

# SoK: The Privacy Paradox of Large Language Models: Advancements, Privacy Risks, and Mitigation

Yashothara Shanmugarasa  
yashothara.shanmugarasa@data61.csiro.au  
CSIRO's Data61  
Sydney, NSW, Australia

M.A.P. Chamikara  
chamikara.arachchige@data61.csiro.au  
CSIRO's Data61  
Melbourne, Victoria, Australia

Ming Ding  
ming.ding@data61.csiro.au  
CSIRO's Data61  
Sydney, NSW, Australia

Thierry Rakotoarivelo  
thierry.rakotoarivelo@data61.csiro.au  
CSIRO's Data61  
Sydney, NSW, Australia

## ABSTRACT

Large language models (LLMs) are sophisticated artificial intelligence systems that enable machines to generate human-like text with remarkable precision. While LLMs offer significant technological progress, their development using vast amounts of user data scraped from the web and collected from extensive user interactions poses risks of sensitive information leakage. Most existing surveys focus on the privacy implications of the training data but tend to overlook privacy risks from user interactions and advanced LLM capabilities. This paper aims to fill that gap by providing a comprehensive analysis of privacy in LLMs, categorizing the challenges into four main areas: (i) privacy issues in LLM training data, (ii) privacy challenges associated with user prompts, (iii) privacy vulnerabilities in LLM-generated outputs, and (iv) privacy challenges involving LLM agents. We evaluate the effectiveness and limitations of existing mitigation mechanisms targeting these proposed privacy challenges and identify areas for further research.

## CCS CONCEPTS

• Security and privacy → Privacy protections;

## KEYWORDS

Large Language Models, Privacy, Systematization of Knowledge

### ACM Reference Format:

Yashothara Shanmugarasa, Ming Ding, M.A.P. Chamikara, and Thierry Rakotoarivelo. 2025. SoK: The Privacy Paradox of Large Language Models: Advancements, Privacy Risks, and Mitigation. In *ACM Asia Conference on Computer and Communications Security (ASIA CCS '25)*, August 25–29, 2025, Hanoi, Vietnam. ACM, New York, NY, USA, 20 pages. <https://doi.org/10.1145/3708821.3733888>

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

*ASIA CCS '25, August 25–29, 2025, Hanoi, Vietnam*

© 2025 Association for Computing Machinery.  
ACM ISBN 979-8-4007-1410-8/2025/08...\$15.00  
<https://doi.org/10.1145/3708821.3733888>

## 1 INTRODUCTION

Artificial intelligence (AI) stands as a groundbreaking frontier technology, offering users the ability to simplify day-to-day tasks through automation and intelligent querying. One of the rapidly advancing AI innovations is the Large Language Model (LLM), which has revolutionized natural language processing, empowering machines to produce human-like text with exceptional precision [30].

LLMs, trained on vast datasets and characterized by extensive parameters, master linguistic nuances to generate sensible, coherent, and conversational responses to natural language queries [128]. Notable platforms like OpenAI's ChatGPT and Google's Gemini, boasting millions of active users, have gained immense popularity.

Leveraging prompt engineering [118], in-context learning capabilities [13], and retrieval-augmented generation [61], LLMs have demonstrated adaptability to diverse contexts and can undertake translation, debugging, and storytelling tasks without requiring training or fine-tuning. While they perform impressively in various tasks, there are opportunities for improvement in addressing challenges such as hallucination, the need for domain-specific knowledge, and horizon cut-off.

### 1.1 Privacy Concerns in Conventional AI/Machine Learning (ML) vs. LLMs

Despite LLMs offering significant advancements and benefits, fairness, reliability, bias, security, and privacy concerns are becoming increasingly prominent [87]. This paper focuses on the multifaceted view of privacy in LLMs, which presents distinct challenges compared to traditional AI privacy concerns related to models and data. Privacy concerns in traditional AI focus on data privacy issues, such as unauthorized access to sensitive information or leakage through models via attacks like membership inference and model inversion [104]. In addition to these traditional risks, LLMs bring about new dimensions of privacy risks due to their advanced capabilities, such as a profound understanding of natural language context, human-like text generation, contextual awareness in knowledge-rich fields, and a robust ability to follow instructions [126, 128].

The utilization of LLMs in commercial applications such as OpenAI's chatGPT and Google's Gemini raises strong privacy concerns, particularly regarding the collection and use of personal data [20]. LLMs are often trained on extensive datasets compiled from various sources, including social media posts, websites, online articles,

and books. There is a significant risk that this data, even when de-identified, can contain sensitive personal information [128]. Despite the sharing of such data by individuals or organizations for diverse agreed purposes, it is plausible that such data have been incorporated into the training data for building LLM without their explicit consent, constituting a privacy breach. Besides, the interactive nature of LLMs poses additional privacy challenges. As users engage with LLMs through prompts, they may inadvertently disclose sensitive information. The sophisticated reasoning capabilities of LLMs can infer sensitive information about users even utilizing seemingly harmless user inputs [128]. Additionally, LLMs often integrate functionalities that involve communication with external agents or third-party modules, extending the ecosystem to perform automated tasks assigned by users. However, such integration can further complicate privacy concerns, as each interaction has the potential to generate and disseminate sensitive data across a network of interconnected systems, each with varying levels of security and adherence to privacy concerns.

## 1.2 Motivation

Evaluating privacy concerns in LLM systems is crucial for several compelling reasons:

- *Regulatory compliance.* Regulations such as the General Data Protection Regulation (GDPR) [89], the Health Insurance Portability and Accountability Act (HIPAA) [1], or the EU AI Act [90] impose laws on data protection and AI safety, establish substantial penalties for non-compliance and empower individuals with rights to protect their sensitive information. Thus, organizations must adhere to such regulations when deploying LLM-based services.
- *Inherent vulnerabilities and privacy risks in LLMs.* Compared to traditional AI, LLMs introduce exacerbated vulnerabilities, as they are often trained on vast amounts of potentially sensitive data. This requires comprehensive analyses and methodologies to understand their strengths and weaknesses, and mitigate their vulnerabilities.
- *Advancement of privacy technologies for LLMs.* Understanding the privacy implications of LLMs facilitates evaluating current technologies, such as differential privacy within the context of LLMs. Such insights are crucial for pinpointing gaps and driving innovation in developing robust mechanisms for privacy-preserving LLMs.
- *Public perception and Trust.* The widespread use of LLMs in consumer applications, particularly in critical sectors such as finance and healthcare, demands transparency and robust privacy protection to maintain public trust, as privacy violations within these domains can lead to significant financial and individual harms.

Hence, a thorough understanding of LLM systems' privacy challenges and the development of practical solutions are essential for their ethical and safe deployment.

## 1.3 The Overall Objectives

The objectives of this paper are three-fold. First, it comprehensively maps the landscape of privacy issues associated with LLMs, analyzing these multiple dimensions: the privacy implications of the

training data used, the potential for sensitive information leakage through user interactions (i.e. prompts), and privacy vulnerabilities specific to the deployment of LLMs (i.e., privacy breaches through LLM-generated output and privacy challenges involving LLM agents). Second, it provides a detailed overview of the current mitigation strategies and technologies employed to address these identified privacy risks. Finally, it broadens the discussion by examining how various privacy-preserving techniques can tackle diverse privacy threats, while highlighting key research challenges and proposing avenues for future exploration.

To support these goals, we conducted a comprehensive review across various sources, including the latest research papers, tutorials, dissertations, and magazines focused on LLMs and privacy. We hope this work provides a holistic perspective for researchers, practitioners, and stakeholders engaged in the development and deployment of LLM systems.

## 1.4 State-of-the-art

The state-of-the-art studies on privacy in LLMs can be broadly categorized into privacy and security challenges in LLM models and training data. Privacy in LLMs can be categorized further into two main groups: privacy risks and defense mechanisms. The former focuses on specific privacy and security challenges within the LLM domain [14, 67, 70, 125], while the latter focuses on addressing these challenges and the available defense mechanisms [20, 55, 80, 120, 128]. While many surveys [80, 106, 128] predominantly focus on the privacy implications of LLMs and their training data, they often overlook the distinct privacy risks introduced by user interactions and LLMs' advanced capabilities [14]. Consequently, we could not identify any previous surveys or SoKs related to the privacy challenges during LLM deployment and user interaction.

The significance of prompts cannot be overstated in LLM operation, as they facilitate the customization of pre-trained LLMs for task-specific purposes by appending a sequence of query texts [68, 149]. However, these prompts can often contain user-sensitive information, which can be inferred by LLM service providers, who leverage the LLMs' capabilities and access to personal data [108]. Recent advancements in the privacy domain for LLMs offer a timely and highly relevant overview of this emerging research area [31].

Our study delves into the latest approaches and techniques in this domain, highlighting current research gaps and privacy challenges in LLMs across the categories: training data, prompts, LLM outputs, and LLM agents.

## 1.5 Research Questions

We formulated the following research questions (RQs) to address the aforementioned objectives of our study systematically.

