IJHSR
ijhighschoolresearch.org
RESEARCH ARTICLE

Decrypting Trust: Public Confidence in AI-Driven Cybersecurity Systems

Aditi B. Kothari

Johns Creek High School, 5575 State Bridge Rd, Johns Creek, GA 30022, USA; abkothari09@gmail.com
Mentor: Rajiv Garg

ABSTRACT: Artificial Intelligence has become increasingly integrated into companies, personal lives, and even entire industries, prompting many to form opinions on the topic. This study tests the hypothesis that as an individual's knowledge of the workings of Artificial Intelligence grows, their trust in the use of this innovation in cybersecurity declines. I analyzed its relationship in four main fields: general information, non-sensitive information, private data, and financial data. The results of this study show that most individuals did not change their opinions after learning more about the AI connection and seemed to trust humans to encrypt their data over AI-generated code. However, the linear regression analysis also displays that American Males often changed their opinion as they learn more about the specifications of AI-generated code.

KEYWORDS: Systems Software, Cybersecurity, AI Trust; Data Encryption, Human vs. AI.

Introduction

Today, Artificial Intelligence, also known as AI, is being implemented for use by both corporations and individuals. Numerous companies have manufactured their own AI tools, but the most widely used by the public remains ChatGPT, a chatbot developed by OpenAI. However, the question of whether and how to integrate it into cybersecurity is common, as questions about the integrity and effectiveness of AI come to the forefront of many conversations.

Currently, cybersecurity— also known as the protection of electronic data—faces ongoing problems that include human error, a shortage of professionals, and the constant evolution of cyberattacks.[1] Cyberattacks have become more complicated to test, which is only amplified by the lack of professionals, because weak points in the online system are harder to track.

This field also offers a new way for AI to be implemented, improving threat detection, automating responses, and providing business insights.[1] However, the probability of implementing widespread AI use in cybersecurity would be low if people luck trust its use, causing a ripple effect towards organizations and other systems.

Nevertheless, AI is still currently being used in cybersecurity. Specifically, it is most used in a dynamic threat detection system that isolates and responds to the anomalies that attack the system.[2] It is also essential for improving incident response time and predictive analysis to fight against cyberattacks preemptively. Arif and Khan also identify the prospects of artificial intelligence in this field, which include enhancing real-time detection and reaction response to changing risks as continuous learning becomes more advanced.

This paper examines the relationship between the public's understanding of how AI is utilized in cybersecurity and their opinions on the topic. This is to assess the feasibility of AI in cybersecurity, as if the public does not prefer it, then its implementation may harm company profits. Specifically, I study the impact of increased knowledge about how connections to the ChatGPT API in cybersecurity use correspond with people's level of trust compared to a human-generated code.

To determine the public's opinion and trust in using AI over humans, I conducted a survey where applicants had to answer questions about specific instances when they would prefer an AI or a Human generate a code to encrypt their messages. The survey consist-ed of 4 sections, each revealing more information about the differences between AI-generated code for encryption (AI-enabled encryption) and Human generated code for encryption (Human-enabled encryption), including statistics such as speed tests, encryption codes, and internet usage. Linear regression models were employed to analyze the survey results.

These steps were taken to test the hypothesis that people's trust in AI decreases as they learn more about how it works. In other words, people are more likely to trust an AI over a human to manage their cybersecurity if they do not know how the AI works. This hypothesis was founded in other research that shows how trust erodes as algorithmic awareness increases due to perceived risk. An article published by the Association for Computing Machinery found that there were higher privacy concerns and perceived risks with higher awareness.[3] These concerns are also associated with less trust. Shin *et al.* After conducting this study, however, I found that most people generally did not change their opinion on the role of AI in cybersecurity, as most of them preferred to have human developers handle their data security. American males acted as the exception to this trend.

