Vision: Comparison of AI-assisted Policy Development Between Professionals and Students

Rishika Thorat
Purdue University
rthorat@purdue.edu

Tatiana Ringenberg
Purdue University
tringenb@purdue.edu

***Abstract*—AI-assisted cybersecurity policy development has the potential to reduce organizational burdens while improving compliance. This study examines how cybersecurity students and professionals develop ISO29147-aligned vulnerability disclosure policies (VDPs) with and without AI. Through this project, we will evaluate compliance, ethical accountability, and transparency of the policies through the lens of Kaspersky's ethical principles.**

Both students and professionals will produce policies manually and with AI, reflecting on utility and reliability. We will analyze resulting policies, prompts, and reflections through regulatory mapping, rubric-based evaluations, and thematic analysis. This project aims to inform educational strategies and industry best practices for integrating AI in cybersecurity policy development, focusing on expertise, collaboration, and ethical considerations.

We invite feedback from the Usable Security and Privacy community on participant recruitment, evaluation criteria, ethical frameworks, and ways to maximize the study's impact on academia and industry.

I. Introduction

Around 70% of countries across the world have enacted data protection and privacy laws (Apacible-Bernardo & Fischer, 2024), with frameworks such as the General Data Protection Regulation (GDPR) in the European Union and the newer Cyber Resilience Act (CRA) setting high standards for compliance and governance. These, and other emerging state regulations within the United States, require organizations not only to protect sensitive data, but also to quickly adapt policies to meet and align with the rapidly evolving regulatory landscape. This evolution challenges organizations in developing and maintaining compliant cybersecurity policies.

The field of artificial intelligence offers potential solutions to these challenges, particularly through large language models (LLMs) like ChatGPT. LLMs analyze regulatory data and generate policy drafts faster than humans, as seen in studies leveraging GPT-4 for the generation of governance, risk and compliance (GRC) policies (McIntosh et al., 2023) and ChatGPT for parliament governance frameworks (Lucke & Sander, 2024). Similarly, initiatives such as the UK Labour Party's Parlex AI (Smyth, 2025), which predicts parliamentary reactions to proposed policies, show how AI is being actively integrated into policy-making workflows. These examples show AI's growing role as a practical resource in governmental and industry contexts.

Despite these advancements, the role of AI in policy development has limitations in areas that require contextual understanding, ethical alignment, or user-centric considerations. Although an LLM might generate policies that align with current legal and compliance standards, there is little prior research on how well LLMs can adapt to emerging regulatory changes, especially given the lack of cohesive methodologies for assessing LLM risks in regulatory contexts (Goanta et al., 2023). These gaps emphasize the need for human-AI collaboration, merging ethical reasoning with computational power.

The growing adoption of AI in cybersecurity further highlights the importance of studying its role in policy development. The global market for AI in cybersecurity is projected to grow from approximately $24 billion in 2023 to about $134 billion U.S. dollars by 2030, indicating an increasing reliance on AI to address challenges (Borgeaud, 2024). This growth emphasizes the need to understand AI's integration into cybersecurity practices to prepare organizations and individuals.

This study examines how cybersecurity students and professionals address these challenges both independently and through collaboration with AI, in the context of VDPs. Vulnerability disclosure, governed by the ISO29147 standard, is critical for ensuring transparency, trust, and risk mitigation in cybersecurity. By focusing on ISO29147, this study aligns with an established and widely recognized framework, making its findings directly applicable to real-world practices. Understanding how these two participant groups approach this task is crucial: students represent the future workforce and must be prepared to navigate AI-enhanced environments, while professionals provide insight into current real-world practices and challenges.

The relevance of this research is in its ability to address critical questions about the collaborative role of AI in policy-making. While AI tools like Parlex and ChatGPT can streamline aspects of policy-making, they cannot replace the nuanced judgment and ethical reasoning provided by human experts. This study explores interactions between students, professionals, and AI to identify areas where human expertise and AI capabilities align, diverge, or complement one another.