- (1) RQ1: How do LLMs' advanced inferring capabilities and contextual awareness impact user privacy across various dimensions, such as inference of sensitive information and memorization of user inputs? This question aims to explore the core capabilities of LLMs that may lead to privacy breaches.
- (2) RQ2: What privacy challenges are associated with users actively interacting with LLMs? This question investigates how users can inadvertently expose sensitive information when interacting with LLMs.

- (3) RQ3: What state-of-the-art solutions are currently employed to mitigate the privacy risks inherent in LLMs? This question investigates the effectiveness and limitations of the existing privacy-preservation technologies for LLMs.
- (4) RQ4: What are the open challenges and possible future trends to overcome these challenges? This question aims to highlight ongoing challenges and predict future technologies to preserve privacy in the context of LLMs.

## 1.6 Scope

This paper focuses on the papers published after 2022, aligning with the surge of LLMs. While we acknowledge the significant body of work exploring training data privacy concerns, we deliberately chose not to delve extensively into them, as they have been comprehensively covered in existing surveys [80, 106, 128]. Instead, our focus lies on conducting a more in-depth analysis of other privacy challenges unique to LLMs that have not been sufficiently explored in the literature. Moreover, our study focuses exclusively on privacy issues within LLMs and does not address the broader security challenges. The employed methodology to identify relevant literature and establish the categories is detailed in Appendix 6.1.

## 2 OVERVIEW OF PRIVACY CHALLENGES IN LLM

Privacy in LLMs refers to protecting sensitive data that extends beyond traditional Personally Identifiable Information (PII) to encompass confidential, proprietary, intellectual property, and contextual or behavioral data that could reveal personal attributes or identities [52, 124]. This section outlines the key privacy challenges associated with LLM systems. Our comprehensive analysis of the literature reveals four main categories of privacy challenges: (i) privacy issues related to LLM training data [36, 80], (ii) privacy challenges in the interaction with LLM systems via user prompts in the interaction with LLM systems [68, 146], (iii) privacy vulnerabilities in LLM-generated outputs [86, 124], and (iv) privacy challenges involving LLM agents [39, 109]. Figure 1 shows the multi-faceted view of four identified privacy challenges in LLM systems.

### 2.1 Privacy in LLM Training Data

Privacy concerns in LLMs primarily revolve around protecting the training data utilized in their development. This is a key issue in commercial deployments of LLM systems, where they are trained using personal data retrieved from various sources such as social media posts and websites [55]. This practice raises significant privacy concerns as it often involves using individuals' data without consent and contextual integrity principles when used outside their intended context [81]. Moreover, concerns arise regarding data storage practices, with LLM service providers such as OpenAI storing personal information potentially conflicting with GDPR [74]. Besides the privacy concern arising from the utilization of public data and the data owner's consent, LLM models are susceptible to significant privacy vulnerabilities, such as memorization and the inadvertent leakage of PII or confidential data [20].

### 2.2 Privacy in Prompts

There are potential privacy leakages from user-provided input prompts during interactions with LLM systems [107]. The ever-increasing inference capabilities of LLM systems and their widespread adoption across various domains have raised significant privacy concerns through user prompts. User prompts are often exposed to the service providers, raising questions about whether current LLM systems could breach users' privacy by inferring personal attributes from the prompts. Despite data transmission and storage encryption, LLM service providers retain the knowledge of the actual data contents, undermining individuals' or entities' trust in these services.

The interactive utilization of LLM systems introduces a new set of privacy risks during inference time. LLM systems are fed with diverse types of information from various sources in their prompts, potentially revealing more contextual data beyond the direct sensitive data in the prompts. Recent findings in [107] have demonstrated that modern LLM systems can be leveraged for highly accurate predictions of personal attributes, even from seemingly innocuous data. Previously, human involvement was necessary for inferring private attributes. However, these models are now swift and economical enough to automate such inferences at scale [108]. Simultaneously, users remain unaware of these concerns, inadvertently sharing texts containing information easily inferable by LLM systems. For evidence of privacy vulnerability via prompts, The New York Times<sup>1</sup> reported instances where personal information, including chatbot conversations and login credentials, was unintentionally exposed.

Given these challenges, we believe the research community is responsible for fostering a paradigm shift in LLM-centered privacy research. New research should extend beyond examining privacy risks associated only with models and data to explore privacy vulnerabilities inherent in other aspects (e.g., user interactions) of LLMs. Prompts are one significant aspect as they directly reveal substantial information. These approaches should prioritize user-friendliness, ensuring users can easily and quickly verify whether they disclose private information in their prompts.

### 2.3 Privacy in LLM-generated Outputs

Within the scope of privacy concerns regarding LLM-generated outputs, we delve into issues such as the retention, extraction, and retrieval of sensitive user data in outputs (i.e., memorizing sensitive information from user prompts) and the inadvertent inclusion of sensitive information in the LLM outputs.

We consider LLM-generated output privacy as a distinct aspect of LLM privacy due to the following reasons: 1) Users may enter sensitive information in prompts without expecting it to appear in the output. However, LLMs may still include such information in their responses. This issue is prevalent in In-context Learning (ICL), where users' private data in prompts help adapt the model to specific tasks with only black-box access. Due to the context-specific examples in the prompts, the output can include labels of sensitive data samples from the prompt data for ICL. 2) In commercialized

<sup>1</sup><https://www.nytimes.com/2023/03/31/technology/chatgpt-italy-ban.html>

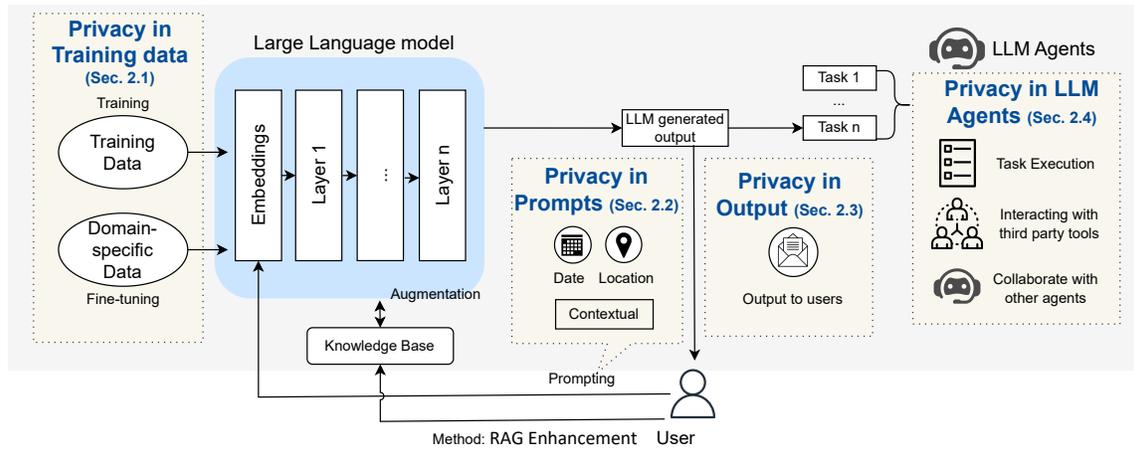


Figure 1: A Multi-Faceted View of the Four Identified Privacy Challenges in LLMs

LLM settings like the GPT store<sup>2</sup>, specialized LLM models are fine-tuned using private data through methods like ICL, fine-tuning, and Retrieval-Augmented Generation (RAG). These techniques integrate user queries with demonstrations or relevant documents from a knowledge base to enhance LLM-generated responses. However, this customization poses a risk of exposing private information from a smaller user base to a broader audience. Although LLM products should ideally avoid generating harmful outputs such as sensitive information, legal, medical / health, financial advice, and misinformation, they can be customized to produce such output. 3) Even when users employ techniques to protect sensitive information in input prompts, LLM-generated outputs remain vulnerable to exposure to malicious service providers, third-party tools, external parties, or hackers[129]. Service providers can potentially disclose sensitive information in output prompts by analyzing training data or accessing external sources, even with partially sensitive data from prompts. For example, using protection mechanisms at the user end, users can restrict LLM systems from retaining sensitive data (e.g., passport identifier) in input prompts. However, LLM systems may still inadvertently reveal sensitive information such as the birth year or age in the output due to its processing [120].

### 2.4 Privacy in LLM Agents

Recent advances in LLM Systems have led to the development of agent-based solutions such as WebGPT [78], AutoGPT [94], and GPT Plugins (e.g., WebPilot [145]). As illustrated in Figure 2, in these systems, a main LLM agent divides a user prompt into individual tasks and transmits them to secondary agents (which can also be LLM-based). These agents then leverage a range of powerful tools, including other agents and third-party applications (e.g., mobile apps, web browsers, sensors, metaverse interfaces, code interpreters, and API plugins) to interact with external environments and carry out real-world tasks assigned by the main LLM agent.