My research builds upon past studies by applying them to a specific field of study. In the past, papers have primarily focused on the general public's perception of Artificial Intelligence. Kelley *et al.* conducted public opinion polls to determine the perception of AI in different countries and across various sentiments.[4] AI was described as exciting, useful, worrying, or

DOI: 10.36838/IJHSR86.93

futuristic. In their study, they find that 22.7% of their participants described AI as worrying. This trend was particularly prevalent in developed countries like Australia, Canada, the United States, and France. In many developing countries, such as Nigeria and India, however, the perception of AI is majorly positive, with most describing it as "exciting" or "futuristic."[5] These statistics show the wide range of perceptions of the future of Artificial Intelligence, possibly affecting further depictions of its use in even more serious situations.

Other studies focus on how Artificial Intelligence is written in association with certain words, such as "Ethical Concerns for AI" or "Work." In a study by Fast and Horvitz, they investigate these correlations and find that a "Loss of Control" has recently become more prevalent in articles about AI.[5] The findings of their study include an increase in the concern that AI use for human work is seen in a negative connotation, but also positive implications, as concerns over a lack of progress are becoming less prevalent. My research will build upon this by reflecting the opinions of people as they learn more about AI. Additionally, this research will apply this concept to a specific field of study, cybersecurity.

Research shows varied perspectives on how different genders percieve AI, with some reporting that women can be more skeptical of AI in high-stakes environments. In a general note, women have been reported to have significantly higher AI anxiety and overall lower positive attitudes towards AI, with there being a gender gap in how people perceive AI.[6] Generally, women are less confident in this situation and have less favorable attitudes towards it.[3] Bialy *et al*. show that AI is also perceived as having more risk in domains associated with control, surveillance, or replacing human judgment. These would all take place with an AI-driven cybersecurity system eliciting more concern for it.[7] These varying results can show an interesting way to see how individuals think of AI in a cybersecurity world, especially as they learn more about it throughout the course of the study.

Methods

To determine human trust in AI before and after they learn how it works, we began by creating code to gather statistics on the difference between AI and human-developed code. Both codes attempted to encrypt the word "Aditi" with a 16-bit key of "1234567890123456" in Advanced Encryption Standard, also known as AES. However, the AI code utilized the ChatGPT Application Programming Interface (API) to execute the encryption, whereas the human-generated code was fully executed locally. According to the results of the speed test (Table 1), the human-generated encryption code appears to be faster, likely due to the API connection to ChatGPT taking longer. However, the AI code generation was also significantly more volatile, ranging from 3.42 seconds to 0.014 seconds.

Table 1: Encryption Speed Comparison: Human-Generated Code vs. AI-Generated Code. This table displays the encryption speed in seconds for five tests of both code versions. It also includes the download and upload speed during the test in Mbps. The time for the AI encryption does not include the time it took to connect to the API, and every encryption test was run with the word "Aditi" with a 16-bit key of 1234567890123456. All API connection encryption speeds are adjusted to account for latency generated by the network.

Test Number	Human-Generated Encryption Speed	AI API Connection Encryption Speed	Download Speed (mbps)	Upload Speed (mbps)
1	0.0049	1.20	179.48	22.95
2	0.0038	1.11	309.28	23.08
3	0.0030	3.42	252.86	23.29
4	0.0038	0.014	251.41	23.11
5	0.0040	1.65	251.41	23.11

I also ran the AI Encryption tests using Ollama on my local machine, which is run and developed by OpenAI. Over the course of 5 tests, the times for an AES encryption averaged around 11.962 seconds. This is slower than the API connection encryption speed because my local machine does not have the necessary infrastructure to run an LLM model quickly. It is important to note that if companies were to use AI to encrypt their data at a larger level, they would need to invest in adequate infrastructure. Due to these facts, we ran the following studies using the statistics located in the table.

I created a Google Form to collect individual results from participants on their trust in AI to encrypt general data, non-sensitive data, private information, and sensitive financial information. The participants went through 4 different rounds, gaining more information with each round. All questions followed a similar format, asking if the respondent would "trust a human-generated algorithm or an AI-generated algorithm to encrypt" specific types of information. For instance, a question they were asked was, "Would you trust a human-generated algorithm or an AI-generated algorithm to encrypt information?" For each of the following questions, the word "information" was changed to the specific category that the question pertained to. The forms began with a baseline set of questions that repeated as more information was gathered on the speeds, the code, and how the internet is used for AI encryption.

