From Gaze to Data: Privacy and Societal Challenges of Using Eye-tracking Data to Inform GenAI Models

Yasmeen Abdrabou
Human-Centered Technologies for Learning, Technical University of Munich
Munich, Germany
yasmeen.abdrabou@tum.de

Suleyman Ozdel
Human-Centered Technologies for Learning, Technical University of Munich
Munich, Germany
ozdelsuleyman@tum.de

Virmarie Maquiling
Human-Centered Technologies for Learning, Technical University of Munich
Munich, Germany
virmarie.maquiling@tum.de

Efe Bozkir
Human-Centered Technologies for Learning, Technical University of Munich
Munich, Germany
efe.bozkir@tum.de

Enkelejda Kasneci
Human-Centered Technologies for Learning, Technical University of Munich
Munich, Germany
enkelejda.kasneci@tum.de

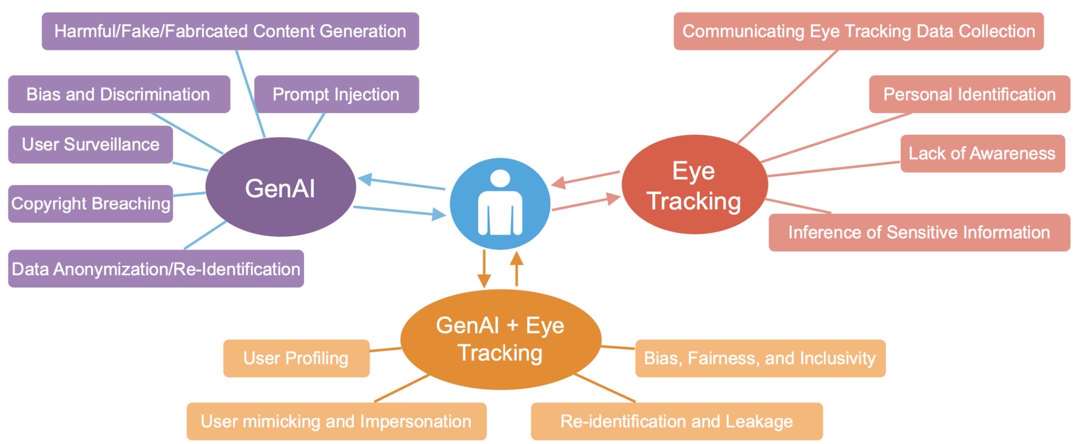

Figure 1: Overview of the privacy and societal risks associated with 1) eye-tracking technology, 2) Generative AI, and 3) eye tracking to inform generative AI with a primary focus on the user.

Abstract

Eye-tracking technology is increasingly integrated into smart glasses and wearable devices, becoming more prevalent in daily life. Meanwhile, generative artificial intelligence (GenAI) has the potential to transform user experiences through personalization and adaptive interactions. The integration of these technologies offers a novel opportunity to refine GenAI models by leveraging human gaze data in adaptive interfaces, personalized content generation, and human-computer interaction. However, gaze data is highly sensitive and can reveal several user attributes, such as cognitive states, emotions, and even medical conditions. Therefore, the use of gaze data to inform GenAI models raises significant privacy concerns. Hence, in this paper, we highlight the implications of using eye gaze information in GenAI models with privacy and societal considerations and discuss strategies to mitigate potential privacy violations. By addressing these issues, we can ensure a trade-off between technological advancements and privacy protection.

CCS Concepts

• **Computing methodologies**; • **Security and privacy**;

Keywords

Privacy, Ethics, Eye Tracking, GenAI, Conversation Agents (CA)

ETRA '25, Tokyo, Japan

ACM Reference Format:
Yasmeen Abdrabou, Suleyman Ozdel, Virmarie Maquiling, Efe Bozkir, and Enkelejda Kasneci. 2025. From Gaze to Data: Privacy and Societal Challenges of Using Eye-tracking Data to Inform GenAI Models. In *2025 Symposium on Eye Tracking Research and Applications (ETRA '25), May 26–29, 2025, Tokyo, Japan.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3715669.3726788

1 Introduction

1.1 Opportunities and Risks in Eye Tracking

Eye movements can reveal more about an individual than they may willingly disclose. Beyond simply tracking where a person looks, gaze patterns can provide insights into cognitive load [Gao et al. 2021], emotional states [Alghowinem et al. 2014], personality traits [Berkovsky et al. 2019], and even medical conditions [Shic 2016]. These capabilities have made eye tracking technology an invaluable tool in fields such as psychology [Rayner 2009; Wolf and Ueda 2021], marketing [Kim and Lee 2021; Wedel et al. 2008], education [Byrne et al. 2023; Ferdinand et al. 2024; Gao et al. 2022], and accessibility research [Edughele et al. 2022; Fischer-Janzen et al. 2024]. In human-computer interaction (HCI), eye tracking enables precise monitoring of user gaze patterns, fixations, and pupil dilation, providing valuable insights into user behavior, attention, and cognitive states [Duchowski and Duchowski 2017]. These capabilities support the evaluation of user interfaces [Novák et al. 2024], improvements in accessibility [Edughele et al. 2022; Sunny et al. 2021], and enhancements in user experience through the analysis of visual attention [Bergstrom and Schall 2014].

1.2 Rise of GenAI and its Convergence with Gaze Data

The emergence of generative artificial intelligence (GenAI) further expands the potential applications of eye-tracking technology. Recent research has begun to explore the intersection of eye tracking and GenAI, focusing on how eye movement data can inform AI models [Kędras and Sobecki 2023]. These models, such as Generative Adversarial Networks (GANs) [Goodfellow et al. 2014] and Transformer-based architectures [Vaswani et al. 2017], rely on large swathes of data to improve their predictive and generative capabilities. By incorporating gaze data, AI systems can potentially better understand human visual processing and attention, enabling more effective personalization of content and more intuitive user interfaces. For instance, GenAI could dynamically modify its content in real time based on where the user focuses or for how long, potentially making interactions more efficient and engaging [Hsieh et al. 2024]. This process would allow GenAI to better understand the user's needs and preferences, offering a more intuitive interaction that responds to attention and emotional states.

1.3 Societal and Privacy Concerns

While the fusion of these two technologies offers exciting possibilities, it also brings additional societal and privacy concerns. While eye-tracking data offers insights into human behavior, its integration into AI systems could create deeply invasive systems, manipulating users' cognitive and emotional states without their complete understanding or consent. For instance, if an AI system could detect a user's emotional state through gaze patterns and adjust its output based on this, it could be seen as manipulative or exploitative. Furthermore, merging sensitive gaze data with GenAI models creates the risk of unauthorized data usage, breaches of informed consent, and the potential for exploitation. The enhanced ability to personalize content could easily violate ethical guidelines, particularly when users are not fully aware of the extent to which their gaze data is being used. GenAI itself also poses many concerns, including privacy, especially regarding its potential to generate content that is indistinguishable from reality [Bozkir et al. 2024]. This ability raises issues related to misinformation, deepfakes, and spreading harmful content [Karnouskos 2020]. Additionally, the black-box nature of many generative models complicates the ability to trace how decisions are made, reducing transparency and trust in AI systems [Bender et al. 2021; Golda et al. 2024].