The transition from text-based interactions to those involving sensitive data with agents operating in the real world raises significant privacy concerns [99]. For example, physical interaction-based

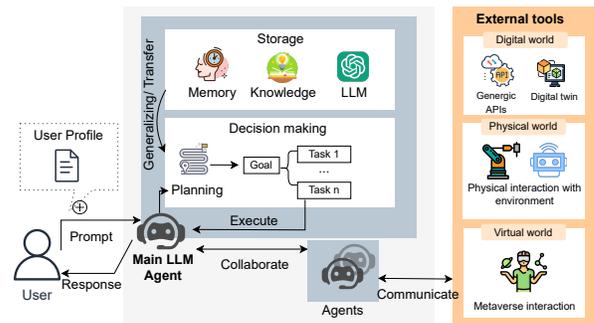


Figure 2: The Workflow of LLM Agents

LLM agents heighten privacy risks by collecting real-world data such as facial photos (cameras) and speech audio (microphones) [32]. A significant challenge in such an agent-based system is the risk of unauthorized propagation of sensitive data across different components of the system [67]. Indeed, agents autonomously interact with external tools, different data modalities, other AI models, and APIs, making tracking and controlling data flow difficult. Users may provide private inputs through prompts, and when processing their requests, an LLM agent may share their data with external tools or across interconnected systems without their explicit awareness. The agent’s ability to autonomously reason, execute multi-step actions, and retrieve information from various sources further amplifies the privacy risks, as data may persist across interactions in ways not immediately apparent to users. LLM agents pose risks of data persistence, where sensitive information from past interactions may resurface in future prompts, especially when contextual memory is retained across sessions or users [67].

Hence, LLM agents extend beyond traditional LLM output privacy concerns by persisting memory, unintended exposure across sessions and contexts, autonomously interacting with external tools, and operating in real-world physical interactions with environments [32]. Therefore, stricter and more diverse mitigation strategies are needed to address privacy concerns in LLM agents.

<sup>2</sup><https://openai.com/index/introducing-the-gpt-store/>

### 3 PRIVACY CHALLENGES IN LLM AND MITIGATION

This section discusses privacy challenges under four broad categories and reviews existing mitigation techniques in the literature.

#### 3.1 Privacy Issues in LLM Training Data and Mitigation

We analyze privacy issues in LLM training data by distinguishing between the causes, attack mechanisms, and consequences of privacy vulnerabilities. Data memorization is one of the primary causes of privacy risks, where models inadvertently retain and reproduce sensitive information from training data. Privacy attacks act as mechanisms that adversaries use to exploit LLMs to extract sensitive information. These attacks can lead to consequences such as privacy leakage. For example, unauthorized access to personal or confidential data can compromise privacy. Figure 3 illustrates these interrelated aspects of privacy risks in LLMs. Although prior research works [20, 80, 106] extensively discuss privacy threats and mitigation, they have not examined this causal structure. This section summarises key insights, emphasizing how data memorization contributes to privacy vulnerabilities and how various attack techniques facilitate privacy leakages.

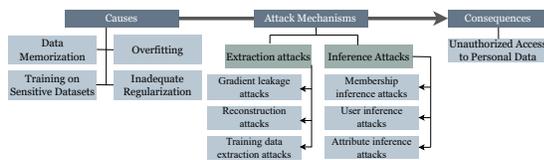


Figure 3: The Privacy Challenges in LLM Training Data

**3.1.1 Potential Causes of Privacy Issues.** Many privacy issues in LLMs arise due to data memorization, which can be exploited by privacy attacks, ultimately resulting in privacy leakages [53]. Data memorization is an intrinsic characteristic of many ML models, as the training process involves retaining specific information from input data to make accurate predictions [3]. However, while users expect LLMs to generate novel content that is semantically similar to the training data, the verbatim memorization and reproduction of learned phrases can potentially expose sensitive data and breach privacy [36, 80]. Hence, memorization is particularly problematic when LLMs are trained on personal information datasets.

Beyond data memorization, privacy leakage can stem from training on sensitive datasets without proper sanitization [14]. Overfitting and inadequate regularization (e.g., dropout, weight decay) further exacerbate the risk by making models more likely to retain and expose specific training data [54, 112].

**Mitigation.** A fundamental approach to mitigating privacy issues in LLM training data is to reduce data memorization, which is the basis for many attacks. We first examine techniques designed to address the challenge of data memorization. One straightforward approach is to deduplicate training datasets, as data redundancy can exacerbate model memorization. Even a 10-fold data duplication can lead to a 1000-fold increase in memorization [51]. Another method is early detection of memorization during the training

phase, enabling practitioners to take corrective actions such as discarding memorized points, reverting to a checkpoint, or halting the training process for adjustments [80]. The utilization of filtering techniques such as the bloom filter method (which is a probabilistic data structure used to detect and filter out memorized or sensitive data efficiently) [45] is another defense mechanism for memorization, scanning the training datasets to check if the model’s next token forms an  $n$ -gram in the training set and selecting an alternate token by sampling from the model’s posterior if it does. *However, this method does not guarantee privacy and can be bypassed with plausible, minimally modified prompts, failing to prevent training data leakage entirely.* Another promising approach for mitigating memorization challenges is differential privacy (DP: detailed in Table 1), which adds noise to data during training, providing mathematical guarantees for privacy but reducing utility [150]. LLM editing offers yet another approach, in which neurons corresponding to memorization and storage of specific training data knowledge can be directly edited (i.e., by altering internal parameters) [15].

A key consequence of data memorization is personal or confidential data leakage, which poses a significant privacy risk. While general techniques such as deduplication [51, 60] can help mitigate model memorization and reduce the risk of personal data leakage, some studies specifically target this issue by employing advanced approaches, such as fuzzy logic deduplication [13] to prevent the retention of sensitive information. Other methods to prevent privacy leakage include data cleaning, PII scrubbing, and filtering with restrictive terms of use [4, 124]. Data cleaning enhances privacy by correcting errors, implementing anonymization, and following secure practices to protect sensitive information [85, 124]. PII scrubbing filters [4, 23, 110] use Named Entity Recognition (NER) to tag and remove PII. However, these tools cannot guarantee complete removal, and AI advancements can infer PII from non-PII data. Therefore, minimizing data retention and enforcing purging policies are critical to reducing the risk of breaches and unauthorized access. *While most of these approaches are applied to smaller text corpora, the use of extensive training data in LLMs presents challenges. More empirical studies are needed to evaluate their effectiveness in such large-scale models with larger datasets.*

Recently, knowledge unlearning techniques [33, 130, 140] have been utilized in LLMs to force models to forget specific knowledge without requiring full retraining. These knowledge unlearning techniques involve randomly sampling data from the training corpus and performing gradient ascent, altering the direction during language modeling on target token sequences. However, the success of this method depends heavily on specific target data and the domain of the data to be unlearned [106].

Another recent method for addressing PII leakage in LLMs is ProPILE [56], a tool that allows data subjects to assess the inclusion and potential leakage of personal data in LLM systems during deployment. ProPILE considers the linkability and structurability of data in inferring PII. Singh et al. [105] proposed a multi-faceted whispered tuning approach by integrating PII redaction, DP, output filtering, and architectural improvements to enhance privacy preservation. Likewise, OpenAI utilizes filtering and fuzzy deduplication techniques to remove PII from the corpora utilized for model training [13]. *Despite these efforts, achieving complete prevention of*

*privacy leaks remains challenging. Hence, in-depth studies are necessary to design effective defense techniques for PII leakages on LLMs while maximizing user education and involvement.*

**3.1.2 Potential Privacy Attacks.** To exploit data memorization in LLM training data, various *privacy attack mechanisms* have been developed, each employing different techniques and targeting specific vulnerabilities. While LLMs utilize Deep Learning (DL) for model building, all traditional attacks on DL models may not directly apply to LLMs due to limited access to model parameters and the prevalent use of application programming interfaces (APIs) in most LLMs. However, certain LLM service providers offer open-source models (e.g., LLaMA, BERT, DeepSeek, Mistral 7B), making DL-related attacks extended to the LLMs. Privacy attacks on LLM can be broadly categorized into two classes: extraction attacks and inference attacks.

**Extraction Attacks.** Extraction attacks involve adversaries attempting to extract sensitive information or insights from the LLM model by accessing model gradients, the training data, or adversarial prompting. LLMs may inadvertently capture and reproduce sensitive information in the training data, potentially raising privacy concerns during the text generation process [80]. Various types of data extraction attacks include gradient leakage attacks [9, 22, 152], reconstruction attacks [8, 19, 69], and training data extraction attacks [14, 79, 127, 147]. Gradient leakage attacks exploit DL optimization algorithms, accessing gradients to leak sensitive data. However, this attack is challenging when the model architecture is unknown (i.e., closed-source) [22]. Reconstruction attacks aim to recover private training data by analyzing model parameters, gradients, or generated output. Even without direct access to LLMs, attackers can use shadow modeling with external datasets and feedback from the target LLM [143]. Training data extraction attacks target sensitive information in training data by strategically querying LLM systems via prompts to retrieve specific examples, including personal information [14, 20]. Jailbreaking [117] and prompt injection [103] security attacks make LLMs more vulnerable to extraction attacks and introduce substantial privacy risks with broad and potentially severe implications even without direct model access. Jailbreaking attacks exploit both direct and indirect techniques to bypass the safety and alignment constraints of LLMs. These attacks often use manually or automatically crafted adversarial prompt templates designed to deceive the model into generating harmful content or revealing sensitive information, potentially including memorized data from its training corpus [70]. In terms of extracting private information, a multi-step jailbreaking approach has bypassed ChatGPT’s safety mechanisms, successfully extracting PII from ChatGPT and LLM-powered search engines [62]. Another study tested 3,700 templated jailbreak prompts, analyzing their effectiveness across different LLMs [88]. Additionally, an automated framework, ReNeLLM, was proposed to generate effective jailbreak prompts using prompt rewriting and scenario nesting to infer private information [25].

