are Less Likely to Study Machine Learning and Artificial Intelligence. In *Proceedings of the 26th ACM Conference on Innovation and Technology in Computer Science Education V. 1 (ITiCSE '21)*. Association for Computing Machinery, New York, NY, USA, 457–463. <https://doi.org/10.1145/3430665.3456332>

[4] Sapna Cheryan, Victoria C. Plaut, Caitlin Handron, and Lauren Hudson. 2013. The Stereotypical Computer Scientist: Gendered Media Representations as a Barrier to Inclusion for Women. *Sex Roles* 69, 1 (July 2013), 58–71. <https://doi.org/10.1007/s11199-013-0296-x>

[5] Sapna Cheryan, Sianna A. Ziegler, Amanda K. Montoya, and Lily Jiang. 2017. Why are some STEM fields more gender balanced than others? *Psychological Bulletin* 143, 1 (2017), 1–35. <https://doi.org/10.1037/bul0000052> Place: US Publisher: American Psychological Association.

[6] CloudPassage. 2016. CloudPassage Study Finds U.S. Universities Failing in Cybersecurity Education. <https://www.globenewswire.com/news-release/2016/04/07/1312702/0/en/CloudPassage-Study-Finds-U-S-Universities-Failing-in-Cybersecurity-Education.html>.

[7] Code.org. 2022. 2022 State of Computer Science Education. <https://advocacy.code.org/stateofcs>.

[8] Computing Research Association. 2021. The CRA Taulbee Survey. <https://cra.org/resources/taulbee-survey/>.

[9] Cyberseek. [n. d.]. <https://www.cyberseek.org/>

[10] Amanda B. Diekman, Elizabeth R. Brown, Amanda M. Johnston, and Emily K. Clark. 2010. Seeking Congruity Between Goals and Roles: A New Look at Why Women Opt Out of Science, Technology, Engineering, and Mathematics Careers. *Psychological Science* 21, 8 (Aug. 2010), 1051–1057. <https://doi.org/10.1177/0956797610377342> Publisher: SAGE Publications Inc.

[11] Brianna Dym, Namita Pasupuleti, Cole Rockwood, and Casey Fiesler. 2021. "You don't do your hobby as a job": Stereotypes of Computational Labor and their Implications for CS Education. In *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education (SIGCSE '21)*. Association for Computing Machinery, New York, NY, USA, 823–829. <https://doi.org/10.1145/3408877.3432396>

[12] Catherine Good, Aneeta Rattan, and Carol S Dweck. 2012. Why do women opt out? Sense of belonging and women's representation in mathematics. *Journal of personality and social psychology* 102, 4 (2012), 700.

[13] Damian Gordon. 2010. Forty years of movie hacking: considering the potential implications of the popular media representation of computer hackers from 1968 to 2008. *International Journal of Internet Technology and Secured Transactions* 2, 1/2 (2010), 59. <https://doi.org/10.1504/IJITST.2010.031472>

[14] Michael Hewner. 2014. How CS undergraduates make course choices. In *Proceedings of the tenth annual conference on International computing education research (ICER '14)*. Association for Computing Machinery, New York, NY, USA, 115–122. <https://doi.org/10.1145/2632320.2632345>

[15] Yin Kiong Hoh. 2009. Using Notable Women in Environmental Engineering to Dispel Misperceptions of Engineers. *International Journal of Environmental and Science Education* 4, 2 (April 2009), 117–131. <https://eric.ed.gov/?id=EJ884388> Publisher: International Consortium for the Advancement of Academic Publication.

[16] (ISC)2. 2020. *Women in Cybersecurity: Young, Educated, and Ready to Take Charge*. Technical Report. <https://www.isc2.org/-/media/isc2/research/isc2-women-in-cybersecurity-report.ashx> Library Catalog: www.isc2.org.

[17] J John and Martin Carnoy. 2017. Race and gender trends in computer science in the Silicon Valley from 1980–2015. (2017). https://cepa.stanford.edu/sites/default/files/JohnCarnoy_Sept2017.pdf

[18] Moritz Kreinsen and Sandra Schulz. 2021. Students' Conceptions of Artificial Intelligence. In *The 16th Workshop in Primary and Secondary Computing Education*. Number 14. Association for Computing Machinery, New York, NY, USA, 1–2. <https://doi.org/10.1145/3481312.3481328>

[19] Ram Shankar Siva Kumar, Magnus Nyström, John Lambert, Andrew Marshall, Mario Goertzel, Andi Comissoneru, Matt Swann, and Sharon Xia. 2020. Adversarial machine learning-industry perspectives. In *2020 IEEE Security and Privacy Workshops (SPW)*. IEEE, 69–75.