1.4 Scope of This Paper

Despite growing research into gaze-based systems and GenAI, few works have systematically examined the privacy and societal challenges at their intersection. In particular, there is a lack of frameworks addressing the unique challenges of multimodal generative systems that utilize sensitive gaze data.

As both technologies evolve, it is essential to highlight privacy risks and establish guidelines to protect user privacy responsibly. Integrating these technologies must prioritize transparency and data security, and ensure that the benefits of personalized experiences are realized without compromising user rights. Hence, this paper explores privacy risks of using eye-tracking data in GenAI models, focusing on the potential consequences of using gaze data to inform genAI systems. In addition, we also discuss the trade-off between technological advancement and responsible usage, highlighting possible strategies to mitigate potential privacy violations. In the following sections, we highlight the privacy and societal challenges associated with 1) eye tracking data, 2) GenAI models, and 3) when eye tracking is used to inform GenAI. We also show the overall structure of the risks in Figure 1.

2 Privacy Risks and Societal Challenges Associated with Eye-Tracking Data

Eye-tracking data uniquely provides deep insights into individuals' perceptions and behaviors which has led to its being studied in various fields such as education, market research, and human-computer interaction. Numerous studies have shown that gaze patterns can reveal a person's level of expertise in a given task [Castner et al. 2020b] with applications in areas such as driving skill assessment [Lang et al. 2018] and medical training [Kok and Jarodzka 2017]. Eye tracking is also widely used in market research as it can offer valuable clues about customer preferences and shopping behavior [Kim and Kim 2024; Kim and Lee 2021]. In educational settings, it has played a large role in assessing students' cognitive load [Gao et al. 2021] and detecting mind wandering [Bühler et al. 2024]. It has also shown great potential for personalizing the user experience [Bergstrom and Schall 2014], user identification [Abdrabou et al. 2024a], and has been used in analyzing and influencing user attention [Bailey et al. 2009; Maquiling et al. 2024c; Ungureanu et al. 2017]. However, while such research has led to advancements in various fields, underlying privacy concerns have, until recently, largely been overlooked. We list possible privacy and societal risks associated with collecting, processing, and using eye-tracking data in the following:

Personal Identification. Gaze patterns can serve as a powerful biometric identifier, enabling accurate user identification even with minimal contextual data [Bozkir et al. 2021; Holmqvist and Andersson 2017; Steil et al. 2019a]. Studies have demonstrated that simple eye-tracking features, such as fixations and saccades, are sufficient for distinguishing individuals with high accuracy. Several works have explored user identification through gaze behavior [Abdrabou et al. 2024a; Liebers et al. 2021], biometric authentication [Yoon et al. 2014], and even implicit recognition during passive activities like viewing images [Abdrabou et al. 2024a]. Moreover, combining gaze data with other biometric information (e.g., facial expressions, voice, or heart rate) can lead to highly detailed personal profiles, increasing privacy risks [Kröger et al. 2020]. While gaze-based authentication can enhance security by restricting device access to unauthorized users, it also introduces significant privacy concerns. Specifically, gaze-based identification extends beyond single-use applications, potentially enabling user re-identification across multiple datasets [Al Zaidawi et al. 2022], raising risks of deanonymization and unauthorized tracking.

Inference of Sensitive Information. Eye tracking can reveal a surprising amount of personal information about an individual–perhaps more than they may be comfortable disclosing. Research has shown that gaze patterns can indicate mental states [Vidal et al. 2012], cognitive load [Bühler et al. 2024; Zagermann et al. 2016], preferences [Cheng and Liu 2012]. Furthermore, eye movements have been linked to the diagnosis of health conditions such as neurological disorders [Kröger et al. 2020; Liebling and Preibusch 2014; Tao et al. 2020; Terao et al. 2017]. Beyond the health domain, eye tracking data has been used to infer personality traits [Berkovsky et al. 2019; Hoppe et al. 2018], physical attributes such as age [Erbilek et al. 2013], and gender [Dibeklioğlu et al. 2012], cultural background [Haensel et al. 2022], ongoing activities [Bulling et al. 2009, 2010], passwords [Abdrabou et al. 2024b; Wang et al. 2024], security decisions [Abdrabou et al. 2022, 2021] and even their level of expertise [Castner et al. 2020a,b]. As the field ventures outside traditional video oculography with the emergence of novel systems such as event-based eye trackers [Angelopoulos et al. 2021], micro electro mechanical (MEMS) based systems [Meyer et al. 2020; Pomianek et al. 2021], and retinal eye trackers [Bartuzel et al. 2020], the potential to infer even more sensitive information is increasing, introducing new challenges and privacy concerns regarding the extent of personal data that can be extracted. For a more detailed exploration of these findings, refer to [Kröger et al. 2020].

Communicating Eye-Tracking Data Collection. Another significant privacy risk is the lack of transparent communication regarding eye-tracking data collection. For instance, if a website passively collects eye-tracking data in the background, it may obtain users' consent once, but over time, users may forget that their gaze is still being monitored, leading to unintended privacy invasions. This risk is even greater when video-based eye trackers are used, as they not only capture eye movements but also record facial features, further highlighting privacy concerns. Additionally, eye-tracking technology does not only affect the primary user—it can extend to bystanders in the scene, raising various challenges and legal concerns [Alsakar et al. 2023]. Moreover, wearable eye trackers typically feature a world-facing camera, which can seamlessly record bystanders, further highlighting privacy risks. Without clear disclosure and proper privacy measures, these technologies show serious implications for personal privacy and consent.

Lack of Awareness. Users may not be fully aware of when and how their gaze data is being collected, making it challenging to obtain truly informed consent [Steil et al. 2019a,b]. Eye tracking involves multiple stages of data collection, including raw gaze capture (typically, infrared eye images for Video-oculography (VOG) based eye trackers), preprocessing, feature extraction, and higher-level data analysis. Each step potentially introduces additional privacy risks. The lack of transparency in these processes further exacerbates concerns, as users may not fully understand the extent to which their gaze data can be analyzed or interpreted.

3 Privacy Challenges Associated with GenAI

GenAI systems are designed to generate content such as text, images, videos, 3D models, audio, and code by recognizing and extrapolating patterns from extensive training data [Feuerriegel et al. 2024]. As technology evolves rapidly, it brings transformative potential across various domains. The wide-ranging applicability of genAI has significantly impacted domains such as healthcare [Cascella et al. 2023; Zhang and Kamel Boulos 2023], finance [Chen et al. 2023; Li et al. 2023] and education [Kasneci et al. 2023], facilitating more efficient systems. These models rely on substantial amounts of data to effectively learn and generalize patterns, typically sourced indiscriminately from the web without rigorous selection or curation [Bender et al. 2021]. However, GenAI introduces multiple critical risks, encompassing privacy violations, security vulnerabilities, and ethical dilemmas [Das et al. 2025; Yao et al. 2024].