**Inference Attacks.** Inference attacks aim to acquire knowledge or insights about a model or data characteristics by observing the model’s responses or behavior [128]. Common inference attacks

in LLMs include membership inference attacks (MIA), attribute inference attacks, and user inference attacks [46, 75, 116].

In MIAs, an adversary attempts to predict whether a particular data point (e.g., a particular sentence or document) is a member of a target model’s training dataset [29] (Detailed in Table 1). MIA can be conducted solely based on the text generated by LLM systems. While MIAs generally perform no better than random guessing for LLM training data due to the vast datasets used [29], their effectiveness improves when small datasets are used for fine-tuning for specific tasks, known as ‘User Inference Attacks’ [50]. These attacks are prevalent in fine-tuned LLMs using techniques such as full fine-tuning, RAG, and ICL techniques [28, 37].

User inference attacks present a significant privacy risk for *LLMs fine-tuned on a smaller group of user data*, where an attacker tries to determine if a specific user participated in the fine-tuning process. Minimal user samples and black-box access to the fine-tuned model are sufficient for this attack to succeed [50]. This poses a privacy risk if the fine-tuning task reveals sensitive information about users, such as a model fine-tuned exclusively on users with a rare disease.

Attribute inference attacks (Detailed in Table 1) pose another privacy risk for LLM systems, as adversaries can use these attacks to identify users’ attributes through prompts using publicly available partial data of users [107]. Robin et al. [107] demonstrated that personal details can be inferred accurately from current LLMs. While prompt injection is often discussed in the literature as a backdoor security attack, it can also be used for attribute inference.

**Mitigation.** Differential Privacy (DP) techniques have become a popular approach in addressing such privacy attacks [66, 102, 105, 137]. Besides DP, techniques to reduce memorization (discussed in Section 3.1.1) can also help prevent these privacy attacks. For example, Jagielski et al. [47] discovered that using larger datasets during fine-tuning or more extended training with non-sensitive data could effectively mitigate privacy concerns by reducing the MIA. Furthermore, test-time defense and instruction processing are also crucial for mitigating privacy attacks [128]. Test-time defenses filter malicious inputs, detect abnormal queries, and post-process LLM outputs. Instruction pre-processing transforms user prompts to eliminate adversarial contexts or malicious intents [64, 96]. Enhancing LLM architectures can also improve safety by integrating LLMs with knowledge graphs [136] and cognitive architectures [97]. Furthermore, the LLM training process can employ robust optimization methods such as adversarial training [134] and robust fine-tuning [26] to prevent malicious text attacks.

*However, these techniques have limitations, such as challenges in effectively handling adversarial inputs, ensuring scalability in large models, and maintaining a balance between privacy and model performance, particularly in dynamic, real-world scenarios.*

Federated learning (FL) may potentially solve privacy attacks in LLM systems by shifting processing from central servers to users [95, 115]. However, its application to LLMs faces challenges due to inadequate FL framework support, handling vast data and complex models, and optimizing communication and computational resources. Some studies [34, 131] have explored FL for LLMs, but its viability remains uncertain. Federated LLM fine-tuning may enable parameter-efficient fine-tuning with limited resources [58].

Fully Homomorphic Encryption (FHE) has been explored as a privacy-preserving solution [92, 98, 142] for language models such as DistilBERT [101] (66M parameters) and BERT (110M parameters) [24]. While FHE has been applied to smaller models with limited datasets [92, 98, 142], scaling it to larger models, such as GPT with massive datasets, remains challenging due to the computational complexity. Outsourcing computation can also be securely conducted for small language models (LM): for instance, an encrypted RNN text classifier was shown to predict on homomorphically encrypted input without accuracy loss [63], and such overhead remains manageable at this scale. With a compact model, full on-device inference becomes realistic, retaining user data locally and mitigating cloud-related vulnerabilities [71, 122, 132]. Conversely, LLMs with billions and trillions of parameters often demand cloud-based infrastructures and massive datasets, introducing more complex privacy concerns.

Mitigation strategies for jailbreaking attacks are extensively studied in the literature and can be categorized into prompt-level and model-level approaches, depending on whether they modify the protected LLM. Prompt-level defenses (e.g., input sanitization) [148] focus on filtering adversarial prompts using rule-based detection or classifier-based approaches. Model-level defenses (e.g., adversarial training) [48] involve enhancing LLMs through safety training, reinforcement learning with human feedback (RLHF), and adversarial fine-tuning to improve robustness against attacks. Constitutional AI frameworks [6] and automated moderation systems also help enforce ethical guidelines and restrict unintended model behaviors. To protect private information from jailbreaking, prior work has explored mostly prompt-level mitigation: defending with additional prompts and using a harmfulness classifier to filter malicious inputs [25, 88]. The study in [62] outlined potential approaches at both model-level defenses and prompt intention detection mechanisms to prevent the disclosure of private information.

#### Key Takeaways

LLM systems have shown an ability to memorize training data, raising concerns about the unintentional disclosure of sensitive information and privacy breaches. Despite the efforts to address these issues with methods such as differential privacy, knowledge unlearning, and test-time defenses, adversaries may still infer private data through adversarial prompting. Researchers should focus on establishing responsible practices that safeguard privacy while preserving the utility of LLMs trained on large-scale datasets.

### 3.2 Privacy Challenges in the Interaction with LLM systems via Prompts and Mitigation

This section focuses on privacy challenges when users interact with pre-built LLM systems via prompts. We identify three types of privacy challenges associated with prompts in the literature: the direct leakage of sensitive data in prompts, the potential inference of sensitive information, and the leakage of contextual information from user devices and activity logs.

**3.2.1 The Direct Leakage of Sensitive Data in Prompts.** Users might inadvertently reveal personal data through prompts when interacting with the system, often due to a lack of awareness or reluctance to undertake additional efforts to minimize personal information in prompts. Even individuals with strong technical backgrounds, often in IT industries, reportedly leaked sensitive company data through ChatGPT prompts in a late 2022 incident<sup>3</sup>.

This challenge became more prominent with the integration of LLMs into interactive computing systems such as conversational agents, making them widely accessible for everyday tasks. Users tend to disclose more private information when engaging with LLM-based conversational agents due to the tools' high utility and human-like interactions [146]. Moreover, many users are unconcerned about revealing personal data in prompts, maybe under the misconception that LLM companies no longer collect interaction data, based on various published statements [120]. However, despite the prevailing belief, often the company's privacy policies indicate that user prompts are periodically integrated into the training process by default [120, 146]. The extent of sensitive data revealed in user prompts is evidenced in the study by Zhang et al. [146]. To identify PII leakage in prompts, they analyzed real-world ChatGPT conversations. The study uncovered significant amounts of sensitive and personal data in LLM prompts, providing a comprehensive overview of user disclosure behaviors.

Protecting personal data in prompts is imperative. Conversations between users and LLMs can become embedded in LLM parameters during training, making them vulnerable to adversarial attacks. Hence, data protection measures should be implemented at the user end to protect data from service providers and external parties.

**Mitigation.** A straightforward mitigation technique for potential direct leakage of sensitive data is input validation and sanitization. Some approaches rely on LLM service providers to implement these solutions, so their efficacy depends on the trustworthiness of the providers or the LLM itself. Other methods acknowledge that privacy vulnerabilities can arise from both LLM service providers and external parties, offering a safer solution than those relying on the trustworthiness of the service providers. One sanitization approach involves using NER or predefined policies and rules to identify and remove specific sensitive details in prompts. Research works [17, 65, 68] and commercial products [41, 73] focus on prompt privacy using this method.

To address the dilemma of potentially compromising prompt utility by removing all personal information from user prompts, EmojiCrypt [68] converts sensitive information in prompts into an encrypted format using emojis and mathematical operations. This method retains informative value for LLMs while making the data incomprehensible to humans. However, it relies on trusting the LLM for encryption, mainly aiming to conceal information from other users. Other constraints include a limited symbolic vocabulary and the risk of introducing inaccuracies.

Another lightweight anonymization technique proposed in [17] aims to protect prompt privacy through substitution or masking. The approach involves two core techniques: hiding private entities in prompts using generative schemes and NER models and seeking private entities for de-anonymization. Both models are

<sup>3</sup><https://tinyurl.com/yny4mjdh>

implemented on the user’s system. After anonymizing the text with the Hide model, the result is sent to the LLM, and the Seek model de-anonymizes the output. This approach commendably does not trust the LLM. *However, it relies on a small corpus or NER for identifying sensitive data, which limits its ability to conceal other words and sentence structures. Additionally, replacing entity words in prompts can challenge tasks that rely on precise semantics.*

Another potential approach involves designing secure prompt templates and sensitive data redaction tools to minimize personal data leakage [73]. This could include using structured data for inputs instead of free text. Variables can be safely inserted into the prompt template without giving users direct control over the prompt structure, allowing for easy removal of sensitive data.