[20] Audra Lane, Ruth Mekonnen, Catherine Jang, Phoebe Chen, and Colleen M Lewis. 2021. Motivating literature and evaluation of the teaching practices game: Preparing teaching assistants to promote inclusivity. In *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education*. 816–822.

[21] Sarah-Jane Leslie, Andrei Cimpian, Meredith Meyer, and Edward Freeland. 2015. Expectations of brilliance underlie gender distributions across academic disciplines. *Science* 347, 6219 (Jan. 2015), 262–265. <https://doi.org/10.1126/science.1261375> Publisher: American Association for the Advancement of Science.

[22] Colleen M. Lewis, Ruth E. Anderson, and Ken Yasuhara. 2016. "I Don't Code All Day": Fitting in Computer Science When the Stereotypes Don't Fit. In *Proceedings of the 2016 ACM Conference on International Computing Education Research (ICER '16)*. Association for Computing Machinery, New York, NY, USA, 23–32. <https://doi.org/10.1145/2960310.2960332>

[23] Colleen M Lewis, Ken Yasuhara, and Ruth E Anderson. 2011. Deciding to major in computer science: a grounded theory of students' self-assessment of ability. In *Proceedings of the seventh international workshop on Computing education research*. 3–10. <https://doi.org/10.1145/2016911.2016915>

[24] LinkedIn. 2022. US Jobs on the Rise Report. <https://business.linkedin.com/talent-solutions/resources/talent-acquisition/jobs-on-the-rise-us>

[25] Jane Margolis. 2017. *Stuck in the Shallow End, updated edition: Education, Race, and Computing*. MIT press.

[26] Sharan B Merriam and Elizabeth J Tisdell. 2015. *Qualitative research: A guide to design and implementation*. John Wiley & Sons.

[27] National Center for Women and Information Technology. 2021. By The Numbers. <https://ncwit.org/resource/bythenumbers/>.

[28] William Newhouse, Stephanie Keith, Benjamin Scribner, and Greg Witte. 2017. National Initiative for Cybersecurity Education (NICE) Cybersecurity Workforce Framework. *NIST Special Publication 800*, 2017 (2017), 181. <https://doi.org/10.6028/NIST.SP.800-181r1>

[29] Department of Homeland Security. 2021. <https://www.dhs.gov/news/2021/11/15/dhs-launches-innovative-hiring-program-recruit-and-retain-world-class-cyber-talent>

[30] Anne Ottenbreit-Leftwich, Krista Glazewski, Minji Jeon, Cindy Hmelo-Silver, Bradford Mott, Seung Lee, and James Lester. 2021. How do Elementary Students Conceptualize Artificial Intelligence?. In *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education (SIGCSE '21)*. Association for Computing Machinery, New York, NY, USA, 1261. <https://doi.org/10.1145/3408877.3439642>

[31] Leila Ouchchy, Allen Coin, and Veljko Dubljević. 2020. AI in the headlines: the portrayal of the ethical issues of artificial intelligence in the media. *AI & SOCIETY* 35, 4 (Dec. 2020), 927–936. <https://doi.org/10.1007/s00146-020-00965-5>

[32] Ehsan Saffari, Seyed Ramezan Hosseini, Alireza Taheri, and Ali Meghdari. 2021. "Does cinema form the future of robotics?": a survey on fictional robots in sci-fi movies. *SN Applied Sciences* 3, 6 (May 2021), 655. <https://doi.org/10.1007/s42452-021-04653-x>

[33] Johnny Saldaña. 2021. *The coding manual for qualitative researchers*. SAGE publications Ltd.

[34] Andrew K. Shenton. 2004. Strategies for ensuring trustworthiness in qualitative research projects. *Education for Information* 22, 2 (Jan. 2004), 63–75. <https://doi.org/10.3233/EFI-2004-22001> Publisher: IOS Press.

[35] Jonathan Sillito. [n. d.]. <http://www.saturateapp.com/>

[36] David S Yeager, Paul Hanselman, Gregory M Walton, Jared S Murray, Robert Crosnoe, Chandra Muller, Elizabeth Tipton, Barbara Schneider, Chris S Hulleman, Cintia P Hinojosa, et al. 2019. A national experiment reveals where a growth mindset improves achievement. *Nature* 573, 7774 (2019), 364–369.