Privacy risks emerge primarily from the extensive collection and storage of potentially sensitive personal data, which can inadvertently expose individuals to misuse, unauthorized access, or unintentional disclosures. In parallel, large generative models enable the creation of hyper-realistic misinformation, often referred to as hallucinations, deepfakes, or other malicious content, posing significant threats to societal trust and stability. Additionally, ethical concerns emerge as these models can amplify existing biases, issues of fairness and accountability, and transparent and responsible deployment of generative technologies. In this context, we identify the aforementioned risks categorized as follows:

Copyright Breaching. The training datasets used for GenAI models may include copyrighted material, raising significant legal and ethical concerns. As a result, AI-generated content can unintentionally replicate, incorporate, or closely resemble copyrighted works present in the training data [Carlini et al. 2023; Sag 2023; Samuelson 2023], potentially leading to intellectual property disputes. This issue is further complicated by the difficulty of tracing specific copyrighted elements within most datasets, making it challenging to determine whether AI outputs constitute breaches. Additionally, the lack of clear guidelines on ownership and attribution of AI-generated content further exacerbates these legal ambiguities.

Prompt Injection. Prompt injection risk happens when an attacker manipulates a GenAI by providing specially crafted inputs, causing the model to unknowingly carry out the attacker's intended actions [Abdali et al. 2024]. Such prompts could also extract sensitive information, disrupt services, or produce harmful and misleading content [Liu et al. 2023, 2024]. In addition, these models can be repurposed for malicious goals, straying far from their intended use and serving harmful objectives [Achiam et al. 2023; Brown et al. 2020]. This possibility makes them vulnerable to manipulation for things like spreading misinformation, committing fraud, or other unethical activities. Without proper oversight and safeguards in place, such misuse can lead to reputational damage, legal troubles, and wider societal harm—ultimately shaking public confidence in AI technologies.

Harmful/Fake/Fabricated Content Generation. Building on the concept of prompt injection, GenAI systems can also produce fake or harmful content when exposed to malicious prompts. Without proper moderation or constraints, these systems may unintentionally generate content that is misleading, offensive, harmful, or even dangerous [Achiam et al. 2023; Weidinger et al. 2021]. Such outputs can fuel the spread of misinformation, amplify hate speech, cause psychological harm, or exacerbate societal tensions. However, the true risk here is that this could be supported by fabricated references, deepfake footage, or even TV snippets, which will lead to user manipulation, causing actual harm. This highlights the urgent need for robust content governance frameworks to mitigate these challenges and ensure the responsible use of AI technologies.

Data Anonymization/Re-Identification. Even data that has been anonymized can sometimes be reverse-engineered or linked to external datasets to re-identify individuals [Kibriya et al. 2024; Patsakis and Lykousas 2023]. This phenomenon threatens user privacy and increases the vulnerability of individuals whose data is presumed secure. GenAI models may inadvertently memorize and subsequently expose sensitive data embedded in their training datasets [Carlini et al. 2023; Zhou et al. 2025]. Such leakage risks jeopardize personal or confidential information, posing severe threats to user privacy and data security. Additionally, with membership inference attacks, adversaries can determine whether a person's data was used in training based on the model's output [Hayes et al. 2017; Mireshghallah et al. 2022]. Furthermore, privacy challenges extend beyond training data, also arising during real-time interactions between users and generative conversational agents. Users often disclose personally identifiable information (PII) and sensitive details about third parties, particularly in domains such as healthcare, finance, or emotional support, without fully realizing the potential privacy challenges [Zhang et al. 2024a]. Lack of user awareness regarding the storage, processing, or reuse of their data by these systems increases the risk of unintended exposure.

User Surveillance. GenAI can also be leveraged for user surveillance and monitoring [Singh 2023]. With its ability to analyze and learn user behavior, thought patterns, and interactions, it has the potential to function as a sophisticated surveillance tool. While such capabilities could be applied to beneficial use cases, such as detecting and preventing extreme or harmful activities, they also pose significant privacy challenges for everyday users. The ability to track, predict, and infer personal behaviors raises serious concerns about consent, data misuse, and the erosion of privacy in both digital and physical spaces.

Bias and Discrimination. Like all trained models, genAI systems inherit biases present in the real-world data they are trained on [Bender et al. 2021; Gallegos et al. 2024; Motoki et al. 2024; Nadeem et al. 2020]. As a result, they can unintentionally reinforce and amplify existing societal biases, leading to discriminatory results or unfair treatment of certain demographic groups. These biases appear in multiple forms, from skewed representations in generated content to biased decision-making in AI-driven applications, ultimately embedding inequality and reducing trust in AI.

4 Privacy Risks and Further Challenges Associated with GenAI and Eye Tracking

The integration of eye tracking with genAI has demonstrated significant potential across various domains, including enhancing reading comprehension [Yang et al. 2024; Zhang et al. 2024b], improving virtual reality experiences [Lau et al. 2024], and refining vision models [Maquiling et al. 2024a,b; Yan et al. 2024]. However, this integration also introduces new risks. As discussed in previous sections, both eye tracking and GenAI present independent privacy and ethical challenges. When combined, these risks are not only amplified but also evolve into novel concerns.

Eye-tracking technology records highly detailed physiological and behavioral data, while GenAI can model, predict, and adapt to user behavior. The combination of these technologies enables AI systems to leverage gaze data at different stages—either during training to refine the model's understanding of gaze dynamics or in real-time applications where gaze patterns directly influence generative outputs [Konrad et al. 2024]. This capability can enhance personalization, encourage more empathetic human-computer interactions, and optimize user experiences by assessing cognitive load, detecting mood fluctuations, and identifying task-related difficulties in real time.

Despite these advantages, the integration of eye tracking with GenAI raises privacy and ethical concerns. Furthermore, genAI systems often incorporate multiple data modalities, where eye-tracking data may be combined with textual, visual, or conversational inputs. This multimodal integration heightens the risk of extracting unintended and potentially harmful inferences about users, raising pressing questions about consent, transparency, and data security. Table 1 shows some real-life scenarios of how such risk could affect the users. In this section, we list possible privacy risks and challenges arising from the integration of both technologies.

User Profiling and Manipulation. The integration of eye-tracking with GenAI heightens the potential for both advanced user profiling and behavioral manipulation. Eye tracking provides an in-depth view into users' cognitive states, emotions, and attention patterns, which, when combined with GenAI's data-processing capabilities, can lead to hyper-personalized experiences. While such integrations can enhance user well-being by offering tailored support, they also come with significant privacy and ethical concerns. Gaze patterns, often collected implicitly, can be processed without explicit consent, leading to unintended profiling and invasive surveillance. Moreover, biases in GenAI—which already lead to discriminatory outcomes—could be amplified by this detailed user data, further deepening personalized and targeted manipulation of users' emotions, decisions, or behaviors. The result is a lack of user agency, as

Table 1: Taxonomy of privacy risks in eye tracking, GenAI, and their integration represented with different risk levels and possible mitigation strategies. Where High risk requires immediate attention due to severe privacy, security, or ethical concerns. Medium risk moderately concerning but can be managed with appropriate measures, and low risk minimal threat but should be considered in privacy/security strategies.

