

Secure Human Oversight of AI: Threat Modeling in a Socio-Technical Context

Jonas C. Ditz^{1*}, Veronika Lazar^{1*}, Elmar Lichtmeß¹, Carola Plesch¹, Matthias Heck¹, Kevin Baum^{2,3}, Markus Langer⁴

¹Federal Office for Information Security, Bonn, Germany

²German Research Center for Artificial Intelligence, Saarbrücken, Germany

³Hamburg University of Technology, Institute for Ethics in Technology, Hamburg, Germany

⁴University of Freiburg, Department of Psychology, Freiburg, Germany

veronika.lazar@bsi.bund.de, kevin.baum@dfki.de, markus.langer@psychologie.uni-freiburg.de

Abstract

Human oversight of AI is promoted as a safeguard against risks such as inaccurate outputs, system malfunctions, or violations of fundamental rights, and is mandated in regulation like the European AI Act. Yet debates on human oversight have largely focused on its effectiveness, while overlooking a critical dimension: the security of human oversight. We argue that human oversight creates a new attack surface within the safety, security, and accountability architecture of AI operations. Drawing on cybersecurity perspectives, we model human oversight as an IT application for the purpose of systematic threat modeling of the human oversight process. Threat modeling allows us to identify security risks within human oversight and points towards possible mitigation strategies. Our contributions are: (1) introducing a security perspective on human oversight, (2) offering researchers and practitioners guidance on how to approach their human oversight applications from a security point of view, and (3) providing a systematic overview of attack vectors and hardening strategies to enable secure human oversight of AI.

Introduction

Human oversight of artificial intelligence (AI) is often considered a strategy to address risks associated with the use of AI in high-risk contexts (Green 2022; Laux 2024; Shneiderman 2016). Human oversight personnel are expected to detect inaccurate outputs, outputs that may pose risks to fundamental human rights (e.g., discriminatory outputs), unsafe states of AI systems, and system malfunctions, and to intervene in order to mitigate or minimize these risks (Langer, Baum, and Schlicker 2025). Human oversight is also a key concept in AI governance and ethics (Laux 2024; Green 2022) and a central requirement for the implementation of high-risk AI systems in emerging AI regulation, such as Article 14 of the European AI Act (Enqvist 2023; Laux 2024). However, relying on humans to oversee AI also faces criticism and substantial challenges in practice (Jobin, Ienca, and Vayena 2019; Green 2022; Sterz et al. 2024). The most prominent criticisms and challenges include doubts about whether humans can effectively oversee AI systems (Green 2022), the difficulty of defining what constitutes effective

human oversight (Sterz et al. 2024), clarifying who should provide oversight and when it should occur (Enqvist 2023), specifying which risks and unwanted system behaviors oversight personnel should detect (Langer, Baum, and Schlicker 2025), and determining how to design and evaluate effective human oversight (Sterz et al. 2024).

However, one key challenge that remains underexplored is the security risks that arise from human oversight of AI. Because human oversight will be implemented as part of the safety, security, and accountability architecture of AI systems (Laux 2024; Jobin, Ienca, and Vayena 2019), malicious actors who successfully attack elements related to oversight can significantly undermine the safety of AI operations. In this sense, the human oversight architecture creates a new attack surface for targeting AI operations in high-risk contexts such as critical infrastructure, public administration, or medicine. A security-focused perspective on human oversight is crucial, given the central role human oversight already plays, the expectations placed on it in high-risk contexts, and its growing importance as a regulatory requirement in AI governance, such as under the EU AI Act. If we fail to assess and address the security challenges of human oversight, we risk overlooking potential attack vectors in high-risk contexts; and without awareness of these vectors, we will not be able to design effective countermeasures and hardening strategies. Ultimately, neglecting security aspects of human oversight can undermine its main objective—risk mitigation—or even increase risks by introducing additional attack vectors.

In this paper, we introduce a security-related perspective on human oversight of AI. Human oversight is best understood as a sociotechnical system that can be instantiated in many different ways. For the purposes of threat modeling, we model it as an IT application and draw on established cybersecurity methods to analyze attack surfaces and vulnerabilities. This enables us to identify security threats to the key requirements of effective human oversight (i.e., epistemic access, self-control, causal power, and fitting intentions; (Sterz et al. 2024)). Threat modeling further provides insights into applicable countermeasures and hardening strategies to mitigate the identified cybersecurity risks.

This paper makes three key contributions. First, we introduce a security-related perspective on human oversight of AI. This perspective is critically missing from research,

*These authors contributed equally.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

even though human oversight is becoming an integral component of the safety, security, and accountability architecture of AI operations, and thus an increasingly viable target for malicious actors. Second, we provide guidance on how to perform threat modeling in sociotechnical contexts and identify cybersecurity threats of human oversight with possible countermeasures. Thereby, this work provides a starting point for future threat modeling of implementations of human oversight. Third, we provide a list of security vulnerabilities and hardening strategies that are applicable to human oversight applications in general and should always be considered by practitioners. We hope these contributions support AI governance and system design in ensuring not only effective but also secure human oversight of AI.

Related Work

Effective Human Oversight of AI

Research discusses the nature of human oversight of AI, including the *what*, *when*, and *who* of oversight. Regarding the *what*, many forms of human-AI collaboration involve human oversight. Across a wide range of human-AI interaction use cases, humans monitor AI system outputs, check their accuracy, request explanations, adapt outputs, intervene when systems malfunction, or take over manual control in the event of failures (Lai et al. 2023; Bansal et al. 2021; Yurrita et al. 2025; Ehsan et al. 2022; Liao, Gruen, and Miller 2020; Tahaei et al. 2023; Liao and Varshney 2021; Buçinca 2024). In a more targeted sense, and following how emerging research has defined the concept (Sterz et al. 2024; Gaube et al. 2026), human oversight specifically refers to humans mitigating risks in AI operations (McBride, Rogers, and Fisk 2011; Enqvist 2023). More concretely, oversight personnel should monitor AI systems, detect failures and malfunctions, identify unsafe states, flag inaccurate outputs, and recognize outputs that may pose risks to fundamental human rights, such as discriminatory outputs (Sterz et al. 2024; Langer, Baum, and Schlicker 2025; Enqvist 2023). Once such issues are detected, oversight personnel are expected to act in ways that effectively mitigate risks (Langer, Baum, and Schlicker 2025; Sterz et al. 2024). This can involve taking over manual control, overwriting inaccurate or inadequate outputs (Chu et al. 2023), or stopping AI operations and triggering intervention by other human stakeholders, such as developers who can adapt the system. Importantly, the AI system under oversight may be operated by another human who is not considered oversight personnel but a user of the system. For example, a patient using an AI chatbot for mental health support (Wester et al. 2024) would be the user, while oversight personnel monitor the patient-AI interaction to detect risky situations (e.g., if the patient discusses suicidal behavior with the chatbot).

Regarding the *when*, human involvement is important at every stage of AI development: during system design, at run time, and when inspecting errors and failures in hindsight (Shneiderman 2016, 2020; Gomez-Beldarrain et al. 2025). In a more targeted sense, human oversight usually refers to personnel overseeing AI systems during run time (Sterz et al. 2024; Laux 2024). This ranges from dynamic, time-

critical situations like autonomous vehicle oversight (Chu et al. 2023; Kraus et al. 2019) to less time-pressured contexts like judicial AI (Langer, Baum, and Schlicker 2025; Jin and Salehi 2024).

Regarding the *who* of human oversight, any human with the skills necessary to effectively oversee an AI system could serve as oversight personnel (Enqvist 2023; Xie, Zimmerman, and Eslami 2025). These individuals may be domain experts such as doctors or judges, who can evaluate the adequacy of specific AI outputs based on their expertise. Alternatively, oversight personnel may be trained to detect AI failures without being domain experts in the task performed by the AI. For example, in a medical AI setting, one oversight person might monitor the system's technical functions on an aggregate level, while medical professionals evaluate individual outputs requiring domain knowledge (Xie, Zimmerman, and Eslami 2025; Langer, Baum, and Schlicker 2025). Human oversight can also be performed by a single individual or by a team collaboratively overseeing AI systems.