To investigate these dynamics, we focus on the following

Symposium on Usable Security and Privacy (USEC) 2025
24 February 2025, San Diego, CA, USA
ISBN 979-8-9919276-5-9
https://dx.doi.org/10.14722/usec.2025.23040
www.ndss-symposium.org, https://www.usablesecurity.net/USEC/

overarching research questions:

- RQ1: How do students' and professionals' manually-developed policies compare in ethical, compliance, and transparency dimensions?
- RQ2: How do the prompts developed by students and professionals differ when constructing these policies using AI?
- RQ3: How do students' and professionals' AI-enhanced policies compare in ethical, compliance, and transparency dimensions?

By addressing these questions, this analysis contributes to the growing body of research on AI-assisted cybersecurity practices, offering insights into the educational and professional applications of AI tools. It will identify educational gaps and opportunities for curriculum improvements and provide actionable recommendations for professional development, ensuring that the next generation of cybersecurity professionals is able to leverage AI to create policies that are compliant, ethical, and transparent.

II. Literature Review

A. AI in Policy Generation

Large language models (LLMs), like GPT-4, have shown potential in streamlining cybersecurity policy generation by addressing the bottlenecks and inefficiencies of traditional approaches. For example, Ferrag et al., 2024 explored how LLMs could integrate real-time threat intelligence into policy drafts, enabling quick responses to emerging vulnerabilities. This capability is especially relevant in fast-evolving fields like cybersecurity, where delays in policy adaptation can amplify risks. Similarly, McIntosh et al., 2023 explored GPT-4's application in Governance, Risk, and Compliance (GRC) frameworks, particularly in ransomware mitigation. They found that, with structured prompts, GPT-4 was able to generate compliant policies that met technical standards, providing a starting point for human improvement.

Despite strengths, LLMs face limitations. Yigit et al., 2024 and Jawhar et al., 2024 highlighted LLM challenges in addressing ethics and aligning with evolving regulations. These findings show that, while AI systems do well with technical precision and efficiency, they tend to fall short with tasks that require long-term adaptability.

B. Human-AI Collaboration

Human expertise enhances AI-assisted policy generation, addressing adaptability and ethical gaps. Fragiadakis et al., 2024 proposed a framework showing how human guidance improves AI's interpretive capabilities, which also helps to align with organizational values. This aligns with Shilton et al., 2020, who, through role-playing simulations, illustrate how embedding human ethical considerations into AI-guided policies improves outcomes for privacy and user rights.

Cai et al., n.d. adds on to these findings by proposing iterative feedback loops between humans and AI systems, highlighting their potential to create sustainable and adaptable policies. This perspective is similar to that of Mitrou et al., 2021, who emphasized the importance of human oversight in resolving ambiguities, particularly when adapting policies to changing regulations. These studies show that AI excels in efficiency and compliance, but human collaboration mitigates LLM limitations to ensure transparency and adaptability.

Looking at these findings together, we see that AI-driven policy development needs to operate within frameworks that prioritize human involvement to address ethical blind spots and improve the overall interpretability of the policies. Our study builds on this prior work by analyzing user-AI policies against Kaspersky ethical principles (Kaspersky, 2020).

C. Compliance, Adaptability, and User Prioritization

Effective cybersecurity policies must balance compliance with accessibility and user-centeredness to improve adherence and trust. Veale and Edwards, 2018 highlighted the General Data Protection Regulation's (GDPR's) "right to explanation" as a mechanism for creating policies interpretable by both technical and non-technical users. This aligns with findings from Da Veiga, 2016, who showed that user-centric policy designs not only improve employee compliance but also contribute to a more stronger security culture of the organization. Perry and Uuk, 2019 discussed how user-centered approaches, based on usability principles, improved the effectiveness of AI-driven frameworks in meeting different stakeholder needs.