Category	Risk	Eye Tracking (Standalone)	GenAI (Standalone)	Eye Tracking + GenAI Combination	Mitigation Strategy	Risk Level
User Profiling & Manipulation	Hyper-personalized behavioral tracking	Tracks user attention, cognitive states, emotions.	AI personalizes content but may manipulate users.	Gaze enables real-time adaptation, raising manipulation concerns.	Data minimization, transparency, opt-out controls, ethical AI design	High
	Reinforcement of hallucinations	Eye movements can be used as feedback mechanism.	AI may reinforce false or misleading information.	Gaze-based feedback loops can amplify AI-generated misinformation.	Fact-checking, user alerts for hallucinations, AI model verification	High
	Emotional influence through AI	Eye movements reveal emotions without explicit consent.	AI detects and exploits emotional responses.	AI can adapt responses based on involuntary gaze reactions.	Emotion-aware consent, transparency labels, user control	Medium
User Mimicking & Impersonation	Biometric authentication risks	Gaze biometrics used for authentication.	AI can be manipulated to bypass security.	AI-generated gaze patterns could impersonate users.	Multi-factor authentication	High
	AI-generated user behavioral clones	Unique gaze dynamics can be mimicked.	Deepfakes can replicate voices, faces, behaviors.	AI could generate gaze-driven deepfakes for fraud.	Watermarking AI-generated content, behavioral detection mechanisms	High
	Phishing attacks leveraging gaze behavior	Gaze tracking reveals reading patterns in real time.	AI-generated models mimic real user patterns.	Attackers use gaze data to create tailored phishing attempts.	AI-based anomaly detection, contextual security prompts	Medium
Bias, Fairness & Inclusivity	Discrimination due to dataset bias	Datasets often lack demographic diversity.	AI models inherit biases from training data.	Biases in both technologies could amplify discrimination.	Prioritization of diverse datasets, and bias mitigation	Medium
	Exclusion from personalized experiences	Gaze tracking fails for individuals with disabilities.	AI struggles with diverse or non-mainstream inputs.	AI may personalize experiences ineffectively for certain groups.	Inclusive dataset training	Low
	Unfair treatment based on gaze behavior	Gaze tracking less accurate for some disabilities.	AI misinterprets non-standard gaze patterns.	AI systems may unintentionally discriminate based on gaze traits.	AI fairness guidelines, gaze-aware accessibility	Medium
Re-identification & Leakage	Unintended user re-identification	Gaze patterns can serve as biometric identifiers.	AI memorizes and leaks sensitive data.	Multimodal fusion increases the risk of user tracking.	Data encryption, user control over gaze data	High
	Data leakage through multimodal AI systems	Gaze data can be combined with other identifiers.	AI models leak training data in responses.	Cross-modal AI processing heightens privacy risks.	Privacy-preserving techniques	High
	Unauthorized third-party access to gaze data	Gaze data stored and processed by multiple platforms.	AI integrates gaze inputs without clear consent policies.	High risk of personal data exposure to untrusted entities.	Secure APIs and data-sharing protocols (GDPR/AI Compliant)	High
	Unintended gaze-based advertising	Gaze tracking enables real-time ad targeting.	AI personalizes ads based on behavior.	AI could hyper-target ads based on microsecond gaze shifts.	Explicit user consent and opt-out controls	Low

they may unknowingly be influenced or even manipulated based on sensitive attributes (e.g., gender, sexual orientation, beliefs) without proper transparency or control. Additionally, gaze can be utilized as an implicit feedback to refine its outputs based on user engagement. However, in the presence of hallucinations, increased user engagement with generated content may reinforce erroneous model outputs. This feedback loop could inadvertently amplify hallucinations, as the system interprets user attention as validation, leading to further propagation of misleading or inaccurate information.

User Mimicking and Impersonation. The detailed behavioral data captured through eye tracking opens a new direction to mimicking and impersonating users in ways that could compromise both security and privacy. GenAI, when coupled with eye-tracking information, can create highly accurate representations of users' unique gaze dynamics, which may then be a victim of unauthorized actions like identity theft or bypassing authentication systems (e.g., bypassing biometric verification that relies on eye movements). This introduces several risks in cybersecurity, as attackers could easily manipulate these systems to gain unauthorized access, especially in areas like digital trust and online interactions. The replication of personal user behaviors opens new directions for deception and fraud, which need to be mitigated by more secure systems.

Bias, Fairness, and Inclusivity. Bias remains a key concern when combining eye-tracking technology with GenAI models. Existing biases within GenAI models and the lack of diversity in eye-tracking datasets, such as underrepresentation of different demographics such as ethnicity, gender, and age [Drozdowski et al. 2020; Prasse et al. 2022] poses a major challenge. This underrepresentation can result in discriminatory outcomes that favor more represented groups, leading to biased AI performance. The absence of diverse, standardized eye-tracking datasets highlights the problem by reinforcing the potential for biased gaze analysis, which could result in misinterpretation of user behavior, leading to unfair treatment. Such biases impact the accessibility, fairness, and inclusivity of

AI systems, limiting the ability of marginalized groups to benefit from these technologies, making them less effective, and potentially perpetuating existing societal inequities.

Re-identification and Leakage. Data leakage risks in GenAI models are amplified when eye-tracking data is included, as gaze patterns can serve as unique identifiers. When combined with other types of personal data (e.g., visual, textual, conversational), this information increases the likelihood of re-identification even in anonymized datasets. Integrating multiple data modalities allows for the reconstruction of detailed user profiles, posing significant risks of data breaches and privacy violations. Even without malicious intent, models may accidentally disclose private user details or trigger unintended exposure due to how deeply integrated eye-tracking data can be within the system. This highlights the importance of robust privacy measures to prevent unauthorized disclosure and protect the user's identity.

5 Mitigation Strategies

To address privacy challenges, while maintaining the benefits of gaze-based GenAI algorithms, we first address the challenges in each technology on its own, and hence, the overall risk could be reduced. Hence, we propose the following mitigation techniques:

Informed Consent and User Communication. : Users must be explicitly informed about how and what their data will be collected, stored, and processed. This step is especially important if the user needs to consent once and the data will be collected every time the application is launched. Hence, applications might include contextual consent prompts during data collection or model interaction. Providing real-time updates on what is tracked, the sensitivity of the information, and how it is utilized (e.g., "Your gaze data is used to personalize your user interface") empowers users to make informed decisions and foster greater trust. Additionally, systems should limit granular inferences and avoid storing specific biomarkers (such as pupil diameters) unless strictly required [Goldsteen et al. 2022].

Raising Users' Awareness. : Most of the users do not know how sensitive their gaze data is, what it can reveal [Alsakar et al. 2023], or how GenAI can profile them [Zhang et al. 2024a]. Hence, we need to raise users' awareness and educate them about eye tracking and its impact on GenAI. In addition, we also need to implement user-controlled data sharing, contextual consent, and transparency for opt-in mechanisms with fine-grained permissions for users to control when and how their gaze data is shared and also the option for opting out.

Privacy-Preserving Data Collection. : As data is the new currency in the world, we need to have privacy-preserving data collection techniques to protect the users without affecting their experience. Techniques such as differential privacy, local processing, and data anonymization can minimize exposure risks and have already shown great potential [Agrawal and Srikant 2000; Bozkir et al. 2021; Steil et al. 2019a]. These approaches ensure that individual gaze trajectories are less distinguishable while still allowing the generative model to learn overall statistical patterns. In addition, encrypting gaze data and ensuring secure access mechanisms can reduce the likelihood of unauthorized breaches [Özdel et al. 2024].