When interpreting the data, the responses were filtered by location, separating those from the USA and International respondents. Most respondents answered 'Human' when asked if they would prefer a human or AI to encrypt their information (Figure 1). However, the numbers roughly evened out when asked how they preferred non-sensitive data to be encrypted (Figure 2). Eventually, however, those statistics on the choice changed as users had more options provided to them (Figures 3 and 4).

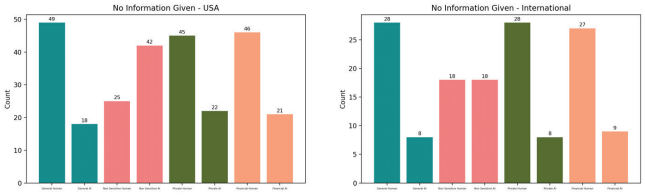

Figure 1 and 2: Number of respondents who answered Human or AI for every type of question without any information, and sorted by location. These graphs display the numbers of those who chose AI or Human and how the numbers for people who chose it changed over time. Both graphs are filtered by location between international and American respondents.

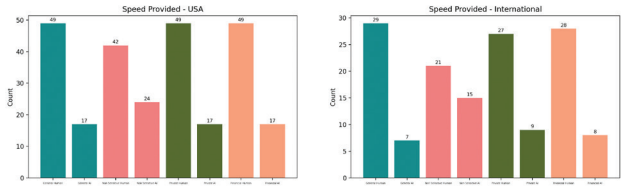

Figure 3 and 4: Number of respondents who answered Human or AI for every type of question, with the speed provided, and sorted by location. This table is similar to Figures 1 and 2. However, it encloses the respondents who chose AI-generated encryption versus Human-generated encryption after being provided the speed it takes to encrypt.

Results and Discussion

The overall results of the linear regression analysis indicate a strong preference for human encryption among respondents, particularly among females and international participants. In addition, these opinions remained largely stable throughout the study. In contrast, American males had the largest variability shifting across from human to AI encryption throughout the survey.

After the data were collected, they were analyzed using a linear regression model to identify patterns in the data. This model was used because it can help analyze the change in one response to the next throughout each type of question, and helps to interpret the relationship better. It helps to determine a baseline in the relationships between the variables, which can be analyzed further later. In addition, the linear model made it easier to combine categories of people, including gender, nationality, and previous responses to a question. It was also important to see the strength of these relationships to determine if there was truly a trend in the data. The categorical options of "Human" and "AI" were translated to 1 and 0, respectively. Additionally, Males were 1, Females were 0, Americans were 1, and international respondents were 0. In the table, the larger the magnitude of the coefficient, or the absolute value of the correlation coefficient, the more likely the group of people was to choose either Human or AI. The preference is determined by whether it is negative (AI) or positive (Human).

Table 2: Regression of Trust in Humans vs. AI based on Speed in Correlation with Demographics. The default response to the benchmark questions was "Human," indicating that every positive correlation coefficient is attributed to a higher likelihood of choosing "Human" over "AI." The opposing values to "United States" and "Male" are "International" and "Female" respectively. The table contains primarily the statistically significant results of the model.

	(1)	(2)	(3)	(4)
VARIABLES	General	Non-sensitive	Private	Sensitive
United States	0.400***	0.562***	0.400***	0.400***
	(0.146)	(0.103)	(0.143)	(0.137)
Male	0.25	0.333***	0	0
	(0.163)	(0.106)	(0.185)	(0.154)
United States x Male	-0.496**	-0.511***	-0.165	-0.15
	(0.237)	(0.168)	(0.247)	(0.220)
Observations	103	103	103	103
R-squared	0.87	0.748	0.872	0.884

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 2 specifically attempts to determine the correlation between the demographics of responders and the answers they chose. Throughout the majority, however, it is seen that most Americans prefer to use an encryption system created by a human over one done by AI. However, changes become more apparent when the parameters are made much more specific. For instance, American males generally preferred using an AI encryption system over the alternative. This implies that American women prefer to use human encryption much more. This relationship is amplified much more when looking at the non-sensitive information, as American males are much more likely to put AI than for any other category. While this may be a correlation, it also suggests less importance is given to that region, as non-sensitive information included things already readily available online, such as Social Media posts. However, generally, throughout this question, the primary answer from Americans and males, exclusive of others, favoring encryption developed by humans over that generated by humans.