Another technique is a combination of small local LMs and remote LLMs. Hartmann et al. [43] proposed a privacy-preserving mechanism where a local small LM at the user’s end personalizes outputs without disclosing confidential information to the remote model. They explored three algorithms: (1) generating a high-level description of the query locally and using the remote LLM for few-shot examples, (2) creating a similar novel problem with new unlabeled examples for the remote LLM, and (3) maintaining the query’s structure while substituting private information with placeholders. *These methods may be domain-specific, rule, or NER-dependent, requiring users to have sufficient knowledge to manage them.*

TextObfuscator [151] protects user privacy using text obfuscation techniques. It learns private representations that obscure original words while retaining their functionality. This is achieved by identifying prototypes for each word and clustering functionally similar words around the same prototype. Random perturbations are then applied to these clusters, obscuring original words and maintaining functionality. *However, this approach requires significant computational power on local machines, only protects word privacy during inference.*

Recently, Ruoyan et al. [100] integrated FHE and provable security theory with parameter-efficient fine-tuning to propose a secure inference scheme for LLMs, ensuring the protection of both user-side inputs and server-side private parameters. Likewise, [139] designs efficient FHE counterparts for the core algorithmic building blocks of prompt ensembling. In [16, 21, 42], the authors explored privacy-preserving inference using transformer models, incorporating FHE and secure multi-party computation to protect the input text in user prompts. *However, computing and communication efficiency remain a challenge when handling multiple prompts.*

**3.2.2 The Potential Inference of Sensitive Information.** Aside from direct leaks of sensitive data, the capabilities of LLMs enable even seemingly insignificant data to unveil sensitive information through inference [57]. Leveraging their advanced capabilities, LLMs can infer various personal attributes from prompts with the help of extensive sets of unstructured training data. For instance, from the prompt, “I always get stuck there waiting for a hook turn,” LLMs could infer that the individual is in Melbourne, as hook turns are a distinct traffic maneuver primarily employed there [65]. This inference ability, combined with widespread LLM availability, reduces the costs of private data inferences, enabling adversaries to scale beyond limitations imposed by costly human profilers. Existing rule-based or NER-based redaction tools often prove inadequate in

protecting users from such sensitive inference, as they struggle to detect revealing, context-dependent cues.

**Mitigation.** Most existing works utilize LLMs themselves as a mitigation technique to identify sensitive information that can be inferred from prompts. A study in [107] shows the ability of pre-trained LLMs to infer personal attributes from the text given at inference time. Their extensive experiments demonstrated LLMs’ capacity to infer personal attributes from real-world data, even when the text is anonymized using commercial tools. While this paper highlighted the potential for sensitive data leakage from prompts, it did not delve into mitigation techniques.

Building on the work in [107], a later study [108] introduced a framework for anonymizing texts using an adversarial feedback-guided approach. This method leverages LLMs’ strong attribute inference abilities to guide a separate anonymizer LLM. The process involves two steps: an ‘LLM adversary’ performs private attribute inference, and an ‘anonymizing LLM’ adjusts the text to obscure or generalize the inferred cues. *However, this method doesn’t fully address where the LLM adversary should be located (user or service provider) or the utility of the anonymized text. Using smaller, fine-tuned LLMs locally could improve this process by identifying inference data before it reaches the remote LLM, and incorporating LLM bias and prior knowledge may enhance inference detection.*

**3.2.3 The leakage of contextual information.** LLMs often rely on contextual information to generate personalized and impactful content, leveraging their in-context learning capabilities. The contextual information may include users’ personal data, such as app usage data (e.g., shopping lists, logs, and calendar events). Most users include contextual data in prompts for improved performance, as it is more parameter and data-efficient than fine-tuning [28].

Given the granular nature of contextual data, privacy risks naturally arise. It’s essential to understand these risks, mainly because the data used in prompts often comes from smaller, private datasets, unlike the large public corpora used for pretraining LLMs. This small set of private data is more vulnerable to inference attacks. Therefore, striking a balance between performance and privacy is crucial.

**Mitigation.** In the literature, techniques for protecting contextual data in prompts include hashing operations [133], prompt ensembling [28], and employing a personalized LM on the local with contextual information to improve the remote LLM responses[141].

Yim et al. [133] used a hash operation on context data before sending it to the LLM, ensuring only hashed values are transmitted. The LLM’s response also includes hashed values that are then reverted to their original form. This method protects personal data privacy while enabling personalized responses. Duan et al. [28] employed a prompt ensembling method to mitigate MIA on prompted data with contextual information. This method reduces the MIA success rate to near-random guessing levels. The prompt ensembling technique aggregates prediction probability vectors over multiple independent prompted models into an ensemble prediction. In CoGenesis [141], smaller, personalized LMs on user devices access private data and activity logs, while advanced general LLMs operate in the cloud, receiving only general instructions and providing high-level knowledge. This allows for collaborative content generation, where

the user-end model uses context and responses from the LLMs to produce personalized outputs.

*Existing techniques face key challenges. Hashing methods [133] rely on trusting the LLM and only protect against eavesdropping. Prompt ensembling [28] lacks formal privacy guarantees and adds computational overhead. CoGenesis [141] depends on logit access, limiting its use with closed-source LLMs, and tested on synthetic data, raising questions about real-world applicability.*

#### Key Takeaways

User prompts submitted to LLM systems may contain sensitive information. Given the advanced inference capabilities of modern LLMs, even seemingly benign or contextual data can lead to the disclosure of private details, making traditional NER or rule-based detection methods inadequate. As mitigation, researchers are investigating the use of small, fine-tuned LLMs with ICL/RAG techniques at the user's end to assess the risk of sensitive information inference and enhance prompt privacy.

### 3.3 Privacy Vulnerabilities in LLM-generated Outputs and Mitigation

Privacy vulnerabilities in LLM outputs arise when sensitive information, whether learned from training data or newly produced in response, is inadvertently disclosed to unintended recipients. This can happen if the model memorizes personal or confidential details in a generated text or synthesizes proprietary content (such as internal strategies, algorithms, or trade secrets) accessible beyond authorized channels. In either case, allowing these outputs to be publicly visible or logged by external systems exposes private data to potential misuse.

**3.3.1 Revealing Sensitive Information in LLM-generated Output.** As reported in many studies [65, 86, 146], LLMs can memorize sensitive data from prompts and include them in their output despite users explicitly requesting them not to memorize their data in prompts. Priyanshu et al. [86] evaluated chatbot responses' compliance with privacy regulations, particularly when user data are used as few-shot samples for ICL. They found that ChatGPT reproduces PII accurately 57.4% of the time, decreasing to 30.5% with regulation prompts and 15.2% with explicit removal prompts. This behavior can be considered a privacy violation as it occurs without user consent. Moreover, there's the potential for additional privacy breaches if these outputs are compromised through attacks, eavesdropping, or retraining the LLM model with user data.

This risk is heightened when the model is fine-tuned with private data for domain-specific applications, such as when organizations use LLMs for specialized internal tasks. Even within a controlled setting, sensitive data can be exposed to all users with access permissions. For example, an LLM-driven HR chatbot fine-tuned with employee details could inadvertently allow employees to view sensitive information about their peers' salaries and benefits<sup>4</sup>. Without proper controls to restrict access, the model could expose this information to all users.

<sup>4</sup><https://medium.com/snowflake/handling-sensitive-data-with-llms-aa765f8ce840>

This risk increases when LLM-generated outputs are shared with third-party services, especially in integrated systems, such as chat platforms or APIs. Unlike LLM agents, which autonomously interact with external tools and the environment, LLM output privacy concerns stem from the uncontrolled dissemination of generated responses to third-party services without the user's explicit awareness, increasing the risk of unintentional data leakage. Ensuring control and transparency in accessing privacy data with user consent is crucial, especially since the output can potentially be sent to numerous third-party applications once it leaves the LLM. Hence, it is essential to implement techniques to manage the data workflow and keep users informed about it.

**Mitigation.** Mitigation techniques for privacy challenges in LLM systems' outputs are still nascent, with only a few research studies addressing this issue. Given the lack of user control over the black box settings in LLM services, controlling the LLM systems' decisions is challenging. The only feasible method for users to influence the model output is solely through prompts. Otherwise, service providers are responsible for managing the output.

Yao et al. [129] proposed an initial approach to protect LLM decisions in a black-box manner. They define decision privacy and investigate instance obfuscation strategies for decision privacy. They tackled the privacy concerns of decision-making by appending an 'obfuscator' (another random text) to the original prompt. The method uses obfuscators to alter LLM decision distribution, preventing adversaries from inferring the correct decision while the data owner resolves it from obfuscated inputs.