Inclusive Datasets: The absence of inclusive datasets for both GenAI and eye-tracking technologies significantly contributes to bias and discrimination. To address these issues, it is essential to prioritize the development of diverse, representative eye-tracking datasets and ensure that AI models are trained on data that accurately reflects the diversity of users. Without these foundational steps, the integration of GenAI with eye-tracking systems could exacerbate existing biases, hindering efforts to promote inclusivity and potentially limiting the effectiveness of these technologies in delivering fair user experiences. Furthermore, building a robust framework for compliance, transparent documentation of training data sources, and clear guidelines for authorship and liability are crucial. Implementing mechanisms such as dataset curation, licensing agreements, and AI watermarking can also help mitigate challenges related to copyright concerns.

Regulatory Compliance. : Existing laws, such as the General Data Protection Regulation (GDPR) in Europe and the California Consumer Privacy Act (CCPA) in the U.S., mandate strict guidelines on data collection and user privacy [Parliament 2016]. Recently, the European Union AI Act established a risk-based framework directly relevant to integrating eye-tracking data with genAI. Biometric identification, including gaze-based biometrics, may be classified as high-risk, requiring oversight, transparency, and risk management, with restricted applications [AIA 2025]. In summary, adherence to these regulatory frameworks is crucial for the responsible integration of eye-tracking data with genAI. Ensuring compliance with evolving laws will help mitigate ethical and legal risks while enhancing user trust and transparency in AI-driven systems.

6 Conclusion

Using gaze behavior as a source of information for generative algorithms offers exciting opportunities but also raises significant privacy risks. As AI continues to advance, it is important to establish clear ethical guidelines and implement robust privacy measures to protect users' personal data. By enhancing transparency, ensuring regulatory compliance, and prioritizing user control, we can help balance innovation with responsibility, ensuring that gaze-based AI applications are developed and deployed responsibly. In this paper, we have examined the privacy and societal challenges associated with both eye-tracking technologies and genAI, as well as the additional risks arising when eye-tracking data informs AI models. We have also proposed several mitigation strategies to enhance technology usage and guide future research. We are confident that this paper serves as both an awareness tool for the community and a starting point for future work, sparking exploration and encouraging responsible development for the collective benefit of society.

Acknowledgments

The project is supported by the Chips Joint Undertaking (Chips JU) and its members, including top-up funding by Denmark, Germany, Netherlands, Sweden, under grant agreement No. 101139942.

References

2025. EU Artificial Intelligence Act. https://artificialintelligenceact.eu/ Accessed: 2025-03-17.

Sara Abdali, Richard Anarfi, CJ Barberan, and Jia He. 2024. Securing large language models: Threats, vulnerabilities and responsible practices. *arXiv preprint arXiv:2403.12503* (2024).

L Abdrabou, Mariam Hassib, Shuqin Hu, Ken Pfeuffer, Mohamed Khamis, Andreas Bulling, and Florian Alt. 2024a. EyeSeeIdentity: Exploring Natural Gaze Behaviour for Implicit User Identification during Photo Viewing. (01 2024). doi:10.14722/usec.2024.23057

Yasmeen Abdrabou, Tatiana Omelina, Felix Dietz, Mohamed Khamis, Florian Alt, and Mariam Hassib. 2024b. Where Do You Look When Unlocking Your Phone?: A Field Study of Gaze Behaviour During Smartphone Unlock. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI EA '24)*. Association for Computing Machinery, New York, NY, USA, Article 374, 7 pages. doi:10.1145/3613905.3651094

Yasmeen Abdrabou, Johannes Schütte, Ahmed Shams, Ken Pfeuffer, Daniel Buschek, Mohamed Khamis, and Florian Alt. 2022. "Your Eyes Tell You Have Used This Password Before": Identifying Password Reuse from Gaze and Keystroke Dynamics. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 400, 16 pages. doi:10.1145/3491102.3517531

Yasmeen Abdrabou, Ahmed Shams, Mohamed Omar Mantawy, Anam Ahmad Khan, Mohamed Khamis, Florian Alt, and Yomna Abdelrahman. 2021. GazeMeter: Exploring the Usage of Gaze Behaviour to Enhance Password Assessments. In *ACM Symposium on Eye Tracking Research and Applications* (Virtual Event, Germany) *(ETRA '21 Full Papers)*. Association for Computing Machinery, New York, NY, USA, Article 9, 12 pages. doi:10.1145/3448017.3457384

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774* (2023).

R. Agrawal and R. Srikant. 2000. Privacy-Preserving Data Mining. In *Proceedings of the ACM SIGMOD Conference on Management of Data.*

Sahar Mahdie Klim Al Zaidawi, Martin H.U. Prinzler, Jonas Lührs, and Sebastian Maneth. 2022. An extensive study of user identification via eye movements across multiple datasets. *Signal Processing: Image Communication* 108 (2022), 116804. doi:10.1016/j.image.2022.116804

Sharifa Alghowinem, Majdah AlShehri, Roland Goecke, and Michael Wagner. 2014. *Exploring Eye Activity as an Indication of Emotional States Using an Eye-Tracking Sensor.* Springer International Publishing. doi:10.1007/978-3-319-04702-7_15

Noora Alsakar, Yasmeen Abdrabou, Simone Stumpf, and Mohamed Khamis. 2023. Investigating Privacy Perceptions and Subjective Acceptance of Eye Tracking on Handheld Mobile Devices. *Proc. ACM Hum.-Comput. Interact.* 7, ETRA, Article 164 (May 2023), 16 pages. doi:10.1145/3591133

Anastasios N Angelopoulos, Julien NP Martel, Amit P Kohli, Jorg Conradt, and Gordon Wetzstein. 2021. Event-Based Near-Eye Gaze Tracking Beyond 10,000 Hz. *IEEE transactions on visualization and computer graphics* 27, 5 (2021), 2577–2586.

Reynold Bailey, Ann McNamara, Nisha Sudarsanam, and Cindy Grimm. 2009. Subtle gaze direction. *ACM Transactions on Graphics (TOG)* 28, 4 (2009), 1–14.

Maciej M Bartuzel, Krystian Wróbel, Szymon Tamborski, Michał Meina, Maciej Nowakowski, Krzysztof Dalasiński, Anna Szkulmowska, and Maciej Szkulmowski. 2020. High-resolution, ultrafast, wide-field retinal eye-tracking for enhanced quantification of fixational and saccadic motion. *Biomedical Optics Express* 11, 6 (2020), 3164–3180.

Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big?. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency.* 610–623.

Jennifer Romano Bergstrom and Andrew Schall. 2014. *Eye tracking in user experience design.* Elsevier.

Shlomo Berkovsky, Ronnie Taib, Irena Koprinska, Eileen Wang, Yucheng Zeng, Jingjie Li, and Sabina Kleitman. 2019. Detecting personality traits using eye-tracking data. In *Proceedings of the 2019 CHI conference on human factors in computing systems.* 1–12.

Efe Bozkir, Onur Günlü, Wolfgang Fuhl, Rafael F. Schaefer, and Enkelejda Kasneci. 2021. Differential privacy for eye tracking with temporal correlations. *PLOS ONE* 16, 8 (2021). doi:10.1371/journal.pone.0255979

Efe Bozkir, Clara Riedmiller, Athanassios N. Skodras, Gjergji Kasneci, and Enkelejda Kasneci. 2024. Can You Tell Real from Fake Face Images? Perception of Computer-Generated Faces by Humans. *ACM Transactions on Applied Perception* 22, 2 (2024). doi:10.1145/3696667

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.