Korobenko et al., 2024 emphasized privacy-aware governance in ethical AI development. Similarly, Kelly et al., 2024 analyze the EU AI act, showing how user-centered compliance frameworks simplify regulatory requirements for nonexpert stakeholders. These studies highlight user-centered approaches as key to accessible and practical cybersecurity policies.

D. Prompt Engineering

Prompt engineering is a technique for customizing LLMs for policy-generation tasks, enhancing domain-specific effectiveness. Sahoo et al., 2024 highlighted few-shot and zero-shot prompting, which allow LLMs to meet regulatory requirements without extensive fine-tuning. Building on this, Wang et al., 2024 introduced LangGPT, a modular framework for structured prompts that improves the consistency and adaptability of policy templates.

Trad and Chehab, 2024 compare prompt engineering with fine-tuning for phishing detection, showing that prompts are resource-efficient but lack precision for specialized contexts. Combining prompt engineering with frameworks like LangGPT enables compliant, adaptable, and efficient policy generation, but human validation is needed to address ethical gaps and ensure policy longevity.

Our study will compare how cybersecurity professionals and students choose to develop prompts to generate policies.

III. Proposed Methods

This study will examine how policy development and AI-assisted policy development differ between professionals and students in cybersecurity. For the purposes of this study, we have chosen to have participants design vulnerability handling

policies which align with ISO29147. ISO29147 was chosen for this analysis because it provides clear, structured guidelines that are widely recognized in cybersecurity, making it an effective baseline for assessing policy development and compliance. This study will also evaluate how participants incorporate principles from Kaspersky's Ethical Principles of Vulnerability Disclosure, focusing on key aspects of transparency, ethical accountability, and predictability. By comparing professionals and students, this study aims to explore how expertise influences the process, outcomes, and perceptions of AI in policy development. The comparison between students and professionals aims to provide a dual perspective: understanding current practices among professionals while preparing students for the evolving demands of AI-assisted policy development.

A. Course Description

Students in this study will come from an existing Cyber Law and Ethics course at a large university. Students in the course are part of a cybersecurity undergraduate degree program, generally at the junior or senior level, with little to no prior training on policy development. There are approximately 100 students enrolled in the course each semester.

One of the primary aims of this course is to help students develop and adapt policies and security controls in response to emerging cybersecurity laws. As part of the course, students complete a module on developing policies and tracing them to specific regulations.

For their final project, students create VDPs aligned with the ISO29147 standard. Policies are created in groups to enhance the alignment with industry practices. This study will not only contribute to the broader literature on policy development with AI but also informs the instructional design of the Cyber Law and Ethics course.

Approval from the university's Institutional Review Board has been granted to collect the data for the study's purpose.

B. Professional Participant Recruitment

To assess policy development and adaptation in security professionals, we will identify participants with prior experience in policy development, compliance, or vulnerability disclosure or management. Criteria for selecting professionals will include years of experience, current and prior roles, and educational background.

We plan to recruit professionals through the university's existing industry networks and councils, professional organizations, and outreach through social media. As we aim to recruit participants with domain-specific policy experience, we will use a snowball sampling method to engage additional qualified participants beyond our initial outreach.

C. Policy Development Activities

For this activity, we seek to develop tasks aligning with real-world policy development challenges. To do this, both students and professionals will be presented with the same scenario designed to simulate a realistic cybersecurity context. The scenario will feature a detailed description of a fictitious company, including its industry, size, and key cybersecurity challenges. This ensures that all participants work from the same baseline which will allow meaningful comparisons of their approaches to developing VDPs aligned with the ISO29147 standard.

Once presented with the scenario, participants will be given access to the ISO29147 standard. Using the standard, participants will develop their own VDP for the fictitious company which aligns with ISO29147. Cybersecurity professionals will work alone, while the cybersecurity students will work in groups of four. The researchers will collect the resulting policies.