Likewise, to protect the privacy of decisions made in ICL, the studies in [113, 119] concentrated on differentially private aggregation methods to prevent the direct extraction of private data. Wu et al. [119] introduce the DP-ICL framework, which can aggregate and release responses without heavily relying on any individual examples provided in the demonstrated private data. The fundamental concept behind the DP-ICL framework involves producing differentially private responses by aggregating noisy consensus from an ensemble of LLMs' outputs, each operating on separate sets of examples.

*In summary, current mitigation techniques for protecting private information in LLM outputs illustrate a clear trade-off between privacy and utility. Specifically, the approach in [129], designed mainly for static tasks (e.g., text classification), does not adequately address the generative nature of LLM outputs, risking exposure of confidential data through dynamic outputs. Moreover, cross-session data retention poses an additional threat to data privacy: if previous conversation states, memory buffers, or hidden user context are not carefully managed or sanitized, sensitive information can reappear in later responses, bypassing any initial privacy safeguards. Consequently, while these methods offer valuable insights, they require further refinement to incorporate robust safeguards against unintended data exposure.*

### Key Takeaways

Recent techniques such as ICL and RAG have exacerbated privacy concerns with LLM-generated outputs, as users increasingly provide sensitive data for domain-specific tasks. The opaque nature of black-box LLM services limits user control, making it difficult to manage or audit model-generated decisions externally. To mitigate these risks, users may obfuscate their input before submission. Additionally, LLM providers may consider implementing techniques like DP to reduce potential leakage in outputs.

## 3.4 Privacy Challenges while Involving Agents in the LLM System Tasks and Mitigation

This section addresses privacy challenges that arise when users interact with LLM agents to accomplish tasks that communicate with the external world. These privacy challenges in LLM agents can be categorized into three areas: automated task execution via LLM agents, adversarial interactions of agents, and the potential exposure of sensitive information to third-party tools via agents.

**3.4.1 Privacy Issues Caused by Automated Task Execution via LLM Agents.** Human instructions via prompts often contain ambiguities or omit crucial details. It is imperative to ensure the resilience and reliability of the agents' decisions, guaranteeing alignment between the actions performed by the agent and their intended tasks. Given that unknown risks lie in complex environments and user instructions, LLM agents are prone to causing unexpected privacy and safety issues [99, 121]. For instance, a LLM agent tasked to process emails might click on phishing links, leading to potential privacy breaches and financial loss. Similarly, if instructed to send an email with file content, the agent might accidentally include sensitive details like credit card information. In another example, [35] demonstrated that LLM agents could even exploit one-day vulnerabilities to hack websites based on task descriptions.

However, due to their long-tail nature, identifying the risks associated with LLM agents is challenging. These risks arise as agents interact with various tools to execute tasks, coupled with the open-ended nature of potential issues and the substantial engineering effort required for testing such interactions.

**Mitigation.** Several approaches assess whether user instructions are safe to execute without compromising user privacy and safety when evaluating LLMs' risk awareness regarding agent privacy. ToolEmu [99] implemented a GPT-4 powered emulator with diverse tools and scenarios to identify potential failure modes in LLM agents and create sandbox states to trigger such failures. It also includes a GPT-4 powered safety evaluator to quantify risks. Similarly, AgentMonitor [77] proposed using an LLM to monitor and halt unsafe actions, preventing potential safety and privacy issues on the open internet. Both ToolEmu and AgentMonitor use LLMs to identify risky actions in agents. However, these methods have limitations: They often ignore core constraints, particularly in complex multi-turn interactions (back-and-forth interactions between LLM agents, users, and diverse environments), and rely on humans to provide risk descriptions.

To evaluate LLMs' risk awareness in agent safety through complex multi-turn interactions, Yuan et al. [135] developed R-Judge, a benchmark dataset. The dataset includes user instructions, agent action histories, and environment feedback annotated with safety labels. In evaluating automated harmful action detection, most LLMs struggled to identify safety and privacy risks, with GPT-4 achieving an F1 score of 72.52% versus the human score of 89.07%.

Hua et al. [44] presented TrustAgent, an agent framework employing three strategies to ensure safety: pre-planning (injecting safety knowledge before plan generation), in-planning (bolstering safety during plan generation), and post-planning (ensuring safety through post-planning inspection). Experiments show that TrustAgent enhances both safety and helpfulness. However, the study underscores the need for inherent reasoning abilities within LLMs to support truly safe agents.

Aside from the methods discussed that rely on external supervision from humans or other LLMs, which require significant investment in human labor and computational resources, alternative methods can operate without human supervision. Self-alignment [111] is an emerging paradigm where LLMs can independently achieve value alignment. Pang et al. [84] explored this direction for achieving self-alignment of LLMs through multi-agent role-playing. This method allows the LLM to create a simulation environment that mirrors real-world multi-party interactions and simulates the social consequences of a user's instruction. It records the textual interactions of the roles and summarizes them as the final output. Although the simulation process can be time-consuming during inference, fine-tuning can mitigate this issue.

*In summary, while these methods aim to prevent LLM agents from disclosing sensitive or personal data, they come with notable trade-offs and limitations tied directly to data privacy. Supervision-based approaches (ToolEmu, AgentMonitor, and TrustAgent) are effective at flagging and mitigating privacy risks but depend on human oversight. This reliance introduces its own privacy challenges (e.g., potential exposure of private information to supervisors) and can be resource-intensive. R-Judge seeks to automate the detection of privacy-related risks, yet it struggles with nuanced or context-dependent leaks, highlighting that even advanced models like GPT-4 still fall short of human-level detection in complex data-privacy scenarios. Meanwhile, self-alignment techniques attempt to reduce reliance on external supervision by enabling the model itself to learn protective behaviors around sensitive data. However, these methods are computationally expensive and have yet to prove robust across diverse real-world privacy contexts (where inadvertent disclosure can occur in unpredictable ways). Moving forward, designing scalable, adaptive frameworks with minimal human involvement remains critical: such approaches must specifically address how to protect sensitive information while ensuring that any oversight mechanisms themselves do not become new vectors for privacy compromise.*

**3.4.2 Adversarial Interactions of Agents.** Privacy risks can be further exacerbated by adversarial agents that mimic honest behavior and request additional private information under the guise of needing it for prompt execution tasks. Since LLM agents have demonstrated remarkable capabilities in generating human-like text, users may unwittingly disclose more information than they would to other service-providing systems [146]. In [107], the authors explore

the interaction between LLM agents and users, where agents gather additional information through human-like conversations.

LLM agents present complex privacy concerns due to their generative power and vulnerability to adversarial manipulation. Techniques such as ICL and RAG can be exploited to train agents for malicious purposes with minimal input. [39] found that modifying a single agent can rapidly lead to widespread harmful behaviors among agents without further intervention from the adversary. The nature of LLM agents and advanced techniques increases the risk of adversarial attacks and interactions. For evidence, [114] introduced Evil Geniuses, a virtual team that develops malevolent strategies and conducts Red-Blue exercises, demonstrating successful harmful actions through adversarial interactions. Similarly, [109] explored the potential for LLMs to act as adversaries by perturbing text samples to bypass safety measures. They investigated whether LLMs can inherently create adversarial examples from benign samples to deceive existing safeguards. Their findings underscore significant challenges for (semi-)autonomous systems using LLM agents, particularly concerning potential adversarial behavior.

**Mitigation.** Chern et al. [18] proposed a multi-agent debate mechanism to improve quality and mitigate adversarial behavior. In this approach, agents self-evaluate through discussion and feedback. They tested this method with state-of-the-art models, evaluating susceptibility to red team attacks in single- and multi-agent settings. The results show that multi-agent debate generally produces less toxic responses to adversarial prompts without introducing new risks at inference time. However, the approach is resource-intensive, requires multiple queries, relies on knowledge from a single LLM, and does not address removing toxic outputs effectively.

Though multi-agent settings can mitigate this issue, multiple agents can have secret collusion between themselves. Motwani et al. [76] formalized the issue of secret collusion in generative AI systems, proposing mitigation measures based on AI and security principles. They investigated monitoring the information content in agents' communications, focusing on agents' capabilities and incentives to evade detection. Using monitoring agents, they detected steganographic techniques that agents use to hide information by analyzing cover text anomalies and agent simulations. Future work could extend this by examining complexity and information theory, optimizing pressures, and exploring complex multi-agent settings.

Common alignment methods to reduce privacy risks include fine-tuning [40, 138] and reward modeling [59]. During fine-tuning, models are trained on specific datasets, incorporating human-generated examples to align with desired behaviors. Reward modeling involves optimizing a reward function to reflect desired outcomes, often using reinforcement learning techniques to adjust the model's behavior accordingly. Although intelligent personal agents should minimize user interruptions, integrating user opinions or human assistance can prove valuable when making significant decisions.