Babette Bühler, Efe Bozkir, Hannah Deininger, Patricia Goldberg, Peter Gerjets, Ulrich Trautwein, and Enkelejda Kasneci. 2024. Detecting Aware and Unaware Mind Wandering During Lecture Viewing: A Multimodal Machine Learning Approach Using Eye Tracking, Facial Videos and Physiological Data. In *Proceedings of the 26th International Conference on Multimodal Interaction.* ACM. doi:10.1145/3678957.3685710

Andreas Bulling, Jamie A Ward, Hans Gellersen, and Gerhard Tröster. 2009. Eye movement analysis for activity recognition. In *Proceedings of the 11th international conference on Ubiquitous computing.* 41–50.

Andreas Bulling, Jamie A Ward, Hans Gellersen, and Gerhard Tröster. 2010. Eye movement analysis for activity recognition using electrooculography. *IEEE transactions on pattern analysis and machine intelligence* 33, 4 (2010), 741–753.

Sean Anthony Byrne, Nora Castner, Ard Kastrati, Martyna Beata Płomecka, William Schaefer, Enkelejda Kasneci, and Zoya Bylinskii. 2023. Leveraging eye tracking in digital classrooms: A step towards multimodal model for learning assistance. In *Proceedings of the 2023 Symposium on Eye Tracking Research and Applications.* 1–6.

Nicolas Carlini, Jamie Hayes, Milad Nasr, Matthew Jagielski, Vikash Sehwag, Florian Tramer, Borja Balle, Daphne Ippolito, and Eric Wallace. 2023. Extracting training data from diffusion models. In *32nd USENIX Security Symposium (USENIX Security 23).* 5253–5270.

Marco Cascella, Jonathan Montomoli, Valentina Bellini, and Elena Bignami. 2023. Evaluating the feasibility of ChatGPT in healthcare: an analysis of multiple clinical and research scenarios. *Journal of medical systems* 47, 1 (2023), 33.

Nora Castner, Lea Geßler, David Geisler, Fabian Hüttig, and Enkelejda Kasneci. 2020a. Towards expert gaze modeling and recognition of a user's attention in realtime. *Procedia Computer Science* 176 (2020).

Nora Castner, Thomas C Kuebler, Katharina Scheiter, Juliane Richter, Therese Eder, Fabian Huettig, Constanze Keutel, and Enkelejda Kasneci. 2020b. Deep semantic gaze embedding and scanpath comparison for expertise classification during OPT viewing. In *ACM Symposium on Eye Tracking Research and Applications.* ACM. doi:10.1145/3379155.3391320

Boyang Chen, Zongxiao Wu, and Ruoran Zhao. 2023. From fiction to fact: the growing role of generative AI in business and finance. *Journal of Chinese Economic and Business Studies* 21, 4 (2023), 471–496.

Shiwei Cheng and Ying Liu. 2012. Eye-tracking based adaptive user interface: implicit human-computer interaction for preference indication. *Journal on Multimodal User Interfaces* 5 (2012), 77–84.

Badhan Chandra Das, M Hadi Amini, and Yanzhao Wu. 2025. Security and privacy challenges of large language models: A survey. *Comput. Surveys* 57, 6 (2025), 1–39.

Hamdi Dibeklioğlu, Theo Gevers, Albert Ali Salah, and Roberto Valenti. 2012. A smile can reveal your age: Enabling facial dynamics in age estimation. In *Proceedings of the 20th ACM international conference on Multimedia.* 209–218.

Pawel Drozdowski, Christian Rathgeb, Antitza Dantcheva, Naser Damer, and Christoph Busch. 2020. Demographic bias in biometrics: A survey on an emerging challenge. *IEEE Transactions on Technology and Society* 1, 2 (2020), 89–103.

Andrew T Duchowski and Andrew T Duchowski. 2017. *Eye tracking methodology: Theory and practice.* Springer.

Hilary O Edughele, Yinghui Zhang, Firdaus Muhammad-Sukki, Quoc-Tuan Vien, Haley Morris-Cafiero, and Michael Opoku Agyeman. 2022. Eye-tracking assistive technologies for individuals with amyotrophic lateral sclerosis. *IEEE Access* 10 (2022), 41952–41972.

Meryem Erbilek, Michael Fairhurst, and Márjory Cristiany Da Costa Abreu. 2013. Age prediction from iris biometrics. In *5th International Conference on Imaging for Crime Detection and Prevention (ICDP 2013).* IET, 1–07.

Joseph Ferdinand, Hong Gao, Philipp Stark, Efe Bozkir, Jens-Uwe Hahn, Enkelejda Kasneci, and Richard Göllner. 2024. The impact of a usefulness intervention on students' learning achievement in a virtual biology lesson: An eye-tracking-based approach. *Learning and Instruction* 90 (2024), 101867.

Stefan Feuerriegel, Jochen Hartmann, Christian Janiesch, and Patrick Zschech. 2024. Generative ai. *Business & Information Systems Engineering* 66, 1 (2024), 111–126.

Anke Fischer-Janzen, Thomas M Wendt, and Kristof Van Laerhoven. 2024. A scoping review of gaze and eye tracking-based control methods for assistive robotic arms. *Frontiers in Robotics and AI* 11 (2024), 1326670.

Isabel O Gallegos, Ryan A Rossi, Joe Barrow, Md Mehrab Tanjim, Sungchul Kim, Franck Dernoncourt, Tong Yu, Ruiyi Zhang, and Nesreen K Ahmed. 2024. Bias and fairness in large language models: A survey. *Computational Linguistics* 50, 3 (2024), 1097–1179.

Hong Gao, Efe Bozkir, Lisa Hasenbein, Jens-Uwe Hahn, Richard Göllner, and Enkelejda Kasneci. 2021. Digital Transformations of Classrooms in Virtual Reality. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems.* ACM. doi:10.1145/3411764.3445596

Hong Gao, Lisa Hasenbein, Efe Bozkir, Richard Göllner, and Enkelejda Kasneci. 2022. Evaluating the effects of virtual human animation on students in an immersive vr classroom using eye movements. In *Proceedings of the 28th ACM Symposium on Virtual Reality Software and Technology.* 1–11.

Abenezer Golda, Kidus Mekonen, Amit Pandey, Anushka Singh, Vikas Hassija, Vinay Chamola, and Biplab Sikdar. 2024. Privacy and security concerns in generative AI: a comprehensive survey. *IEEE Access* (2024).

Abigail Goldsteen, Gilad Ezov, Ron Shmelkin, Micha Moffie, and Ariel Farkash. 2022. Data minimization for GDPR compliance in machine learning models. *AI and Ethics*

2, 3 (2022), 477–491.

Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. *Advances in neural information processing systems* 27 (2014).

Jennifer X Haensel, Tim J Smith, and Atsushi Senju. 2022. Cultural differences in mutual gaze during face-to-face interactions: A dual head-mounted eye-tracking study. *Visual Cognition* 30, 1-2 (2022), 100–115.

Jamie Hayes, Luca Melis, George Danezis, and Emiliano De Cristofaro. 2017. Logan: Membership inference attacks against generative models. *arXiv preprint arXiv:1705.07663* (2017).