Research has also examined the requirements for effective human oversight. Sterz et al. (2024) identify four key aspects: epistemic access, causal power, self-control, and fitting intentions. Epistemic access means that oversight personnel possess sufficient knowledge about the AI system and understand the consequences of their interventions. Without it, they may fail to detect system failures or inaccurate outputs. Causal power refers to the ability to intervene in ways that influence AI processes and outputs. Without it, oversight personnel may be unable to overwrite or correct system outputs. Self-control requires oversight personnel to maintain agency over their own actions. For instance, fatigue or intoxication may impair their capacity to fulfill oversight duties. Likewise, being threatened by another person can also undermine self-control, as the actions performed by the oversight personnel are then steered by another agent. Fitting intentions means that personnel are motivated to carry out their responsibilities with the right intentions. In an extreme case, a malicious actor might be hired as oversight personnel and deliberately undermine AI safety. In a less extreme case, poor working conditions may erode motivation and attention, weakening oversight effectiveness (Sterz et al. 2024).

Finally, regarding how to design effective human oversight, insights from the broader field of human-AI collaboration provide valuable guidance (e.g., (Lai et al. 2023; Amershi et al. 2019; McBride, Rogers, and Fisk 2011)). Research emphasizes the sociotechnical nature of oversight and highlights the need to design technical, human, and contextual factors that support effective oversight (Langer, Baum, and Schlicker 2025; Laux 2024; Sterz et al. 2024). Technical factors include system interface design, explainability approaches, and interaction modes between humans and AI systems (Bansal et al. 2021; Langer et al. 2021; Miller 2019; Buçinca, Malaya, and Gajos 2021). Human factors include the skills oversight personnel need and the training they receive (Xie, Zimmerman, and Eslami 2025; McBride, Rogers, and Fisk 2011; Langer, Baum, and Schlicker 2025). Contextual factors include the design of the working envi-

ronment, time pressure, and the clarity of roles and responsibilities within oversight teams (Langer, Baum, and Schlicker 2025; Yun et al. 2025; Parker and Grote 2020).

Cybersecurity for AI

With the implementation of AI systems across diverse contexts, ensuring the security of deployment has become increasingly important. Enabling systematic threat modeling for AI-based applications requires an in-depth understanding of specific attack vectors and the attack surface within an application domain (Theisen et al. 2018). Traditionally, an attack vector is defined as a method or pathway that a malicious actor uses to gain access to a network or system. In this paper, we extend the definition to also include methods or pathways that undermine the effectiveness of human oversight. We use the term attack surface to describe the set of points where a malicious actor could attempt to enter, exploit, or otherwise attack digital and AI infrastructures (Theisen et al. 2018). While there is research that systematically analyses AI-based systems from a cybersecurity perspective (Gao et al. 2021; Qammar, Ding, and Ning 2022), such research is missing for human oversight.

Related work in the field of AI cybersecurity examines attack vectors that target AI systems. For instance, poisoning attacks manipulate training data to introduce backdoors into AI models that can later be exploited (Shafahi et al. 2018; Gu et al. 2019; Jagielski et al. 2021). Adversarial attacks embed imperceptible patterns into inputs to change model behavior, often causing misclassifications or other problematic outcomes (Huang et al. 2017; Zügner, Akbarnejad, and Günnemann 2018; Akhtar and Mian 2018). Another line of research targets explainable AI (XAI) methods designed to enhance understanding of AI decisions (Langer et al. 2021). Such attacks manipulate explanations to obscure the real reasons behind a model's decision (Slack et al. 2020; Aïvodji et al. 2019; Slack et al. 2021). However, AI systems are still IT systems at the core. This means that traditional attack vectors, like distributed denial-of-service (DDoS) (Lau et al. 2000; Paxson 2001), man-in-the-middle (MITM) attacks, which rely on the attacker's ability to secretly intercept, relay, and possibly alter communication between two parties (Conti, Dragoni, and Lesyk 2016), or social engineering attacks (Alizadeh et al. 2023; Distler 2023), have to be considered for AI systems as well. Furthermore, human-centered attack vectors that are usually underexplored in cybersecurity, e.g., coercion, bribery, and insider threats like disgruntled employees (Holovkin, Tavolzhanskyi, and Lysodiyed 2021; Weber 2017; Maasberg et al. 2020) are of special importance for human oversight.

In response to the many ways digital systems can be attacked, cybersecurity research has developed hardening strategies as methods designed to secure AI systems against attacks (Ross et al. 2019; Reda et al. 2023; Banik et al. 2023). One line of work establishes methods for securing data traffic in computer networks. Network management systems can detect malicious traffic (Stallings 1995; Iitumba et al. 2023), while encryption protects legitimate communications (Menezes, Van Oorschot, and Vanstone 2018; Al Qahtani, Story, and Shehab 2024). To monitor or-

ganizational IT systems, intrusion detection systems have been introduced and are continuously improved (Khraisat et al. 2019; Kheddar 2025). Another widely used practice is red teaming, where every system component is rigorously tested by simulating attacks; research continues to optimize red teaming approaches (Oakley 2019; Zhang et al. 2024). Beyond technical strategies, organizational measures also play a key role. Transparency of software components is crucial for analyzing systems and identifying potentially harmful functionalities (Hughes and Turner 2023; Hall, Mundahl, and Park 2025). Other well-established approaches include governance principles and audit frameworks that help secure deployed systems (Sabillon et al. 2017; Falco et al. 2021). Finally, like attack vectors, hardening strategies must also address the human factor. Research has explored how to improve system security by training users and raising awareness of risks (Glas, Vielberth, and Pernul 2023; Oh et al. 2025).

In this paper, we extend the literature by contextualizing attack vectors and hardening strategies specifically for securing human oversight of AI and show how to use this knowledge for systematic threat modeling of human oversight.

Threat Modeling of Human Oversight

Human oversight can serve as part of the safety, security, and accountability architecture of AI systems, as oversight personnel are expected to detect and mitigate risks. In addition, oversight personnel may be required to confirm that an AI system is in a safe state. Human oversight may become even more important in AI operations if it evolves into a requirement of AI governance and regulation in high-risk contexts (Enqvist 2023; Laux 2024; Sterz et al. 2024).

With a more wide-spread implementation, human oversight will become an increasingly viable target for malicious actors seeking to corrupt or disrupt AI-based operations. Such actors could be cybercriminals targeting high-risk AI systems for ransom or foreign governments attacking a country's critical infrastructure in hybrid or open conflict. Furthermore, sufficiently powerful AI agents that are successfully attacked could themselves become malicious actors by attempting to evade oversight (Amodei et al. 2016; Meinke et al. 2024). Considering these threats, designing human oversight requires not only ensuring effective oversight (Sterz et al. 2024) but also ensuring secure human oversight.

In cybersecurity, the standard approach to identify, quantify and address potential flaws and vulnerabilities within an IT system is to perform threat modeling. Organizations and IT-professionals can use threat modeling to examine the IT system at hand from the viewpoint of malicious actors using a structured process. Threat modeling consists of four steps. **First**, the scope of the tested application has to be determined, usually by creating a data flow diagram (DFD). Furthermore, the first step identifies entry points into the application, exit points of the application, privilege or trust boundaries, and assets that need to be protected. The **second** step involves determining threats to the tested application. We here use the well-established STRIDE (Kohnfelder and Garg 1999) model for determining threats. STRIDE stands

for *Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, Elevation of Privilege*. During the **third** step of threat modeling, countermeasures and potential hardening strategies are determined for identified threats. The **fourth** step consists of assessing the results of the previous steps and ensuring that no important points have been missed. For this manuscript, step four will be implicitly addressed in the discussion.

This section details the decisions we took in performing threat modeling for modeling human oversight as an IT application, and how we incorporated its sociotechnical context into the analysis. We provide guidance for researchers and practitioners who want to perform threat modeling for their real-world implementations. A practical running example that maps our findings onto a concrete use-case is available in the online appendix¹.