Following initial policy development, participants will construct and execute a prompt in ChatGPT to generate a ISO29147-aligned policies for their fictitious company. Students and professionals will use the same ChatGPT model to ensure consistency. The researchers will collect both the prompts and policies developed by the students and professionals.

Participants will next be asked to review Kaspersky's Ethical Principles of Vulnerability Disclosure, focusing on transparency, ethical accountability, and predictability. They will then be asked to modify their manually-developed policies manually to align with these principles. They will then use ChatGPT to modify their AI-enhanced policies based on the same principles. These principles were chosen because they are directly tailored to vulnerability disclosure while aligning with broader GRC principles commonly used in cybersecurity policy development. Although Kaspersky's framework is domain-specific, its core principles of trust, fairness, and transparency overlap with well-established ethical guidelines for AI and policy governance, like the Ada Lovelace Institute's work on participatory data stewardship. This framework provides practical and ethical guidance that is directly relevant to the policies created by participants. The researchers will collect the resulting modified policies.

Following completion of the activity, participants will be asked to reflect on the usability and reliability of ChatGPT for developing these policies. The reflections will then be collected by the researchers. The workflow of the participant activity is summarized in Figure 1.

All data collected from participants will be anonymized prior to analysis.

D. Data Analysis

The policies and prompts produced by participants will be evaluated using a combination of regulatory mapping and qualitative analysis. While the qualitative analyses in this study will primarily be conducted by the first author, approximately 20% of the data will be coded by a second annotator to ensure reliability. The annotators will meet consistently to discuss and address any disagreement in the qualitative results.

The participant-generated policies will be evaluated through mapping the requirements of ISO29147 to elements of the students' and professionals' policies, noting gaps. For this analysis, we will use a content analysis driven by the ISO29147 requirements.

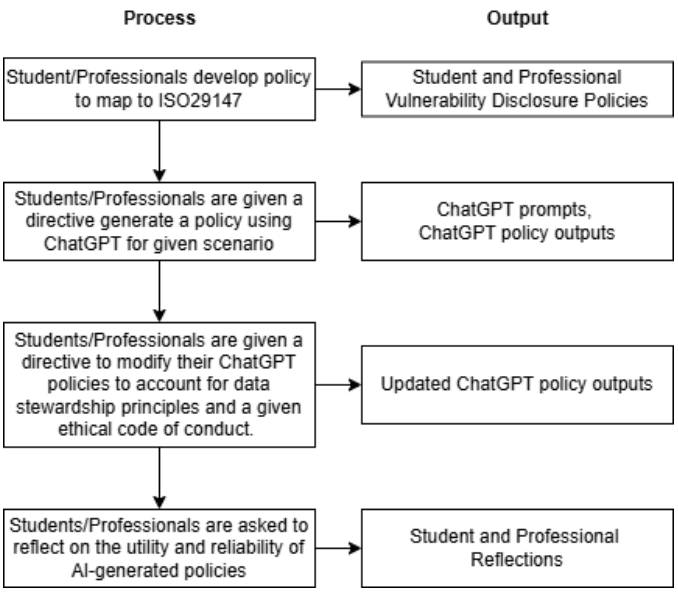

Fig. 1. Workflow of student and professional policy development activity.

Additionally, we will evaluate the prompts students and professionals used to generate policies using ChatGPT. We will note the characteristics of the fictitious scenario and ISO standard which each group of participants includes.

We will also analyze differences in how students and professionals adapt their policies to the chosen Kaspersky ethical principles. Changes will be assessed using a content analysis in which we deductively identify the aspects of the principles present in the document. Finally, we will conduct a thematic analysis to assess students' and professionals' perceptions of the utility and reliability of collaborating with AI to develop and adapt the VDPs.

E. Expected Outcomes

This study will provide valuable insights into policy development and education, including:

- Comparison of professionals' and students' approaches to developing vulnerability disclosure policies both with and without AI.
- The influence of expertise on the compliance of cybersecurity policies to standards and regulations.
- Comparison of how professionals and students interact with AI tools for policy creation.
- Participant perceptions of AI integration into policy development.
- Participant perceptions of the usability of AI for policy development.