Table 3: Regression of Trust in Humans vs. AI based on Speed in Correlation to Baseline Questions. The default response to the benchmark questions was "Human," which means that every positive correlation coefficient is attributed to a higher likelihood of choosing "Human" rather than "AI." The opposing values to "United States" and "Male" are "International" and "Female" respectively. The answers under Baseline are set to filter to those who chose "Human" in the question before. The table contains primarily the statistically significant results of the model.

	(1)	(2)	(3)	(4)
VARIABLES	General	Non-sensitive	Private	Sensitive
Baseline General	1***			
	(0.109)			
Baseline General x United States x Male	0.587**			
	(0.288)			
Baseline Non-sensitive		0.900***		
		(0.130)		
Baseline Non-sensitive x US		-0.462**		
		(0.220)		
Baseline Private			0.875***	
			(0.113)	
Baseline Private x US			-0.328*	
			(0.197)	
Baseline Sensitive				1***
				(0.109)
Baseline Sensitive x US				-0.453**
				(0.189)
Observations	103	103	103	103
R-squared	0.87	0.748	0.872	0.884

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 3 presents the regression analysis of the respondents' answers to the "Speed" question in comparison to their responses to the Baseline questions. The purpose of this graph is to show the potential changes in opinion to help prove or disprove the hypothesis. However, most of the respondents' answers did not change with the addition of the "speed" timings. There are changes where specific sections of people, such as Americans who initially answered "Human" for the Baseline Non-sensitive, changed their answer to "AI." This can likely be attributed to people who thought that a slower time meant a more in-depth analysis, but there is no way to prove this thought.

Table 4: Regression of Trust in Humans vs. AI based on Code Correlation with Demographics. The default response to the benchmark questions was "Human," indicating that every positive correlation coefficient is attributed to a higher likelihood of choosing "Human" over "AI." The opposing values to "United States" and "Male" are "International" and "Female" respectively. The table contains primarily the statistically significant results of the model.

	(5)	(6)	(7)	(8)
VARIABLES	General	Non-sensitive	Private	Sensitive
United States	0.400**	0.437***	0.400***	0.400***
	(0.155)	(0.0987)	(0.148)	(0.146)
Male	0.500***	0.267**	0.333*	0.25
	(0.174)	(0.102)	(0.191)	(0.163)
US x Male	-0.592**	-0.396**	-0.498*	-0.337
	(0.252)	(0.162)	(0.255)	(0.233)
Observations	103	103	103	103
R-squared	0.859	0.747	0.865	0.873

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Similar patterns to the first question can be seen in this table (Table 4). Generally, the answers from the United States and

DOI: 10.36838/IJHSR86.93

Males (exclusively) indicated that they preferred a Human to encrypt their information over an AI. However, when examining the answers of the American Male population, they are more likely to report that they prefer AI encryption.

Table 5: Regression of Trust in Humans vs. AI based on Code in Correlation to Baseline Answers. The default response to the benchmark questions was "Human," which means that every positive correlation coefficient is attributed to a higher likelihood of choosing Human" rather than "AI." The opposing values to "United States" and "Male" are "International" and "Female" respectively. The answers under Baseline are set to filter to those who chose "Human" in the question before. The table contains primarily the statistically significant results of the model.

VARIABLES	(5) General	(6) Non-sensitive	(7) Private	(8) Sensitive
Baseline General	0.778***			
	(0.122)			
Baseline General x US	-0.578**			
	(0.220)			
Baseline General x Male	-0.580**			
	(0.235)			
Baseline General x Male x US	0.665**			
	(0.322)			
Baseline Nonsensitive		0.700***		
		(0.132)		
Baseline Private			0.750***	
			(0.138)	
Baseline Private x Male			-0.550**	
			(0.240)	
Baseline Sensitive				0.750***
				(0.121)
Base Sensitive x US				-0.550**
				(0.210)
Observations	103	103	103	103
R-squared	0.856	0.741	0.83	0.873

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 5 shows that most respondents had changes in their answers compared to the baseline. Specifically, once again, for the Americans, Males, and American Males. They were more likely to change their choice in areas regarding "Sensitive/ Financial Information," "Private Information," and "General Information."