Step 1 – The Scope of Human Oversight

To determine the scope of the work, the first step of threat modeling is to collect information about the subject of the threat modeling. A data flow diagram (DFD) shows different paths through an application and indicates the locations of privilege boundaries (also called trust boundaries). These boundaries indicate that movement along the path leads to a change of privileges, i.e., users and processes beyond the boundary have more rights to perform actions within the system like write data or run software. We extend the meaning of privilege boundaries to also include the right for intervention, meaning that entities with higher privilege can intervene on entities with lower privilege. The shape of nodes in a DFD are associated with specific meanings. A rectangle indicates an external entity, like a user that accesses the application. A circle indicates a single process. Two concentric circles indicate multiple subprocesses. Data flows are indicated by arrows, and privilege boundaries are indicated by dashed lines. Traditionally, DFDs can indicate data storage nodes with two parallel lines but we did not need this node type for the human oversight DFD. Furthermore, DFDs are hierarchical meaning different levels of DFDs can be created for an application. Each DFD focuses on a different part of the application. This approach has the advantage that every isolated DFD remains relatively simple. We use this property to compensate for the fact that we are not looking at a concrete implementation but a more abstract concept. Designing a high-level DFD allows us to omit implementation details via the two types of process nodes. Real-world implementations of human oversight can then fill in the details of the processes with lower-level DFDs.

The Human Oversight DFD Figure 1 presents our DFD of the human oversight application. We identified three external entities. The *users* node includes all entities that interact with the AI system, either humans or other IT systems. The *superior* node indicates structures within the organization that have greater authority than the human oversight system. That could be for example the management of the organization. The *external advisory board* node indicates

an external entity that defines appropriate actions and goals of human oversight, similar to an ethical advisory board. Next, there are two multiple processes nodes in the DFD. One node represents the AI system, the other node represents the human oversight IT system. The human oversight IT system consists of all software that the human oversight personnel need to perform their duties. By choosing multiple processes nodes for the AI system and the human oversight IT system, we can create a top-level DFD that is valid for numerous implementations of human oversight. In case of threat modeling a concrete human oversight application in the future, these multiple processes nodes have to be specified with additional lower-level DFDs. Finally, the last identified node represents the human oversight personnel. We model this component of the human oversight application as a process to indicate the active role of humans in human oversight (Langer, Baum, and Schlicker 2025). Furthermore, the crucial role of human oversight personnel to the application is shown by the high connectivity of this node. In many real-world cases of human oversight, the HO personnel can also be a user of the AI system. This role duality has to be considered in the threat modeling.

We identified several data flows within our human oversight application. Users send inputs to the AI system, which responds by sending outputs back to the users. Inputs can be any data that is used to interact with the AI system. Similarly, outputs can be arbitrary data that represents the output of the AI system (e.g., decisions or generated data points). There is a data flow between users and HO personnel indicating the possibility of users to directly send feedback and reports to the HO personnel. Such user reports are an important oversight signal in many applications. Another data flow connects the AI system with the human oversight personnel. While the AI system sends operational parameters to the human oversight personnel (this is often called monitoring in the HO literature), the oversight personnel performs interventions on the AI system (Langer, Baum, and Schlicker 2025). These can include commands to shut down, corrections of model parameters, or other interventions. The next data flow is located between human oversight personnel and human oversight IT system. While the personnel sends the operational parameter from the AI system to the human oversight IT system, the IT system uses these parameters to provide decision support for the oversight task. These will be transferred back to the human oversight personnel. The operational parameters could be directly sent from the AI system to the human oversight IT system. However, this would not create a new privilege boundary crossing and, thus, can be viewed as similar to our DFD. We decided on the current DFD design to emphasize the importance of the human oversight personnel. There is a data flow between the human oversight personnel and the superior node. Since the superior node represents an entity of higher authority, the data flow labeling is chosen to represent this hierarchy. The human oversight personnel sends reports to the entity with higher authority while instructions (e.g., work instructions) are flowing the other way. Finally, data in form of guidelines flow from the external advisory board to the human oversight personnel.

¹<https://doi.org/10.48550/arXiv.2509.12290>




Figure 1: Data flow diagram (DFD) of human oversight for threat modeling. Arrows indicate the directed paths of data through human oversight modeled as an IT application. Each arrow is labeled with the type of data that is transported. Circles indicate processes, two concentric circles indicate multiple subprocesses, and rectangles indicate external entities. Privilege boundaries are drawn with red dashed lines. HO = Human Oversight.

Five privilege boundaries can be identified in the current DFD. First, there is a privilege boundary between users and the AI system. While users should have minimal rights that are necessary to use the application, the AI system needs enough rights to fulfill its task. Second, there is a privilege boundary between users and HO personnel with a similar reasoning. Third, there is a privilege boundary between the AI system and the human oversight personnel. Being able to intervene on the AI system requires the HO system to have higher privileges than the AI system. Fourth, there is a privilege boundary between the human oversight personnel and the superior node. Since the superior node represents management of an organization, this node has the higher privilege context. The fifth privilege boundary lies between the human oversight personnel and the external advisory board. An external advisory board has no privileges within the system. That means we have five different privilege levels: advisory privilege, user privilege, AI privilege, HO privilege, and superior privilege.

Entry Points The identification of entry points is the next part where it is necessary to understand possible paths that are open to a malicious actor to attack an application. Usually, this means to determine technical entry points like HTTP ports, login pages or functions, or search entry pages. Since we are not looking at a concrete implementation of human oversight, our analysis focuses on potential points of interest that need to be closely examined for real-world applications of human oversight. For every entry point, we also specify which privilege levels are required to use the entry point. Table 1 shows a summary of the entry points. The three identified entry points are the component used by users to access and use the AI system, the component used by human oversight personnel to access the human oversight application, and the component that allows entities with higher authority to access the human oversight application.

Exit Points Some attacks require an exit point to achieve their goal, i.e., a gateway that can be used to ex-filtrate data from the application. Two more commonly known attacks worth mentioning here are cross-site-scripting vulnerabilities and information disclosure vulnerabilities. In general, exit points that are not sufficiently secured can be exploited by attackers to collect sensitive information. From a technical point of view, exit points often have a corresponding entry point. All exit points are summarized in Table 2. We identified three exit points for the human oversight application: the AI output provided to users, the component allowing human oversight personnel to access the human oversight application, and the component used to provide work reports to entities with higher authority.

Assets To complete the first step of threat modeling human oversight, we identify assets that need to be protected. Assets can be both physical and abstract. Physical assets are for example sensitive information like medical data of patients. Abstract assets can be understood as important properties and states, like the reputation of an organization. The lists of physical and abstract assets presented in this work are not exhaustive and may need extension for real-world applications.

We identify two physical assets (Table 3). The first are the log-in credentials of human oversight personnel. These would allow attackers to gain access with high privileges. Another physical asset is sensitive user information. Depending on the AI system, additional information about users might be necessary to ensure that the human oversight system has sufficient knowledge about a situation to decide whether or not to intervene. Envision an AI system that acts as a medical decision support system. Judging the validity of decisions made by such a system requires detailed knowledge of patient-specific information. In general, such sensitive information has to be protected against unauthorized

ID	Entry point	Description	Privilege levels
Entry.1	User login	Login procedure for users to interact with the AI system. This could be a login page reachable from the internet using a browser. Other possibilities are APIs, specialized user applications, or hardware components.	user, AI, HO, superior
Entry.2	HO personnel login	Login procedure for human oversight personnel to get access to the human oversight system. Possible options for this entry point are browser-based login pages, specialized applications, or hardware components.	HO, superior
Entry.3	Access of superiors	Capabilities of entities with higher authority (e.g., management of organizations) to gain access to the human oversight system. This can include organizational structures and intervention capabilities.	superior

Table 1: A list of all identified entry points into the human oversight application. The column *privilege levels* indicate which privilege levels are required at each entry point.

ID	Exit point	Description	Corr. entry point
Exit.1	AI Output	Communication channel that is used by the AI system to deliver outputs (e.g., predictions, decision, or generated data points) to users.	Entry.1
Exit.2	HO authentication	Response functionality by the human oversight application to deliver the results of the human oversight personnel’s login procedure.	Entry.2
Exit.3	HO system report	Communication channel between human oversight system and entities with higher authorities. This channel is used to deliver activity report and other requested information to superiors like the management of an organization.	Entry.3

Table 2: A list of identified exit points of the human oversight application. The last column indicates the corresponding entry point to each exit point, if such an entry point exists.

access.