Kenneth Holmqvist and Richard Andersson. 2017. *Eye-tracking: A comprehensive guide to methods, paradigms and measures.*

Sabrina Hoppe, Tobias Loetscher, Stephanie A Morey, and Andreas Bulling. 2018. Eye movements during everyday behavior predict personality traits. *Frontiers in human neuroscience* 12 (2018), 328195.

He-Yen Hsieh, Ziyun Li, Sai Qian Zhang, Wei-Te Mark Ting, Kao-Den Chang, Barbara De Salvo, Chiao Liu, and HT Kung. 2024. GazeGen: Gaze-Driven User Interaction for Visual Content Generation. *arXiv preprint arXiv:2411.04335* (2024).

Stamatis Karnouskos. 2020. Artificial intelligence in digital media: The era of deepfakes. *IEEE Transactions on Technology and Society* 1, 3 (2020), 138–147.

Enkelejda Kasneci, Kathrin Seßler, Stefan Küchemann, Maria Bannert, Daryna Dementieva, Frank Fischer, Urs Gasser, Georg Groh, Stephan Günnemann, Eyke Hüllermeier, et al. 2023. ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and individual differences* 103 (2023), 102274.

Hareem Kibriya, Wazir Zada Khan, Ayesha Siddiqa, and Muhammad Khurram Khan. 2024. Privacy issues in large language models: a survey. *Computers and Electrical Engineering* 120 (2024), 109698.

Ju Yeon Kim and Mi Jeong Kim. 2024. Identifying customer preferences through the eye-tracking in travel websites focusing on neuromarketing. *Journal of Asian Architecture and Building Engineering* 23, 2 (2024), 515–527.

Nayeon Kim and Hyunsoo Lee. 2021. Assessing consumer attention and arousal using eye-tracking technology in virtual retail environment. *Frontiers in Psychology* 12 (2021), 665658.

Ellen M Kok and Halszka Jarodzka. 2017. Before your very eyes: the value and limitations of eye tracking in medical education. *Medical Education* 51, 1 (2017). doi:10.1111/medu.13066

Robert Konrad, Nitish Padmanaban, J Gabriel Buckmaster, Kevin C Boyle, and Gordon Wetzstein. 2024. Gazegpt: Augmenting human capabilities using gaze-contingent contextual ai for smart eyewear. *arXiv preprint arXiv:2401.17217* (2024).

Jacob Leon Kröger, Otto Hans-Martin Lutz, and Florian Müller. 2020. What does your gaze reveal about you? On the privacy implications of eye tracking. In *IFIP International Summer School on Privacy and Identity Management*. Springer, 226–241.

Maja Kędras and Janusz Sobecki. 2023. What Is Hidden in Clear Sight and How to Find It—A Survey of the Integration of Artificial Intelligence and Eye Tracking. *Information* 14, 11 (2023). doi:10.3390/info14110624

Yining Lang, Liang Wei, Fang Xu, Yibiao Zhao, and Lap-Fai Yu. 2018. Synthesizing Personalized Training Programs for Improving Driving Habits via Virtual Reality. In *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. doi:10.1109/VR.2018.8448290

Ka Hei Carrie Lau, Sema Sen, Philipp Stark, Efe Bozkir, and Enkelejda Kasneci. 2024. Personalized Generative AI in VR for Enhanced Engagement: Eye-Tracking Insights into Cultural Heritage Learning through Neapolitan Pizza Making. *arXiv preprint arXiv:2411.18438* (2024).

Yinheng Li, Shaofei Wang, Han Ding, and Hang Chen. 2023. Large language models in finance: A survey. In *Proceedings of the fourth ACM international conference on AI in finance*. 374–382.

Jonathan Liebers, Patrick Horn, Christian Burschik, Uwe Gruenefeld, and Stefan Schneegass. 2021. Using Gaze Behavior and Head Orientation for Implicit Identification in Virtual Reality. In *Proceedings of the 27th ACM Symposium on Virtual Reality Software and Technology* (Osaka, Japan) *(VRST '21)*. Association for Computing Machinery, New York, NY, USA, Article 22, 9 pages. doi:10.1145/3489849.3489880

Daniel J. Liebling and Sören Preibusch. 2014. Privacy considerations for a pervasive eye tracking world. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication*. ACM. doi:10.1145/2638728.2641688

Yi Liu, Gelei Deng, Yuekang Li, Kailong Wang, Zihao Wang, Xiaofeng Wang, Tianwei Zhang, Yepang Liu, Haoyu Wang, Yan Zheng, et al. 2023. Prompt Injection attack against LLM-integrated Applications. *arXiv preprint arXiv:2306.05499* (2023).

Yupei Liu, Yuqi Jia, Runpeng Geng, Jinyuan Jia, and Neil Zhenqiang Gong. 2024. Formalizing and benchmarking prompt injection attacks and defenses. In *33rd USENIX Security Symposium (USENIX Security 24)*. 1831–1847.

Virmarie Maquiling, Sean Anthony Byrne, Diederick C Niehorster, Marco Carminati, and Enkelejda Kasneci. 2024a. Zero-Shot Pupil Segmentation with SAM 2: A Case Study of Over 14 Million Images. *arXiv preprint arXiv:2410.08926* (2024).

Virmarie Maquiling, Sean Anthony Byrne, Diederick C Niehorster, Marcus Nyström, and Enkelejda Kasneci. 2024b. Zero-shot segmentation of eye features using the segment anything model (sam). *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 7, 2 (2024), 1–16.

Virmarie Maquiling, Li Zhaoping, and Enkelejda Kasneci. 2024c. Imperceptible Gaze Guidance Through Ocularity in Virtual Reality. *arXiv preprint arXiv:2412.09204* (2024).

Johannes Meyer, Thomas Schlebusch, Wolfgang Fuhl, and Enkelejda Kasneci. 2020. A novel camera-free eye tracking sensor for augmented reality based on laser scanning. *IEEE Sensors Journal* 20, 24 (2020), 15204–15212.

Fatemehsadat Mireshghallah, Kartik Goyal, Archit Uniyal, Taylor Berg-Kirkpatrick, and Reza Shokri. 2022. Quantifying privacy risks of masked language models using membership inference attacks. *arXiv preprint arXiv:2203.03929* (2022).

Fabio Motoki, Valdemar Pinho Neto, and Victor Rodrigues. 2024. More human than human: measuring ChatGPT political bias. *Public Choice* 198, 1 (2024), 3–23.

Moin Nadeem, Anna Bethke, and Siva Reddy. 2020. StereoSet: Measuring stereotypical bias in pretrained language models. *arXiv preprint arXiv:2004.09456* (2020).

Jakub Štěpán Novák, Jan Masner, Petr Benda, Pavel Šimek, and Vojtěch Merunka. 2024. Eye tracking, usability, and user experience: A systematic review. *International Journal of Human–Computer Interaction* 40, 17 (2024), 4484–4500.

Süleyman Özdel, Efe Bozkir, and Enkelejda Kasneci. 2024. Privacy-preserving Scanpath Comparison for Pervasive Eye Tracking. *Proceedings of the ACM on Human-Computer Interaction* 8, ETRA (2024). doi:10.1145/3655605

European Parliament. 2016. General Data Protection Regulation (GDPR). Online, Available: https://gdpr.eu.