Table 6: Regression of Trust in Humans vs. AI based on Code in Correlation to Speed. The default response to the benchmark questions was "Human," which means that every positive correlation coefficient is attributed to a higher likelihood of choosing "Human" rather than "AI." The opposing values to "United States" and "Male" are "International" and "Female" respectively. The answers under Baseline are set to filter to those who chose "Human" in the question before. The table contains primarily the statistically significant results of the model.

VARIABLES	(13) General	(14) Non-sensitive	(15) Private	(16) Sensitive
Speed General	1***			
	(0.0726)			
Speed General x US	-0.300**			
	(0.150)			
Speed General x Male	-0.400***			
	(0.142)			
Speed General x United States x Male	0.513**			
	(0.209)			
Speed Nonsensitive		0.889***		
		(0.106)		
Speed Private			1***	
			(0.0926)	
Speed x US			-0.300*	
			(0.178)	
Speed Private x Male			-0.456***	
			(0.168)	
Speed Private x US x Male			0.535**	
			(0.245)	
Speed Sensitive				0.900***
				(0.082)
Speed Sensitive x Male				-0.300*
				(0.155)
Speed Sensitive x US x Male				0.413*
Speed Private x US				(0.227)
Observations	103		103	103 103
R-squared	0.932		0.837	0.905 0.919

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

In comparison to the previous table, which attempted to find the correlation between the answers to the "Baseline" questions and the opinions after learning about the code and its inner workings, Table 6 repeats that process but compares the "Code" section with the "Speed Section." Most opinions remained unchanged throughout this questioning, except for those of males across all four questions. For private and general information, American respondents also showed a change in their opinion, as indicated in the regression model.

Table 7: Regression of Trust in Humans vs. AI based on Internet Correlation with Demographics. The default response to the benchmark questions was "Human," indicating that every positive correlation coefficient is attributed to a higher likelihood of choosing "Human" over "AI." The opposing values to "United States" and "Male" are "International" and "Female" respectively. The table contains primarily the statistically significant results of the model.

VARIABLES	(9) General	(10) Non-sensitive	(11) Private	(12) Sensitive
United States	0.800***	0.438***	0.800***	0.800***
	(0.163)	(0.105)	(0.174)	(0.153)
Male	0.750***	0.333***	0.333	0.500***
	(0.183)	(0.108)	(0.225)	(0.171)
US x Male	-0.935***	-0.309*	-0.545*	-0.737***
	(0.265)	(0.171)	(0.300)	(0.245)
Observations	103	103	103	103
R-squared	0.856	0.741	0.83	0.873

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 7 continues to show the trend in demographics, where, even as individuals gained more information about the insights into how the encryption styles worked, their opinions remained the same throughout. American males continued to prefer an AI encryption method over the alternative provided.

Table 8: Regression of Trust in Humans vs. AI based on Internet Correlation with Baseline Questions. The default response to the benchmark questions was "Human," indicating that every positive correlation coefficient is attributed to a higher likelihood of choosing "Human" over "AI." The opposing values to "United States" and "Male" are "International" and "Female" respectively. The answers under Baseline are set to filter to those who chose "Human" in the question before. The table contains primarily the statistically significant results of the model.

VARIABLES	(9) General	(10) Non-sensitive	(11) Private	(12) Sensitive
Baseline General	0.778***			
	(0.122)			
Baseline General x US	-0.578**			
	(0.220)			
Baseline General x Male	-0.580**			
	(0.235)			
Baseline General x Male x US	0.665**			
	(0.322)			
Baseline Nonsensitive		0.700***		
		(0.132)		
Baseline Private			0.750***	
			(0.138)	
Baseline Private x Male			-0.550**	
			(0.240)	
Baseline Sensitive				0.750***
				(0.121)
Base Sensitive x US				-0.550**
				(0.210)
Observations	103	103	103	103
R-squared	0.856	0.741	0.83	0.873

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

This table confirms the recurring patterns observed in previous tables, such as Tables 3, 4, and 5. While the general direction of respondents' opinions remained unchanged, the more specific categories all experienced a significant shift in opinion. Some instances of this situation have repeatedly occurred in questions regarding the security of private information, where males have changed their opinion from their previous response.