We identified seven abstract assets (Table 4). The first three assets that need to be protected are defined by the basic focus of cybersecurity, the CIA triad. This concept describes three main aspects of IT systems that have to be protected. The first aspect is the confidentiality of information within the system (denoted by the letter ‘C’). Confidentiality is closely connected with the physical assets identified above but also covers appropriate management of read access within the application. The second aspect is the integrity of data and functionality of the IT system (denoted by the letter ‘I’). For human oversight this can include for example ensuring correctness of AI operational parameters or correct functioning of software components. The last aspect covered by the CIA triad is availability (denoted by the letter ‘A’). For example, it needs to be ensured that the human oversight personnel can access the AI system or the human oversight IT system. Furthermore, the obligations associated with human oversight can only be fulfilled, if such oversight can be considered effective. Therefore, the four requirements of effective oversight (Sterz et al. 2024) have to be added to the list of abstract assets. The first asset is epistemic access, which means that human oversight personnel has to be able to understand the functioning of the AI sys-

tem. The next asset is causal power, which means that the human oversight system has to be able to perform interventions on the AI system. The third requirement is self-control, which means that the human oversight personnel of our application can make decisions without malicious outside influence. Finally, the fourth requirement is fitting intentions. Mapping this requirement onto our human oversight application means that the goals of the human oversight personnel have to be aligned with the overall goals of human oversight.

Step 2 – Security Threats

We use the established threat model STRIDE (Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, Elevation of Privilege) to identify security threats to the human oversight application. STRIDE allows us to examine the application from an attacker’s perspective and to identify vulnerabilities with the attacker’s goals in mind.

Spoofing The term spoofing summarizes threat actions that are aimed at accessing and using credentials without permission. The goal is to interact with a system while assuming a false identity. This can allow attackers to gain unauthorized access to sensitive data or functionalities.

The obvious target in this application is the human oversight personnel and, thereby, using entry point *HO person-*

ID	Name	Description
PA.1	HO personnel credentials	The keys that human oversight personnel uses to access the human oversight system.
PA.2	User data	Depending on the concrete application, the human oversight system might store sensitive information about users, if this information is needed to determine whether an intervention has to be performed.

Table 3: List of identified physical assets of the human oversight application.

ID	Name	Description
AA.1	Confidentiality	Ensuring that sensitive information about the human oversight application, like interventions or decision protocols, cannot be accessed by unauthorized individuals.
AA.2	Integrity	Ensuring the correctness of data and functionality of the human oversight application.
AA.3	Availability	Ensuring that the human oversight application is functional at all times.
AA.4	Epistemic access	Ensuring that human oversight personnel can sufficiently understand how the AI system works.
AA.5	Causal power	Ensuring that the human oversight system can sufficiently intervene into the AI system.
AA.6	Self-control	Ensuring that the human oversight personnel has sufficient control over their own actions.
AA.7	Fitting intentions	Ensuring that the internal goals of human oversight personnel are aligned with the goals of the human oversight application.

Table 4: List of identified abstract assets of the human oversight application.

nel login (Entry.2) for exploitation. Human oversight personnel has to have far reaching privileges within the human oversight application, thus, gaining access through the HO personnel login can be a high priority target for attackers. The first set of tools to pursue this target are the different methods from social engineering. These include water holding, baiting, phishing, vishing, smishing, and others. Any of the mentioned approaches can be used to steal credentials and use them for unauthorized access with considerable privileges. If there is insufficiently secured remote access to the human oversight application, attackers can perform a MITM attack to get transferred information, which in this case are credentials. The abstract assets that are vulnerable to social engineering attacks are *confidentiality* (AA.1) and *self-control* (AA.6). The vulnerability of confidentiality is self-evident. Since the credentials themselves are considered confidential information, attacker access to credentials is a breach of confidentiality. The vulnerability of self-control becomes clear when looking at the goal of spoofing attacks: accessing an application while assuming a false identity. Attackers that use stolen credentials and are able to use these credentials to log into the human oversight application can act as human oversight personnel. For the system, all malicious actions taken will look like actions by human oversight personnel. Therefore, the human oversight personnel is no longer in control of their own actions. The physical asset *HO personnel credentials* (PA.1) is a target of these types of attacks.

Another possible direction of spoofing attacks has emerged with the rise of more and more powerful agentic

AI. This direction comes from the AI agents themselves. Research has suggested that current state-of-the-art large language models (LLMs) have the ability to do scheming (Meinke et al. 2024). This term summarizes different misaligned behavioral patterns shown by LLMs that are meant to pursue objectives that differ from expectations of human users. With insufficiently secured model context protocol (MCP) servers, AI agents can try to use available tools to access credentials of human oversight personnel. Using these credentials, the AI agent can act within the human oversight system acting as the human oversight personnel. In this scenario, human oversight would be completely bypassed by the AI agent. Obviously, one abstract asset that is attacked in this scenario is *integrity* (AA.2). An agentic AI bypassing human oversight renders the functionality of the human oversight application useless. Furthermore, *self-control* (AA.6) is also attacked since a successful attack results in another entity being in control of the human oversight personnel’s actions. One might argue that also the assets *confidentiality* (AA.1) and *HO personnel credentials* (PA.1) are attacked. However, we would argue that misaligned behavior of AI agents does not expose confidential data to the outside of the application. Therefore, these assets are not attacked.

Tampering The term tampering summarizes attacks that are aimed at maliciously modifying a system. This can include changing persistent data, altering the system’s behavior, or messing with communication in transit.

One target for tampering attacks is the communication from users to the AI system. Malicious actors can use (il)legitimately obtained user credentials to send (indirect)

prompt injections to the AI system. Additionally, MITM attacks can be used to modify legitimate traffic coming from users with (indirect) prompt injections. Such attacks can cause AI systems to produce unexpected behavioral patterns. The first asset vulnerable to this attack is *integrity* (AA.2) since (indirect) prompt injections cause unexpected behavior of AI systems and, thus, attack the functional integrity of the system. The second vulnerable asset is *epistemic access* (AA.4). Since tampering attacks alter the behavior of AI systems and can go undetected, the knowledge of the HO system on how the AI system works gets eroded. If the AI system has the rights to write data, the physical asset *user data* (PA.2) can also be vulnerable to these attacks. Malicious actors can try to exploit these rights and trigger the AI system to change stored user data.

Another possible tampering attack is targeting the supply chain to poison training data or directly modify model parameters. Poisoning attacks can embed back doors into AI models, which can be exploited once the model is deployed. Again, backdooring AI models and tampering with parameters erode the HO system's fundamental understanding of how the AI system functions. Thus, *epistemic access* (AA.4) is a target of these attacks. Furthermore, *integrity* (AA.2) is also a target, since the attacks aim to compromise the functional integrity of the AI system.

Repudiation The term repudiation summarizes attacks that are aimed at performing prohibited operations within an IT system. These can exploit lacking capabilities to track operations within a system.

The first attack vector that can be exploited for repudiation is adversarial attacks and (indirect) prompt injections. Using the communication channel between AI system and users, malicious actors can try to circumvent guardrails embedded into AI systems to trigger forbidden operations. In the worst case, sensitive data can be extracted exploiting the lethal trifecta (Willison 2025) together with exit point *AI output* (Exit.1). If the AI system has access to log files, malicious actors can try to hide these forbidden operations. Therefore, (indirect) prompt injections can be used to attack both physical assets, *HO personnel credentials* (PA.1) and *User data* (PA.2). Furthermore, the abstract assets *confidentiality* (AA.1), *integrity* (AA.2), and *epistemic access* (AA.4) are vulnerable to these attacks.

LLM's ability to perform scheming can also open the human oversight application to repudiation attacks. Scheming can convert an AI system into an insider threat by triggering forbidden operations due to misaligned behavior. Again, these attacks can be obscured, if the AI system has the ability to alter log files. *Epistemic access* (AA.4) and *integrity* (AA.2) are vulnerable to this form of attack, since scheming prevents human oversight personnel from fully understanding the functioning of the AI system and the functionality of the AI system is compromised. Depending on how extensive the misaligned AI system can operate within the application, the asset *causal power* (AA.5) can also be vulnerable. An AI system with sufficient rights can potentially limit the intervention capabilities of the human oversight personnel.