IV. Limitations and Discussion Request for USEC

While designing the study we have anticipated a set of limitations of the proposed activity. We seek to discuss such limitations within the Symposium on Usable Security and Privacy to identify possible mitigation strategies and improvements.

The first limitation revolves around the grouping of participants. We plan to have individual professionals engage in the exercise, while students will engage in the exercise as part of an overall group project.

Students developing policy in a group aligns more closely with the collaborative nature of the policy development process in industry. However, this makes direct comparison to individual cybersecurity professionals difficult, as group dynamics and group discussions will likely influence how policies are developed and analyzed. This was a practical consideration as placing professionals into groups would significantly reduce the participant pool, as we anticipate cybersecurity professional teams will be difficult to recruit in large numbers. We seek discussion within the symposium on practical methods to mitigate this difference.

Another limitation is the potential generalizability, as this study includes only an assessment of a small subset of cybersecurity students and professionals in the development of VDPs and a specific ISO standard. Thus, the findings in this study are exploratory. While the findings will point to possible patterns and future research directions, the findings are specific to a restricted domain. Thus, results may differ by cybersecurity domain, regulation, or fictitious scenario to which participants must adhere.

Lack of risk involved in the policy development process is another potential limitation. In a professional setting, lack of adherence to a regulation, for instance, could result in fines, decreased reputation of the professional or the company, or impractical processes that impact business continuity. In a simulated environment, such risk is difficult to replicate.

Finally, choice of analysis method may also be a limitation of the study design. While the exploratory nature of the study lends itself well to qualitative analysis, more tailored metrics of assessing difference in compliance, for instance, may be beneficial. The authors would especially like to discuss the evaluation criteria used to assess alignment between the policies ISO29147 with the USEC community.

A. Feasibility of Outcomes and Confounding Factors

This study does not assume that students will under-perform compared to professionals, but instead aims to explore differences in their approaches, with the expectation that students might bring unique perspectives based on their coursework. Comparisons will focus on identifying gaps and strengths instead of analyzing performance.

V. Further Discussions with the USEC Community

In addition to discussion of the limitations of the study, the authors would like to refine the evaluation criteria for the Kaspersky ethical principles and discuss additional frameworks or metrics that could complement the study.

Finally, the authors wish to discuss the potential applicability of this activity and the resulting insights of this study to other universities. Comparisons between students and professionals in this study will result in insights applicable to students beyond the authors' university. The authors would

like to discuss possible avenues of dissemination which would reach a broader audience of students.

VI. Conclusion

The proposed study explores opportunities and challenges associated with integrating AI in the development of cybersecurity policy. We propose an activity and analysis of the interplay between cybersecurity students, professionals, and AI in the developing policies to align with industry standards.

While exploratory in nature, this study has the potential to offer valuable insights into how AI tools can be used effectively by professionals with prior experience and students who are learning to develop industry-compliant policies. Findings from this study will also highlight perceptions of students and professionals who have actively engaged in policy development both with and without AI assistance through this activity. Such insights may help future students and professionals in how they construct their policy development tasks or in determining whether or not to integrate AI into their policy development processes.

Future iterations of this project will extend outside of the vulnerability disclosure domain, assessing applicability of the exploratory findings to a broader set of cybersecurity domains, regulations, and standards.

References

Apacible-Bernardo, A., & Fischer, L. (2024). Identifying global privacy laws, relevant dpas.

Borgeaud, A. (2024). Artificial intelligence (ai) in cybersecurity - statistics facts. *Statista*. https://www.statista.com/topics/12001/artificial-intelligence-ai-in-cybersecurity/#statisticChapter

Cai, W., Pasquale, L., Ramkumar, K., McCarthy, J., Nuseibeh, B., & Doherty, G. (n.d.). Human-ai collaboration for sustainable security: Opportunities and challenges. *Unpublished*.