Table 9: Regression of Trust in Humans vs. AI based on the Internet in Correlation to Code. The default response to the benchmark questions was "Human," indicating that every positive correlation coefficient is attributed to a higher likelihood of choosing "Human" over "AI." The opposing values to "United States" and "Male" are "International" and "Female" respectively. The answers under Baseline are set to filter to those who chose "Human" in the question before. The table contains primarily the statistically significant results of the model.

VARIABLES	(17) General	(18) Non-sensitive	(19) Private	(20) Sensitive
Code General	0.818***			
	(0.0919)			
Code General x US	-0.568***			
	(0.191)			
Code Nonsensitive		0.750***		
		(0.128)		
Code Private			0.889***	
			(0.101)	
Code Private x US			-0.639***	
			(0.194)	
Code Sensitive				0.889***
				(0.0912)
Code Sensitive x US				-0.639***
				(0.175)
Code Sensitive x US x Male				0.523**
				(0.261)
Observations	103	103	103	103
R-squared	0.9	0.808	0.898	0.918

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Throughout this section of the forms, there were more changes in responses as seen in Table 9. Those who initially answered 'Human' in the US for both sensitive and non-sensitive information, which also included males, showed a higher tendency to change their responses from 'Human' to 'AI-encryption' over 'Human-encryption'. Males generally also tended to trust AI more after learning about how the internet connection of the API works over the local system of human encryption.

The results of this study show that females generally do not change their opinion, but they were also more likely to choose humans throughout the study. Specifically, after seeing the code and the speeds at which the AI worked, the original answers led them to choose human encryption over AI. International females were also shown as more likely to answer Human encoding for sensitive information after seeing the code. However, they were more likely to choose AI-encryption for the non-sensitive question in the same section.

Overall, these results indicate a trend of continuing to trust Human-generated code over AI-generated code as they learned more about the problem situation. However, males, and specifically American males, often changed their results from "Human" to "AI" encryption in specific circumstances after learning more about the differences in the encryption methods. This suggests that female participants, including international respondents, were generally consistent in their preferences humans as their responses.

■ Conclusion

This project aims to understand the public's view and trust in AI encryption for cybersecurity. My hypothesis was that people's trust in AI would decrease as they learned more about how the technology works. As technology advances, it is crucial to determine the popular belief towards AI in protecting data and other cybersecurity measures. Specifically, to identify the trends in the data, I coded an AES encryption both locally on my computer and through an API connection to ChatGPT. Information about those codes, which included the speed of encryption, the code, and a description of how the APIs worked, was provided to respondents who completed a form and answered the same questions as they learned more about how each code functioned.

The conclusions of the linear regression analysis do not support my initial hypothesis, as participants were not less trusting of AI as they learned more about it. Throughout the experiment, trust increased in a few participants but overall remained stable. Specifically, American Males acted as the exception and changed their opinions numerous times during the survey, including after they saw the code and the internet connection. Additionally, I found that throughout the survey, people tended to trust human developers over AI for most types of information.

In the professional world, the implications of this study can help inform the use of AI in cybersecurity. In financial organizations, the use of AI can be detrimental to customers' perception of the company, making it essential to emphasize the use of human developers. However, if your target audience is primarily male, emphasizing AI or incorporating it more into cybersecurity would be a better option. Although humans still generally have reservations about trusting AI, the results of this study can help companies decide when and where to emphasize the use of AI with specific target audiences.

A possible limitation of my results may be from the limitations of the forms. On the forms, respondents were only presented with the options of "Human" or "AI," which does not account for those who do not care which section encrypts their information. Additionally, the way this form was written does not account for partial trust, which could be addressed by using a scale to determine the degree of confidence in a certain encryption method. The regression model was also trained on a smaller number of participants, which increases the likelihood of the model being incorrect.