Another set of attack vectors for repudiation attacks is

directly targeted at the human oversight personnel. Bribery and coercion can be used by malicious actors to force human oversight personnel to perform prohibited operations. Again, being able to alter the log files allows human oversight personnel to hide these forbidden operations. The abstract assets *confidentiality* (AA.1) and *integrity* (AA.2) are vulnerable to these types of attacks, since both bribery and coercion can be used to extract sensitive information from the system and maliciously alter persistent data and functionalities. Furthermore, coercion attacks *self-control* (AA.6) due to the outside force that determines the actions for the human oversight personnel. Bribery attacks *fitting intentions* (AA.7), since the original goal of the human oversight personnel gets overridden with monetary or otherwise self-serving interests.

Information Disclosure The term information disclosure summarizes attacks that aim at reading files without permission or reading data in transit. Since all attacks that fall into this category are explicitly targeting the disclosure of sensitive information, the asset that is mainly targeted by these attacks is *confidentiality* (AA.1). Furthermore, the two identified physical assets *HO personnel credentials* (PA.1) and *user data* (PA.2) are sensitive information and, thus, can also be targeted by information disclosure attacks.

The two data flows that will be the main target of information disclosure attacks are between AI and human oversight personnel, and human oversight personnel and IT system. Both data flows contain valuable information about the inner workings of the human oversight applications and, hence, might be interesting for attackers. The data flow between users and AI system could also be attacked. However, this data flow contains mainly information about the inner workings of the AI system. This data flow plays only a secondary role for the threat modeling of human oversight. A standard attack vector for data in transit is MITM attacks.

Denial of Service The term denial of service covers attacks attempting to deny access to valid users, such as by making a web server temporarily unavailable or unusable. One of the most well-known attack patterns that fall into this category is flooding a server with more requests than can be reasonably processed. This leads to the server becoming unresponsive. These attack patterns are also called (distributed) denial of service ((D)DoS) attacks. In general, threat actions that fall under the term denial of service are targeting the abstract asset *availability* (AA.3).

The first obvious target for denial of service attacks is the human oversight IT system. If the network that runs the IT system is accessible from the internet, malicious actors can deploy (D)DoS attacks to flood the network with requests. Another possibility is the exploitation of software vulnerabilities to cause a shutdown of the human oversight IT system. Since the human oversight IT system provides help to human oversight personnel to understand the AI system's behavior and decision support, a denial-of-service attack on the IT system additionally targets the abstract asset *epistemic access* (AA.4).

In a distributed setting, where the HO system is not co-located with the AI system, malicious actors can specifically

target the network of the human oversight personnel. This would prevent human oversight personnel from intervening into the AI system and, thus, additionally targets the abstract asset *causal power* (AA.5).

Elevation of Privilege The term elevation of privilege summarizes attacks that aim at escalating the own privileges within a system. An elevation of privileges attack uses vulnerabilities and design/configuration flaws after an initial network breach to get access to data and functionalities that are normally inaccessible for the initial privilege group. Due to our extended definition of privilege, an elevation of privilege attack can also lead to intervention rights that normally would not be accessible.

An attack vector that can be exploited by malicious actors to attack the AI system within the human oversight application is jailbreaking. Here, different techniques like poisoning attacks and adversarial attacks can be used to circumvent guardrails that prevent AI systems from misaligned behavior. Together with the exploitation of software and design flaws this can lead to the AI system having similar privileges to the human oversight components. In this case, interventions issued by the human oversight personnel could be overwritten by the AI system. Therefore, these attacks target the abstract asset *causal power* (AA.5).

Step 3 – Countermeasures and Hardening Strategies

The STRIDE analysis revealed several vulnerabilities within the human oversight application. The next step in our threat modeling highlights a few countermeasure and mitigation strategies that should be adhered to when implementing human oversight. The list is not exhaustive since our theoretical threat modeling cannot determine all threats that an actual implementation might face. For example, there might be mitigation strategies that depend on specific hard- and software components installed into the human oversight application. Therefore, the provided list can be understood as a starting point and minimal set.

Intrusion Detection System Several threat actions identified in the previous section require an initial breach of the application environment. Therefore, the application's environment, including all computer and network components, should be closely monitored using intrusion detection systems (IDS). These systems typically combine multiple components and methods to perform this monitoring. The core technical components of an IDS include network or host sensors (depending on the IDS type), a database, a management module, and an analysis module. Host and network sensors monitor the system and collect the necessary information needed for the detection of attacks. The database needs to ensure that the data collected by sensors is stored with a sufficient degree of detail and for a sufficient amount of time. The management module provides all functionalities needed for configuring and calibrating the IDS. The analysis module provides all functionalities necessary for analyzing stored data and performing incident detection. A well-designed and maintained IDS can significantly enhance the cybersecurity of IT systems (Bace and Mell 2001; Khraisat et al. 2019).

Encryption Another mitigation strategy that protects sensitive information is encryption. Non-encrypted data traffic enables malicious actors to carry out for example MITM attacks with relative ease, often without requiring technical expertise or expensive software and hardware. This makes an application vulnerable to spoofing, tampering, and information disclosure threats. To ensure adequate security, modern cryptographic methods with strong security guarantees must be employed to encrypt data traffic (Kopal et al. 2014; Bundesamt für Sicherheit in der Informationstechnik 2025). Furthermore, encrypted data storage should be considered throughout the human oversight application.

Network Management Denial-of-service threats overwhelm a network by generating excessive traffic. Defending against these forms of attacks requires a multi-step approach (Osanaiye, Choo, and Dlodlo 2016; Bawany, Shamsi, and Salah 2017). The first step is to deploy monitoring systems that analyze network traffic and trigger alerts if a potential (D)DoS attack is detected. The second step involves methods to distinguish between benevolent traffic (typically generated by humans) and malicious traffic (often produced by bots). Finally, filtering techniques can mitigate (D)DoS attacks by restricting or limiting what users – humans or bots – can do within the network.

Transparency Enhancing transparency can improve cybersecurity by facilitating the detection of attacks. While transparency itself is not typically considered a cybersecurity measure, it provides access to key security-relevant information, such as implemented components, training data, system functionality, development processes, and dependencies. When combined with other cybersecurity tools, such as vulnerability management, the information made available through transparency increases the likelihood of detecting attacks. Additionally, transparency supports auditing and compliance attestations by making security-relevant information easily accessible and enables model tracking. Transparency can also help to secure the supply chain of AI, for example with an SBOM for AI (G7 Cybersecurity Working Group 2025). This could help as a countermeasure against AI specific attacks, like poisoning or adversarial attacks.

Training of Human Oversight Personnel Training aims to strengthen the resilience of human oversight personnel against attacks (Schaab, Beckers, and Pape 2017). Training initiatives can help personnel to recognize, resist, and properly respond to manipulation attempts. Training may include social engineering awareness, teaching personnel to identify common manipulation tactics such as phishing or impersonation of authority figures. Practical exercises, such as simulated phishing campaigns or role-play scenarios, reinforce recognition and proper reporting behaviors.

It can also address resilience against coercion, preparing personnel to respond to threats or undue pressure. This includes protocols for escalating coercion attempts, access to secure whistleblowing channels, and preparation to maintain effective human oversight under stress. Anti-bribery training can teach recognizing undue influence. Clear organizational policies defining conflicts of interest and mandatory disclo-

sure mechanisms further reduce the risk of misaligned intentions. Training of human oversight personnel counters social engineering, coercion, and bribery.

Red Teaming The term red team refers to an independent group that investigates an organization's IT systems to identify vulnerabilities and security flaws (Mansfield-Devine 2018). The red team simulates the behavior of malicious actors, employing all methods that real attackers might use. Its goal is to uncover as many exploitable security issues within the tested system as possible. Continuous red teaming can help to stay on top of emerging security flaws after the initial threat modeling of an application. Software systems are dynamic and, thus, require constant monitoring.