Da Veiga, A. (2016). Comparing the information security culture of employees who had read the information security policy and those who had not: Illustrated through an empirical study. *Information Computer Security*.

Ferrag, M. A., Alwahedi, F., Battah, A., Cherif, B., Mechri, A., & Tihanyi, N. (2024). Generative ai and large language models for cyber security: All insights you need. *arXiv preprint arXiv:2405.12750*.

Fragiadakis, G., Diou, C., Kousiouris, G., & Nikolaidou, M. (2024). Evaluating human-ai collaboration: A review and methodological framework. *arXiv*.

Goanta, C., Aletras, N., Chalkidis, I., Ranchordas, S., & Spanakis, G. (2023). Regulation and nlp (regnlp): Taming large language models. https://arxiv.org/abs/2310.05553

Jawhar, S., Miller, J., & Bitar, Z. (2024). Ai-based cybersecurity policies and procedures. *2024 IEEE 3rd International Conference on AI in Cybersecurity (ICAIC)*.

Kaspersky. (2020). Kaspersky's ethical principles in responsible vulnerability disclosure. https://media.kasperskydaily.com/wp-content/uploads/sites/85/2020/05/18113421/RVD-Ethical-Principles-EN.pdf

Kelly, J., Zafar, S., Heidemann, L., Zacchi, J., Espinoza, D., & Mata, N. (2024). Navigating the eu ai act: A methodological approach to compliance for safety-critical products. *arXiv*.

Korobenko, D., Nikiforova, A., & Sharma, R. (2024). Towards a privacy and security-aware framework for ethical ai: Guiding the development and assessment of ai systems. *arXiv*.

Lucke, J. V., & Sander, F. (2024). A few thoughts on the use of chatgpt, gpt 3.5, gpt-4 and llms in parliaments: Reflecting on the results of experimenting with llms in the parliamentarian context. *Digital Government: Research and Practice*.

McIntosh, T., Liu, T., Susnjak, T., Alavizadeh, H., Ng, A., Nowrozy, R., & Watters, P. (2023). Harnessing gpt-4 for generation of cybersecurity grc policies: A focus on ransomware attack mitigation. *Computers Security*.

Mitrou, L., Janssen, M., & Loukis, E. (2021). Human control and discretion in ai-driven decision-making in government. *14th International Conference on Theory and Practice of Electronic Governance*.

Perry, B., & Uuk, R. (2019). Ai governance and the policymaking process: Key considerations for reducing ai risk. *Big Data and Cognitive Computing*.

Sahoo, P., Singh, A. K., Saha, S., Jain, V., Mondal, S., & Chadha, A. (2024). A systematic survey of prompt engineering in large language models: Techniques and applications. *arXiv*.

Shilton, K., Heidenblad, D., Porter, A., Winter, S., & Kendig, M. (2020). Role-playing computer ethics: Designing and evaluating the privacy by design (pbd) simulation. *Science and Engineering Ethics*.

Smyth, C. (2025). Parlex ai to advise ministers on how policies will be received. *The Times*. https://www.thetimes.com/uk/politics/article/parlex-ai-to-advise-ministers-on-how-policies-will-be-received-99txwlwph

Trad, F., & Chehab, A. (2024). Prompt engineering or fine-tuning? a case study on phishing detection with large language models. *Machine Learning and Knowledge Extraction*.

Veale, M., & Edwards, L. (2018). Clarity, surprises, and further questions in the article 29 working party draft guidance on automated decision-making and profiling. *Computer Law Security Review*.

Wang, M., Liu, Y., Liang, X., Li, S., Huang, Y., Zhang, X., & Shen, S. (2024). Langgpt: Rethinking structured reusable prompt design framework for llms from the programming language. *arXiv*.

Yigit, Y., Buchanan, W. J., Tehrani, M. G., & Maglaras, L. (2024). Review of generative ai methods in cybersecurity. *arXiv*.