Constantinos Patsakis and Nikolaos Lykousas. 2023. Man vs the machine in the struggle for effective text anonymisation in the age of large language models. *Scientific Reports* 13, 1 (2023), 16026.

Mateusz Pomianek, Marek Piszczek, and Marcin Maciejewski. 2021. MEMS mirror based eye tracking: simulation of the system parameter effect on the accuracy of pupil position estimation. *Metrology and Measurement Systems* 28, 4 (2021).

Paul Prasse, David Robert Reich, Silvia Makowski, Lena A Jäger, and Tobias Scheffer. 2022. Fairness in oculomotoric biometric identification. In *2022 Symposium on Eye Tracking Research and Applications*. 1–8.

Keith Rayner. 2009. The 35th Sir Frederick Bartlett Lecture: Eye movements and attention in reading, scene perception, and visual search. *Quarterly journal of experimental psychology* 62, 8 (2009), 1457–1506.

Matthew Sag. 2023. Copyright safety for generative AI. *Hous. L. Rev.* 61 (2023), 295.

Pamela Samuelson. 2023. Generative AI meets copyright. *Science* 381, 6654 (2023), 158–161.

Frederick Shic. 2016. Eye Tracking as a Behavioral Biomarker for Psychiatric Conditions: The Road Ahead. *Journal of the American Academy of Child & Adolescent Psychiatry* 55, 4 (2016), 267–268. doi:10.1016/j.jaac.2016.02.002

Tulika Singh. 2023. AI-Driven Surveillance Technologies and Human Rights: Balancing Security and Privacy. In *International Conference on Smart Systems: Innovations in Computing*. Springer, 703–717.

Julian Steil, Inken Hagestedt, Michael Xuelin Huang, and Andreas Bulling. 2019a. Privacy-aware eye tracking using differential privacy. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications*. ACM. doi:10.1145/3314111.3319915

Julian Steil, Marion Koelle, Wilko Heuten, Susanne Boll, and Andreas Bulling. 2019b. PrivacEye: privacy-preserving head-mounted eye tracking using egocentric scene image and eye movement features. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications* (Denver, Colorado) *(ETRA '19)*. Association for Computing Machinery, New York, NY, USA, Article 26, 10 pages. doi:10.1145/3314111.3319913

Md Samiul Haque Sunny, Md Ishrak Islam Zarif, Ivan Rulik, Javier Sanjuan, Mohammad Habibur Rahman, Sheikh Iqbal Ahamed, Inga Wang, Katie Schultz, and Brahim Brahmi. 2021. Eye-gaze control of a wheelchair mounted 6DOF assistive robot for activities of daily living. *Journal of NeuroEngineering and Rehabilitation* 18 (2021), 1–12.

Ling Tao, Quan Wang, Ding Liu, Jing Wang, Ziqing Zhu, and Li Feng. 2020. Eye tracking metrics to screen and assess cognitive impairment in patients with neurological disorders. *Neurological Sciences* 41 (2020), 1697–1704.

Yasuo Terao, Hideki Fukuda, and Okihide Hikosaka. 2017. What do eye movements tell us about patients with neurological disorders?—An introduction to saccade recording in the clinical setting—. *Proceedings of the Japan Academy, Series B* 93, 10 (2017), 772–801.

Florina Ungureanu, Robert Gabriel Lupu, Adrian Cadar, and Adrian Prodan. 2017. Neuromarketing and visual attention study using eye tracking techniques. In *2017 21st international conference on system theory, control and computing (ICSTCC)*. IEEE, 553–557.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).

Mélodie Vidal, Jayson Turner, Andreas Bulling, and Hans Gellersen. 2012. Wearable eye tracking for mental health monitoring. *Computer Communications* 35, 11 (2012), 1306–1311.

Hanqiu Wang, Zihao Zhan, Haoqi Shan, Siqi Dai, Maximilian Panoff, and Shuo Wang. 2024. GAZEploit: Remote Keystroke Inference Attack by Gaze Estimation from Avatar Views in VR/MR Devices. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security* (Salt Lake City, UT, USA) *(CCS '24)*. Association for Computing Machinery, New York, NY, USA, 1731–1745. doi:10.1145/3658644.3690285

Michel Wedel, Rik Pieters, et al. 2008. Eye tracking for visual marketing. *Foundations and Trends® in Marketing* 1, 4 (2008), 231–320.

Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, Myra Cheng, Mia Glaese, Borja Balle, Atoosa Kasirzadeh, et al. 2021. Ethical and social risks of harm from language models. *arXiv preprint arXiv:2112.04359* (2021).

Alexandra Wolf and Kazuo Ueda. 2021. Contribution of eye-tracking to study cognitive impairments among clinical populations. *Frontiers in psychology* 12 (2021), 590986.

Kun Yan, Zeyu Wang, Lei Ji, Yuntao Wang, Nan Duan, and Shuai Ma. 2024. Voila-A: Aligning Vision-Language Models with User's Gaze Attention. *Advances in Neural Information Processing Systems* 37 (2024), 1890–1918.

Sikai Yang, Gang Yan, and Wan Du. 2024. See Where You Read with Eye Gaze Tracking and Large Language Model. *arXiv preprint arXiv:2409.19454* (2024).

Yifan Yao, Jinhao Duan, Kaidi Xu, Yuanfang Cai, Zhibo Sun, and Yue Zhang. 2024. A survey on large language model (llm) security and privacy: The good, the bad, and the ugly. *High-Confidence Computing* (2024), 100211.

Hong-Jun Yoon, Tandy R Carmichael, and Georgia Tourassi. 2014. Gaze as a biometric. In *Medical Imaging 2014: Image Perception, Observer Performance, and Technology Assessment*, Vol. 9037. SPIE, 39–45.

Johannes Zagermann, Ulrike Pfeil, and Harald Reiterer. 2016. Measuring cognitive load using eye tracking technology in visual computing. In *Proceedings of the sixth workshop on beyond time and errors on novel evaluation methods for visualization*. 78–85.

Peng Zhang and Maged N Kamel Boulos. 2023. Generative AI in medicine and healthcare: promises, opportunities and challenges. *Future Internet* 15, 9 (2023), 286.

Yuhong Zhang, Qin Li, Sujal Nahata, Tasnia Jamal, Shih-Kuen Cheng, Gert Cauwenberghs, and Tzyy-Ping Jung. 2024b. Integrating Large Language Model, EEG, and Eye-Tracking for Word-Level Neural State Classification in Reading Comprehension. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* (2024).

Zhiping Zhang, Michelle Jia, Hao-Ping (Hank) Lee, Bingsheng Yao, Sauvik Das, Ada Lerner, Dakuo Wang, and Tianshi Li. 2024a. "It's a Fair Game", or Is It? Examining How Users Navigate Disclosure Risks and Benefits When Using LLM-Based Conversational Agents. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. ACM. doi:10.1145/3613904.3642385

Xin Zhou, Martin Weyssow, Ratnadira Widyasari, Ting Zhang, Junda He, Yunbo Lyu, Jianming Chang, Beiqi Zhang, Dan Huang, and David Lo. 2025. LessLeak-Bench: A First Investigation of Data Leakage in LLMs Across 83 Software Engineering Benchmarks. *arXiv preprint arXiv:2502.06215* (2025).