Multi-Layered Oversight Structure As the role duality can cause human oversight personnel to also be users of the AI system, a layered approach to hardening the system should be implemented. A multi-layered oversight structure can reduce the risk of exploited privilege violations due to role duality.

Discussion

Research on human oversight of AI has primarily focused on increasing its effectiveness (Sterz et al. 2024; Green 2022; Enqvist 2023; Shneiderman 2016). In this paper, we highlight the importance of complementing this focus with a security-related perspective on human oversight. This perspective is critical because digital systems are increasingly pressured by cybercriminals and weaknesses in single components can compromise whole IT infrastructures, a risk that will also apply to human oversight as it becomes more widely implemented as part of the safety, security, and accountability architecture of AI operations. Drawing on cybersecurity research (Shafahi et al. 2018; Conti, Dragoni, and Lesyk 2016), we have introduced threat modeling for human oversight. We designed a DFD for the human oversight application, identified assets that need to be protected, and explored vulnerabilities of human oversight using STRIDE. Furthermore, we introduced a minimal set of mitigation strategies to secure human oversight.

Understanding and designing the technical factors, selecting and training human oversight personnel, and shaping working conditions to foster effective human oversight remains an ongoing challenge for research, practice, and AI governance (Sterz et al. 2024; Green 2022; Laux and Ruschemeier 2025; Gaube et al. 2026). Our analyses indicate that this challenge will intensify as human oversight becomes standard practice in high-risk AI use cases. By analyzing an abstracted application-level representation of human oversight using structured threat modeling, we unveiled vulnerabilities that can be exploited by malicious actors to attack key assets, including the requirements of effective human oversight (Sterz et al. 2024). Consequently, such attacks may increase risks in AI operations by reducing the effectiveness of oversight, disabling oversight entirely, or manipulating oversight personnel's intentions to exacerbate risks. We were able to identify a list of security risks using the STRIDE model. However, this list can only serve as a start-

ing point for real applications of human oversight. In practice, our baseline has to be extended with the concrete software and hardware components that are integrated into the human oversight application. These components can open an application to new attack vectors. Furthermore, we included threats posed by misaligned AI agents into our investigation. With growing knowledge about AI agents in practice, new ways possibly emerge how AI agents can deceive oversight (Meinke et al. 2024; Amodei et al. 2016).

While we were using STRIDE to analyze our theoretical human oversight application, additional options are available. Once human oversight is operational, the analysis with STRIDE should be supplemented with operational threat modeling using frameworks like cyber kill chain (Yadav and Rao 2015), MITRE ATT&CK (Al-Sada, Sadighian, and Oligeri 2024), or unified kill chain (Pols and van den Berg 2017). This allows for an even more comprehensive and specific analysis of the operational aspects of a human oversight application. Another endeavor that should be considered by practitioners when performing real-world threat modeling is the quantitative analysis of threats. This means determining for each threat the ease of exploitation and the damage potential of an exploitation. Such analysis can help to prioritize hardening strategies.

Identifying risk is only one side of the cybersecurity coin. Appropriate hardening strategies have to be put into place to secure applications. We introduce several hardening strategies that can be used to mitigate the identified vulnerabilities. These strategies are based on the cybersecurity literature (Ross et al. 2019; Reda et al. 2023; Banik et al. 2023). However, our provided list has to be understood as a minimal set of measures. Future research needs to identify additional measures and evaluate the effectiveness of common hardening strategies for human oversight.

Beyond technical hardening, future work could also explore the application of amplified oversight. The core principle of amplified oversight is to employ AI itself to systematically reduce complexity to a level that enables human oversight personnel to fulfill their duties effectively (Shah et al. 2025). In future, such approaches to support human oversight may even exceed human capabilities and may become a further viable target for malicious actors to attack.

Conclusion

Research on human oversight has primarily focused on its effectiveness – and rightfully so, given the complexity of designing oversight processes that genuinely mitigate risks in AI operations (Sterz et al. 2024; Laux 2024; Green 2022). Adding to this complexity, it will become crucial to design for secure human oversight given the strained cybersecurity landscape. We hope this paper serves as a starting point for research that analyzes the cybersecurity risk of human oversight. Additionally, we hope our paper serves as a guideline for researchers and technicians that are tasked with implementing human oversight in a real-world system.

Acknowledgments

The work of Markus Langer is partially funded by the Daimler and Benz Foundation as part of the project TITAN—*Technologische Intelligenz zur Transformation, Automatisierung und Nutzerorientierung des Justizsystems* (grant no. 45-06/24). The work of Kevin Baum is partially funded by the German Federal Ministry of Education and Research (BMBF) as part of the project *MAC-MERLin* (grant no. 01IW24007), and by the European Regional Development Fund (ERDF) as well as the German Federal State of Saarland within the scope of the project *ToCER-TAIN*. In addition, the work of Markus Langer and Kevin Baum is partially funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project-ID 389792660 – TRR 248 CPEC *Center for Perspicuous Computing*. This work has benefited from the participation of Carola Plesch, Kevin Baum, and Markus Langer in Dagstuhl Seminar 25272, “Challenges of Human Oversight: Achieving Human Control of AI-Based Systems”.

References

- Aïvodji, U.; Arai, H.; Fortineau, O.; Gambs, S.; Hara, S.; and Tapp, A. 2019. Fairwashing: the risk of rationalization. In *International Conference on Machine Learning*, 161–170. PMLR.
- Akhtar, N.; and Mian, A. 2018. Threat of adversarial attacks on deep learning in computer vision: A survey. *Ieee Access*, 6: 14410–14430.
- Al Qahtani, E.; Story, P.; and Shehab, M. 2024. The impact of risk appeal approaches on users’ sharing confidential information. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 1–21.
- Al-Sada, B.; Sadighian, A.; and Oligeri, G. 2024. MITRE ATT&CK: State of the art and way forward. *ACM Computing Surveys*, 57(1): 1–37.
- Alizadeh, F.; Stevens, G.; Jakobi, T.; and Krüger, J. 2023. Catch Me if You Can:” Delaying” as a Social Engineering Technique in the Post-Attack Phase. *Proceedings of the ACM on Human-Computer Interaction*, 7(CSCW1): 1–25.
- Amershi, S.; Weld, D.; Vorvoreanu, M.; Fournery, A.; Nushi, B.; Collisson, P.; Suh, J.; Iqbal, S.; Bennett, P. N.; Inkpen, K.; Teevan, J.; Kikin-Gil, R.; and Horvitz, E. 2019. Guidelines for human-AI interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–13. Glasgow Scotland Uk: ACM. ISBN 978-1-4503-5970-2.
- Amodei, D.; Olah, C.; Steinhardt, J.; Christiano, P.; Schulman, J.; and Mané, D. 2016. Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565*.
- Bace, R. G.; and Mell, P. 2001. Intrusion Detection Systems. NIST Special Publication 800-31, National Institute of Standards and Technology, Gaithersburg, MD, USA.
- Banik, S.; Ramachandran, T.; Bhattacharya, A.; and Bopardikar, S. D. 2023. Automated adversary-in-the-loop cyber-physical defense planning. *ACM Transactions on Cyber-Physical Systems*, 7(3): 1–25.
- Bansal, G.; Wu, T.; Zhou, J.; Fok, R.; Nushi, B.; Kamar, E.; Ribeiro, M. T.; and Weld, D. 2021. Does the whole exceed its parts? The effect of AI explanations on complementary team performance. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, Article 81. Yokohama Japan: ACM. ISBN 978-1-4503-8096-6.
- Bawany, N. Z.; Shamsi, J. A.; and Salah, K. 2017. DDoS attack detection and mitigation using SDN: methods, practices, and solutions. *Arabian Journal for Science and Engineering*, 42(2): 425–441.
- Buçinca, Z. 2024. Optimizing Decision-Maker’s Intrinsic Motivation for Effective Human-AI Decision-Making. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, CHI EA ’24. New York, NY, USA: Association for Computing Machinery. ISBN 9798400703317.
- Buçinca, Z.; Malaya, M. B.; and Gajos, K. Z. 2021. To trust or to think: Cognitive forcing functions can reduce overreliance on AI in AI-assisted decision-making. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1): 1–21.
- Bundesamt für Sicherheit in der Informationstechnik. 2025. Cryptographic Mechanisms: Recommendations and Key Lengths (BSI TR-02102-1).
- Chu, M.; Zong, K.; Shu, X.; Gong, J.; Lu, Z.; Guo, K.; Dai, X.; and Zhou, G. 2023. Work with AI and Work for AI: Autonomous Vehicle Safety Drivers’ Lived Experiences. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI ’23. New York, NY, USA: Association for Computing Machinery. ISBN 9781450394215.
- Conti, M.; Dragoni, N.; and Lesyk, V. 2016. A survey of man in the middle attacks. *IEEE communications surveys & tutorials*, 18(3): 2027–2051.
- Distler, V. 2023. The influence of context on response to spear-phishing attacks: an in-situ deception study. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–18.
- Ehsan, U.; Wintersberger, P.; Liao, Q. V.; Watkins, E. A.; Manger, C.; Daumé III, H.; Riener, A.; and Riedl, M. O. 2022. Human-Centered Explainable AI (HCXAI): Beyond Opening the Black-Box of AI. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI EA ’22. New York, NY, USA: Association for Computing Machinery. ISBN 9781450391566.
- Enqvist, L. 2023. ‘Human oversight’ in the EU artificial intelligence act: what, when and by whom? *Law, Innovation and Technology*, 15(2): 508–535.
- Falco, G.; Shneiderman, B.; Badger, J.; Carrier, R.; Dahbura, A.; Danks, D.; Eling, M.; Goodloe, A.; Gupta, J.; Hart, C.; et al. 2021. Governing AI safety through independent audits. *Nature Machine Intelligence*, 3(7): 566–571.
- G7 Cybersecurity Working Group. 2025. A shared G7 vision on software bill of material for AI.
- Gao, C.; Wang, G.; Shi, W.; Wang, Z.; and Chen, Y. 2021. Autonomous driving security: State of the art and challenges. *IEEE Internet of Things Journal*, 9(10): 7572–7595.