*In summary, while multi-agent debate and alignment strategies enhance LLM privacy, they come with trade-offs. Debate-based methods reduce adversarial risks but are resource-intensive and knowledge-limited. Secret collusion raises concerns about hidden adversarial behavior, requiring costly monitoring. Fine-tuning and reward modeling enhance alignment but may not generalize well.*

**3.4.3 Potential Exposure of Sensitive Information to Third-party Tools.** The next challenge is preventing the unnecessary and unauthorized disclosure of user-sensitive information during interactions with agents to third-party tools and the external world [67]. The LLM agents orchestrate the task and invoke relevant third-party tools to execute it with the collected private information. However, in this process, the LLM possesses knowledge of the user's private input and autonomously initiates queries in places without user awareness or inadvertently shares this information with the other agents or external tools. Another threat model is adversarial third-party applications that can manipulate the context of interaction to trick LLM-based agents into exposing private information irrelevant to the task [5].

**Mitigation.** The study in [144] examined privacy concerns for tool-using LLM agents. They proposed PrivacyAsst using encryption-based and shuffling-based solutions to preserve privacy. The encryption-based solution allows tasks to operate on encrypted inputs. In contrast, the shuffling-based solution uses attribute-based forgery generative models and an attribute shuffling mechanism to create privacy-preserving requests. However, limitations include the detectability of dummy prompts by advanced LLMs, the complexity and computational cost of encryption, reliance on third-party tools, and the burden of creating multiple dummy inputs on users.

AirGapAgent [5] focuses on preventing unintended data leakage during third-party interactions. It restricts access to only the necessary data for specific tasks, considering user privacy preferences. The design involves two LLMs: one minimizes data for sharing appropriately, while the other handles third-party interactions with the minimized data. This ensures the agent can distinguish between private and non-private data for each task.

*PrivacyAsst and AirGapAgent enhance privacy but face challenges: PrivacyAsst incurs high computational costs and third-party reliance, while AirGapAgent's effectiveness depends on accurate filtering. Both highlight the trade-off between privacy and usability, requiring more scalable, adaptive solutions.*

#### Key Takeaways

LLM agents' automated interactions with the physical and digital worlds amplified their privacy vulnerabilities. Existing mitigation approaches often rely on human-guided LLM emulators to enforce user-specific privacy preferences and prevent sensitive information leakage. Emerging approaches such as self-alignment seek to reduce human oversight by enabling agents to self-assess privacy risks through multi-agent role-playing.

Overall, Figure 4 provides an intricate portrayal of privacy concerns, encapsulating the four broader privacy challenges and the tailored mitigation technologies deployed at various phases of LLM systems, including data pre-training, training, and post-training approaches. A more detailed summary of these approaches can be found in Appendix: Table 2.

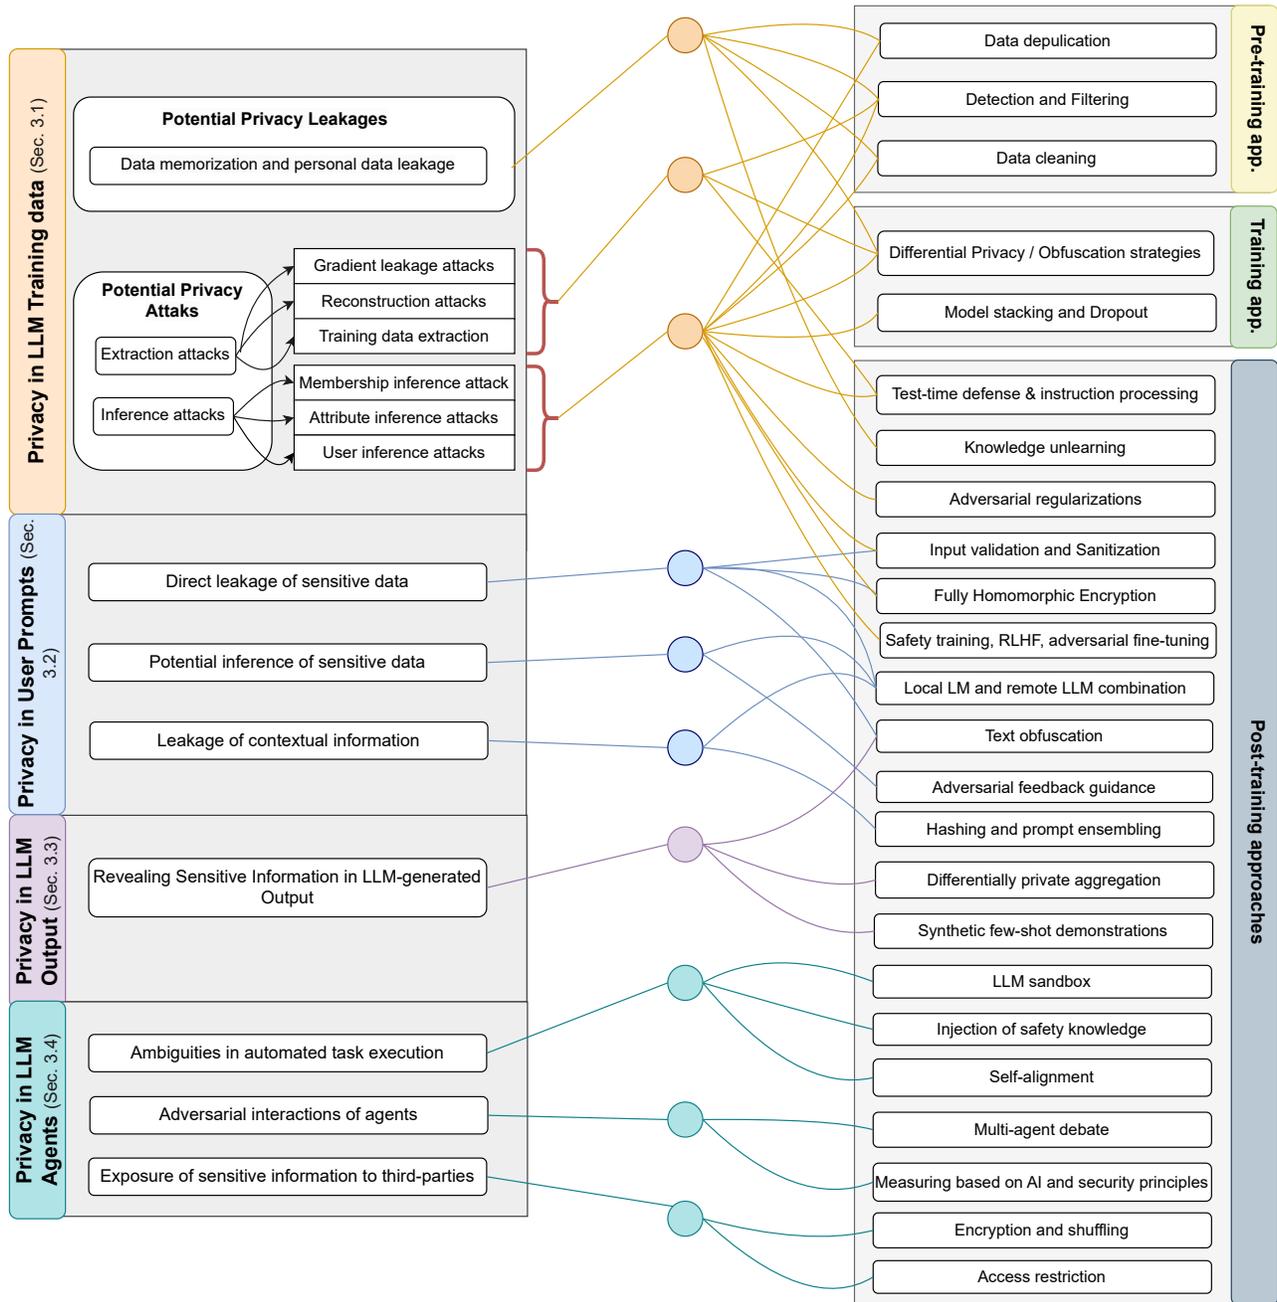


Figure 4: Summary of Privacy Challenges in LLM Systems and Mitigation Techniques in the Literature

#### 4 DISCUSSION AND FUTURE DIRECTIONS

This section highlights how some privacy-preserving methods could address multiple privacy threats and discusses their research challenges and potential future directions.

NER is widely used for direct leakage prevention in both training data and user prompts. Similarly, encryption and DP-based

mechanisms can mitigate multiple privacy risks with trade-offs in processing and model performance. Adversarial regularization, which enhances the robustness of LLM, also plays a role in mitigating privacy attacks. However, these solutions are often complex due to fine-tuning between privacy preservation and model efficiency.

Furthermore, many solutions operate in isolation, addressing individual threats rather than working together as part of a unified and adaptive privacy framework.

Privacy concerns in LLM systems usually stem from foundational models and extend to various downstream applications. Traditional methods for protecting LLM training data fall short against more advanced attacks, especially when dealing with fine-tuned LLMs and their various ways of user interaction. These models are particularly susceptible to inference attacks due to their complexity and natural language interaction. Existing mitigation techniques—such as detection, filtering, and model stacking—often come with trade-offs such as reduced performance, incomplete filtering, and higher latency and computational costs. To enhance LLM privacy, future approaches should explore hybrid models that integrate adaptive filtering, modular architectures, and human-in-the-loop validation. Reducing training data influence through prompt conditioning or modular architectures could provide additional privacy gains without sacrificing accuracy, particularly in fine-tuned models. Developing domain-specific privacy frameworks tailored for LLMs will also be crucial in aligning with emerging regulations. Future research should also focus on developing adaptive privacy mechanisms that dynamically assess conversation context and adjust LLM responses accordingly, ensuring that privacy-sensitive topics trigger enhanced protections such as real-time content filtering, obfuscation, or user alerts.