There are multiple ways to discuss this topic further. To start, individuals can conduct a similar project that expands on why people trust AI over Humans more. To do this, using a multiple-choice questionnaire where respondents were asked after every section why they chose the answers they did could be useful. Additionally, using a scale to determine how likely someone is to trust an AI over a Human for encryption could be useful, as it would also encompass those who do not care about how their information is encrypted. Ultimately, this topic can be revisited in the future to explore its relevance to the evolving field of cybersecurity. It can be researched what the public's opinions are on certain methods of AI encryption. In the future, new methods of AI implementation in this field may exist beyond encryption, which can be utilized for a more comprehensive research process.

■ Acknowledgments

This paper would not have been possible without the support of the following people. Thank you so much for your support and guidance.

First, I would like to thank my family for their continued support and motivation throughout this project. Mom and

DOI: 10.36838/IJHSR86.93

Dad, thank you both for your encouragement and enthusiasm. The time you spent helping me with concepts and getting respondents for surveys. Thank you to my sister Kriti, who has always been a constant source of help and guidance for me.

I would also like to thank all of the respondents to the survey. Without your help, this paper would not be possible.

References

1. Khan, M. I.; Arif, A.; Khan, A. R. A. The most recent advances and uses of AI in cybersecurity. https://www.journal.mediapublikasi.id/index.php/bullet/article/view/4540.
2. Biały, F.; Elliot, M.; Meckin, R. Perceptions of AI Across Sectors: A Comparative Review of Public Attitudes; arXiv preprint arXiv:2509.18233, 2025.
3. Russo, C.; Romano, L.; Clemente, D.; Iacovone, L.; Gladwin, T. E.; Panno, A. Gender differences in artificial intelligence: the role of artificial intelligence anxiety. Frontiers in Psychology 2025, 16, 1559457. https://doi.org/10.3389/fpsyg.2025.1559457.
4. Kelley, P.G. *et al*. (2021). Exciting, useful, worrying, futuristic: Public perception of artificial intelligence in 8 countries, Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society [Preprint]. doi:10.1145/3461702.3462605.
5. Fast, E.; Horvitz, E. Long-Term trends in the public perception of artificial intelligence. Proceedings of the AAAI Conference on Artificial Intelligence 2017, 31 (1). https://doi.org/10.1609/aaai.v31i1.10635.
6. Alasmari, A. A.; Alruwaili, R. F.; Alotaibi, R. F.; Youssef, I. K.; Asklany, S. A. Demographic influences on trust in artificial intelligence across cognitive domains: A statistical perspective. PLoS ONE 2025, 20 (11), e0331003. https://doi.org/10.1371/journal.pone.0331003.
7. Borges, A. F. S.; Laurindo, F. J. B.; Spínola, M. M.; Gonçalves, R. F.; Mattos, C. A. The strategic use of artificial intelligence in the digital era: Systematic literature review and future research directions. International Journal of Information Management 2020, 57, 102225. https://doi.org/10.1016/j.ijinfomgt.2020.102225.
8. SentinelOne. Top 5 Cyber Security Challenges. SentinelOne. https://www.sentinelone.com/cybersecurity-101/cybersecurity/cyber-security-challenges/.
9. Shin, D.; Kee, K. F.; Shin, E. Y. Algorithm awareness: Why user awareness is critical for personal privacy in the adoption of algorithmic platforms? International Journal of Information Management 2022, 65, 102494. https://doi.org/10.1016/j.ijinfomgt.2022.102494.

Authors

Aditi Kothari is a rising junior at Johns Creek High School and has had an interest in Computer Science since she was young. Over the past year, she has competed in a hackathon and other technology-based competitions with FBLA, such as Cybersecurity.

Dr. Rajiv Garg is a professor at Emory University. His research explores the economic and social impact of human-machine interactions. Professor Garg's work investigates information flow in digital platforms and networks, the role of technology in labor markets and entrepreneurship, and the business value of emerging technologies like artificial intelligence.