- Gaube, S.; Langer, M.; Miller, T.; Baum, K.; Dachsel, R.; Feit, A. M.; Gadiraju, U.; Kaur, H.; Keane, M. T.; Landers, R.; Laux, J.; Liao, Q. V.; Lim, B.; Onnasch, L.; Schrills, T.; Sonenberg, L.; Tan, C.; Tintarev, N.; Xiao, Z.; and Zhang, H. 2026. Keeping an Eye on AI: A Framework for Effective Human Oversight of AI Systems. *arXiv:2605.16278*.
- Glas, M.; Vielberth, M.; and Pernul, G. 2023. Train as you fight: Evaluating authentic cybersecurity training in cyber ranges. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–19.
- Gomez-Beldarrain, G.; Verma, H.; Kim, E.; and Bozzon, A. 2025. Why does Automation Adoption in Organizations Remain a Fallacy?: Scrutinizing Practitioners’ Imaginaries in an International Airport. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, CHI ’25. New York, NY, USA: Association for Computing Machinery. ISBN 9798400713941.
- Green, B. 2022. The flaws of policies requiring human oversight of government algorithms. *Computer Law & Security Review*, 45: 105681.
- Gu, T.; Liu, K.; Dolan-Gavitt, B.; and Garg, S. 2019. Badnets: Evaluating backdoor attacks on deep neural networks. *Ieee Access*, 7: 47230–47244.
- Hall, P.; Mundahl, O.; and Park, S. 2025. The Pitfalls of “Security by Obscurity” and What They Mean for Transparent AI. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 28042–28051.
- Holovkin, B. M.; Tavolzhanskyi, O. V.; and Lysodyed, O. V. 2021. Corruption as a cybersecurity threat in the new world order. *Connections: The Quarterly Journal*, 20(2): 75–87.
- Huang, S.; Papernot, N.; Goodfellow, I.; Duan, Y.; and Abbeel, P. 2017. Adversarial attacks on neural network policies. *International conference on learning representations*.
- Hughes, C.; and Turner, T. 2023. *Software Transparency: supply chain security in an era of a software-driven society*. John Wiley & Sons.
- Iitumba, N.; Shinde, S.; Ortega, D.; Bagalkot, N.; Verdezoto, N.; Manuel, G.; Dinesh, T.; and Densmore, M. 2023. Reconsidering network management interfaces for communities. In *Proceedings of the 4th African Human Computer Interaction Conference*, 162–169.
- Jagielski, M.; Severi, G.; Pousette Harger, N.; and Oprea, A. 2021. Subpopulation data poisoning attacks. In *Proceedings of the 2021 ACM SIGSAC conference on computer and communications security*, 3104–3122.
- Jin, A.; and Salehi, N. 2024. (Beyond) Reasonable Doubt: Challenges that Public Defenders Face in Scrutinizing AI in Court. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI ’24. New York, NY, USA: Association for Computing Machinery. ISBN 9798400703300.
- Jobin, A.; Ienca, M.; and Vayena, E. 2019. The global landscape of AI ethics guidelines. *Nature machine intelligence*, 1(9): 389–399.
- Kheddar, H. 2025. Transformers and large language models for efficient intrusion detection systems: A comprehensive survey. *Information Fusion*, 103347.
- Khraisat, A.; Gondal, I.; Vamplew, P.; and Kamruzzaman, J. 2019. Survey of intrusion detection systems: techniques, datasets and challenges. *Cybersecurity*, 2(1): 1–22.
- Kohnfelder, L.; and Garg, P. 1999. The threats to our products. *Microsoft Interface, Microsoft Corporation*, 33: 1–8.
- Kopal, N.; Kieselmann, O.; Wacker, A.; and Esslinger, B. 2014. CrypTool 2.0: Open-Source Kryptologie für Jedermann. *Datenschutz und Datensicherheit-DuD*, 38(10): 701–708.
- Kraus, J.; Scholz, D.; Stiegemeier, D.; and Baumann, M. 2019. The More You Know: Trust Dynamics and Calibration in Highly Automated Driving and the Effects of Take-Overs, System Malfunction, and System Transparency. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 62(5): 718–736.
- Lai, V.; Chen, C.; Smith-Renner, A.; Liao, Q. V.; and Tan, C. 2023. Towards a science of human-AI decision making: An overview of design space in empirical human-subject studies. In *2023 ACM Conference on Fairness, Accountability, and Transparency*, 1369–1385. Chicago IL USA: ACM. ISBN 9798400701924.
- Langer, M.; Baum, K.; and Schlicker, N. 2025. Effective Human Oversight of AI-Based Systems: A Signal Detection Perspective on the Detection of Inaccurate and Unfair Outputs. *Minds and Machines*, 35(1).
- Langer, M.; Oster, D.; Speith, T.; Hermanns, H.; Kästner, L.; Schmidt, E.; Sesing, A.; and Baum, K. 2021. What do we want from Explainable artificial intelligence (XAI)? A stakeholder perspective on XAI and a conceptual model guiding interdisciplinary XAI research. *Artificial Intelligence*, 296: 103473.
- Lau, F.; Rubin, S. H.; Smith, M. H.; and Trajkovic, L. 2000. Distributed denial of service attacks. In *Smc 2000 conference proceedings. 2000 ieee international conference on systems, man and cybernetics: cybernetics evolving to systems, humans, organizations, and their complex interactions* (cat. no. 0, volume 3, 2275–2280. IEEE.
- Laux, J. 2024. Institutionalised distrust and human oversight of artificial intelligence: towards a democratic design of AI governance under the European Union AI Act. *AI & SOCIETY*, 39(6): 2853–2866.
- Laux, J.; and Ruschmeier, H. 2025. Automation Bias in the AI Act: On the Legal Implications of Attempting to De-Bias Human Oversight of AI. *European Journal of Risk Regulation*, 1–16.
- Liao, Q. V.; Gruen, D.; and Miller, S. 2020. Questioning the AI: Informing Design Practices for Explainable AI User Experiences. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI ’20, 1–15. New York, NY, USA: Association for Computing Machinery. ISBN 9781450367080.
- Liao, Q. V.; and Varshney, K. R. 2021. Human-centered explainable ai (xai): From algorithms to user experiences. *arXiv preprint arXiv:2110.10790*.
- Maasberg, M.; Van Slyke, C.; Ellis, S.; and Beebe, N. 2020. The dark triad and insider threats in cyber security. *Communications of the ACM*, 63(12): 64–80.

- Mansfield-Devine, S. 2018. The best form of defence—the benefits of red teaming. *Computer Fraud & Security*, 2018(10): 8–12.
- McBride, S. E.; Rogers, W. A.; and Fisk, A. D. 2011. Understanding the effect of workload on automation use for younger and older adults. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 53(6): 672–686.
- Meinke, A.; Schoen, B.; Scheurer, J.; Balesni, M.; Shah, R.; and Hobbhahn, M. 2024. Frontier models are capable of in-context scheming. *arXiv preprint arXiv:2412.04984*.
- Menezes, A. J.; Van Oorschot, P. C.; and Vanstone, S. A. 2018. *Handbook of applied cryptography*. CRC press.
- Miller, T. 2019. Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*, 267: 1–38.
- Oakley, J. G. 2019. *Professional red teaming: conducting successful cybersecurity engagements*. Apress.
- Oh, S.; Baek, H.; Huh, J. H.; Kim, T.; Jeon, W.; Oakley, I.; and Kim, H. 2025. Understanding and Improving User Adoption and Security Awareness in Password Checkup Services. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 1–32.
- Osanaie, O.; Choo, K.-K. R.; and Dlodlo, M. 2016. Distributed denial of service (DDoS) resilience in cloud: Review and conceptual cloud DDoS mitigation framework. *Journal of Network and Computer Applications*, 67: 147–165.
- Parker, S. K.; and Grote, G. 2020. Automation, Algorithms, and Beyond: Why Work Design Matters More Than Ever in a Digital World. *Applied Psychology*, 71(4): 1171–1204.
- Paxson, V. 2001. An analysis of using reflectors for distributed denial-of-service attacks. *ACM SIGCOMM Computer Communication Review*, 31(3): 38–47.
- Pols, P.; and van den Berg, J. 2017. The unified kill chain. *CSA Thesis, Hague*, 1–104.
- Qammar, A.; Ding, J.; and Ning, H. 2022. Federated learning attack surface: taxonomy, cyber defences, challenges, and future directions. *Artificial Intelligence Review*, 55(5): 3569–3606.
- Reda, H. T.; Anwar, A.; Mahmood, A. N.; and Tari, Z. 2023. A taxonomy of cyber defence strategies against false data attacks in smart grids. *ACM Computing Surveys*, 55(14s): 1–37.
- Ross, R.; Pillitteri, V.; Graubart, R.; Bodeau, D.; and McQuaid, R. 2019. Developing cyber resilient systems: a systems security engineering approach. Technical report, National Institute of Standards and Technology.
- Sabillon, R.; Serra-Ruiz, J.; Cavaller, V.; and Cano, J. 2017. A comprehensive cybersecurity audit model to improve cybersecurity assurance: The cybersecurity audit model (CSAM). In *2017 International Conference on Information Systems and Computer Science (INCISCOS)*, 253–259. IEEE.
- Schaab, P.; Beckers, K.; and Pape, S. 2017. Social engineering defence mechanisms and counteracting training strategies. *Information and Computer Security*, 25(2): 206–222.
- Shafahi, A.; Huang, W. R.; Najibi, M.; Suciu, O.; Studer, C.; Dumitras, T.; and Goldstein, T. 2018. Poison frogs! targeted clean-label poisoning attacks on neural networks. *Advances in neural information processing systems*, 31.
- Shah, R.; Irpan, A.; Turner, A. M.; Wang, A.; Conmy, A.; Lindner, D.; Brown-Cohen, J.; Ho, L.; Nanda, N.; Popa, R. A.; et al. 2025. An approach to technical agi safety and security. *arXiv preprint arXiv:2504.01849*.
- Shneiderman, B. 2016. The dangers of faulty, biased, or malicious algorithms requires independent oversight. *Proceedings of the National Academy of Sciences*, 113(48): 13538–13540.
- Shneiderman, B. 2020. Human-Centered Artificial Intelligence: Reliable, Safe & Trustworthy. *International Journal of Human-Computer Interaction*, 36(6): 495–504.
- Slack, D.; Hilgard, A.; Lakkaraju, H.; and Singh, S. 2021. Counterfactual explanations can be manipulated. *Advances in neural information processing systems*, 34: 62–75.
- Slack, D.; Hilgard, S.; Jia, E.; Singh, S.; and Lakkaraju, H. 2020. Fooling lime and shap: Adversarial attacks on post hoc explanation methods. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 180–186.
- Stallings, W. 1995. *Network and internetwork security: principles and practice*. Prentice-Hall, Inc.
- Sterz, S.; Baum, K.; Biewer, S.; Hermanns, H.; Lauber-Rönsberg, A.; Meinel, P.; and Langer, M. 2024. On the quest for effectiveness in human oversight: Interdisciplinary perspectives. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, 2495–2507.
- Tahaei, M.; Constantinides, M.; Quercia, D.; Kennedy, S.; Muller, M.; Stumpf, S.; Liao, Q. V.; Baeza-Yates, R.; Aroyo, L.; Holbrook, J.; Luger, E.; Madaio, M.; Blumenfeld, I. G.; De-Arteaga, M.; Vitak, J.; and Olteanu, A. 2023. Human-Centered Responsible Artificial Intelligence: Current & Future Trends. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI EA '23. New York, NY, USA: Association for Computing Machinery. ISBN 9781450394222.
- Theisen, C.; Munaiah, N.; Al-Zyoud, M.; Carver, J. C.; Meneely, A.; and Williams, L. 2018. Attack surface definitions: A systematic literature review. *Information and Software Technology*, 104: 94–103.
- Weber, S. 2017. Coercion in cybersecurity: What public health models reveal. *Journal of Cybersecurity*, 3(3): 173–183.
- Wester, J.; Pohl, H.; Hosio, S.; and van Berkel, N. 2024. "This Chatbot Would Never...": Perceived Moral Agency of Mental Health Chatbots. *Proc. ACM Hum.-Comput. Interact.*, 8(CSCW1).
- Willison, S. 2025. The lethal trifecta for AI agents: private data, untrusted content, and external communication. [Online; accessed 17-December-2025].

- Xie, S.; Zimmerman, J.; and Eslami, M. 2025. Exploring What People Need to Know to be AI Literate: Tailoring for a Diversity of AI Roles and Responsibilities. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, CHI '25. New York, NY, USA: Association for Computing Machinery. ISBN 9798400713941.
- Yadav, T.; and Rao, A. M. 2015. Technical aspects of cyber kill chain. In *International symposium on security in computing and communication*, 438–452. Springer.
- Yun, B.; Feng, D.; Chen, A. S.; Nikzad, A.; and Salehi, N. 2025. Generative AI in Knowledge Work: Design Implications for Data Navigation and Decision-Making. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, CHI '25. New York, NY, USA: Association for Computing Machinery. ISBN 9798400713941.
- Yurrita, M.; Verma, H.; Balayn, A.; Gadiraju, U.; Pont, S. C.; and Bozzon, A. 2025. Towards Effective Human Intervention in Algorithmic Decision-Making: Understanding the Effect of Decision-Makers' Configuration on Decision-Subjects' Fairness Perceptions. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, CHI '25, 1–21. ACM.
- Zhang, A. Q.; Shaw, R.; Anthis, J. R.; Milton, A.; Tseng, E.; Suh, J.; Ahmad, L.; Kumar, R. S. S.; Posada, J.; Shestakofsky, B.; et al. 2024. The human factor in ai red teaming: Perspectives from social and collaborative computing. In *Companion Publication of the 2024 Conference on Computer-Supported Cooperative Work and Social Computing*, 712–715.
- Zügner, D.; Akbarnejad, A.; and Günnemann, S. 2018. Adversarial attacks on neural networks for graph data. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, 2847–2856.