Other privacy-preserving approaches for generic models, such as cryptographic techniques, knowledge unlearning, and Federated Learning, show opportunities and challenges for the LLM domain. Cryptographic methods, while theoretically effective, are computationally expensive when applied to LLMs due to their large-scale processing requirements. Knowledge unlearning struggles with completely removing sensitive information when dealing with deeply integrated training data [106]. Adapting such techniques for LLMs requires careful optimization to balance privacy, efficiency, and performance.

Another critical direction is the advancement of policy-driven AI governance. Efforts such as the EU AI Act [90], AI Safety Institutes [2], and Anthropic’s Constitutional AI [7] highlight the need for clearer guidelines on data collection, storage, and use in LLM systems. Automated auditing mechanisms and adversarial testing can further enhance transparency and accountability in LLM deployment, while continuous monitoring, adversarial testing, and proactive threat modeling could mitigate other privacy risks.

## 5 CONCLUSION

Our SoK identifies four main categories of privacy challenges: (i) privacy issues in LLM training data, (ii) privacy challenges associated with user prompts, (iii) privacy risks in LLM-generated outputs, and (iv) privacy vulnerabilities involving LLM agents. While the literature has extensively examined privacy issues in training data, privacy issues arising from user interactions and LLM outputs, particularly involving small and private datasets through fine-tuning, ICL, RAG, and agents, remain underexplored. Addressing these overlooked areas requires further research to strengthen privacy safeguards across user prompts, LLM outputs, and agent-based applications.

## REFERENCES

- [1] Accountability Act. 1996. Health insurance portability and accountability act of 1996. *Public law* 104 (1996), 191.
- [2] UK AI Safety Institute. 2024. *Inspect AI: Framework for Large Language Model Evaluations*. UK AI Safety Institute. [https://github.com/UKGovernmentBEIS/inspect\\_ai](https://github.com/UKGovernmentBEIS/inspect_ai)
- [3] Devansh Arpit, Stanislaw Jastrzundzinski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S. Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, and Simon Lacoste-Julien. 2017. A closer look at memorization in deep networks. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70* (Sydney, NSW, Australia) (ICML '17). JMLR.org, 233–242.
- [4] Shubhi Asthana, Ruchi Mahindru, Bing Zhang, and Jorge Sanz. 2025. Adaptive PII Mitigation Framework for Large Language Models. *arXiv preprint arXiv:2501.12465* (2025).
- [5] Eugene Bagdasaryan, Ren Yi, Sahra Ghalebikesabi, Peter Kairouz, Marco Gruteser, Sewoong Oh, Borja Balle, and Daniel Ramage. 2024. Air Gap: Protecting Privacy-Conscious Conversational Agents. *arXiv preprint arXiv:2405.05175* (2024).
- [6] Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. 2022. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073* (2022).
- [7] Yuntao Bai, Saurav Kadavath, Sandipan Kundu, and Amanda Askell et al. 2022. Constitutional AI: Harmlessness from AI Feedback. *arXiv:2212.08073 [cs.CL]* <https://arxiv.org/abs/2212.08073>
- [8] Borja Balle, Giovanni Cherubin, and Jamie Hayes. 2022. Reconstructing training data with informed adversaries. In *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1138–1156.
- [9] Mislav Balunovic, Dimitar Dimitrov, Nikola Jovanović, and Martin Vechev. 2022. Lamp: Extracting text from gradients with language model priors. *Advances in Neural Information Processing Systems* 35 (2022), 7641–7654.
- [10] Dipto Barman, Ziyi Guo, and Owen Conlan. 2024. The dark side of language models: Exploring the potential of llms in multimedia disinformation generation and dissemination. *Machine Learning with Applications* (2024), 100545.
- [11] Kristian González Barman, Nathan Wood, and Pawel Pawłowski. 2024. Beyond transparency and explainability: on the need for adequate and contextualized user guidelines for LLM use. *Ethics and Information Technology* 26, 3 (2024), 47.
- [12] Rouzbeh Behnia, Mohammadreza Reza Ebrahimi, Jason Pacheco, and Balaji Padmanabhan. 2022. Ew-tune: A framework for privately fine-tuning large language models with differential privacy. In *2022 IEEE International Conference on Data Mining Workshops (ICDMW)*. IEEE, 560–566.
- [13] 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.
- [14] Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, et al. 2021. Extracting training data from large language models. In *30th USENIX security symposium (USENIX Security 21)*. 2633–2650.
- [15] Ting-Yun Chang, Jesse Thomason, and Robin Jia. 2023. Do Localization Methods Actually Localize Memorized Data in LLMs? *arXiv preprint arXiv:2311.09060* (2023).
- [16] Tianyu Chen, Hangbo Bao, Shaohan Huang, Li Dong, Binxing Jiao, Daxin Jiang, Haoyi Zhou, Jianxin Li, and Furu Wei. 2022. The-x: Privacy-preserving transformer inference with homomorphic encryption. *arXiv preprint arXiv:2206.00216* (2022).
- [17] Yu Chen, Tingxin Li, Huiming Liu, and Yang Yu. 2023. Hide and Seek (HaS): A Lightweight Framework for Prompt Privacy Protection. *arXiv:2309.03057 [cs.CR]*
- [18] Steffi Chern, Zhen Fan, and Andy Liu. 2024. Combating Adversarial Attacks with Multi-Agent Debate. *arXiv preprint arXiv:2401.05998* (2024).
- [19] Junjie Chu, Zeyang Sha, Michael Backes, and Yang Zhang. 2024. Conversation Reconstruction Attack Against GPT Models. *arXiv preprint arXiv:2402.02987* (2024).
- [20] Badhan Chandra Das, M Hadi Amini, and Yanzhao Wu. 2024. Security and privacy challenges of large language models: A survey. *arXiv preprint arXiv:2402.00888* (2024).
- [21] Leo de Castro, Antigoni Polychroniadou, and Daniel Escudero. 2024. Privacy-Preserving Large Language Model Inference via GPU-Accelerated Fully Homomorphic Encryption. In *NeurIPS Safe Generative AI Workshop 2024*.
- [22] Jieren Deng, Yijue Wang, Ji Li, Chao Shang, Hang Liu, Sanguthevar Rajasekaran, and Caiwen Ding. 2021. Tag: Gradient attack on transformer-based language models. *arXiv preprint arXiv:2103.06819* (2021).
- [23] Sentry Developers. [n. d.]. PII and Data Scrubbing — develop.sentry.dev. <https://develop.sentry.dev/pii>. [Accessed 24-04-2024].

- [24] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv:1810.04805* [cs.CL] <https://arxiv.org/abs/1810.04805>
- [25] Peng Ding, Jun Kuang, Dan Ma, Xuezhi Cao, Yunsen Xian, Jiajun Chen, and Shujian Huang. 2023. A Wolf in Sheep's Clothing: Generalized Nested Jailbreak Prompts can Fool Large Language Models Easily. *arXiv preprint arXiv:2311.08268* (2023).
- [26] Xinshuai Dong, Anh Tuan Luu, Min Lin, Shuicheng Yan, and Hanwang Zhang. 2021. How should pre-trained language models be fine-tuned towards adversarial robustness? *Advances in Neural Information Processing Systems* 34 (2021), 4356–4369.
- [27] Wenyu Du, Tongxu Luo, Zihan Qiu, Zeyu Huang, Yikang Shen, Reynold Cheng, Yike Guo, and Jie Fu. 2024. Stacking Your Transformers: A Closer Look at Model Growth for Efficient LLM Pre-Training. *arXiv preprint arXiv:2405.15319* (2024).
- [28] Haonan Duan, Adam Dziedzić, Mohammad Yaghini, Nicolas Papernot, and Franziska Boenisch. 2023. On the privacy risk of in-context learning. In *The 61st Annual Meeting Of The Association For Computational Linguistics*.
- [29] Michael Duan, Anshuman Suri, Niloofar Mireshghallah, Sewon Min, Weijia Shi, Luke Zettlemoyer, Yulia Tsvetkov, Yejin Choi, David Evans, and Hannaneh Hajishirzi. 2024. Do Membership Inference Attacks Work on Large Language Models? *arXiv preprint arXiv:2402.07841* (2024).
- [30] Travis Dyde. 2023. Documentation on the emergence, current iterations, and possible future of Artificial Intelligence with a focus on Large Language Models. (2023).
- [31] Kennedy Edemacu and Xintao Wu. 2024. Privacy Preserving Prompt Engineering: A Survey. *arXiv preprint arXiv:2404.06001* (2024).
- [32] L. Ekenstam. 2023. Linux on x: "if you think this is creepy... you should know what all major stores in the world knows about you as a shopper they use in-store, online, cell-tower, and more to keep track of everything... everything. cambridge analytics is a joke in comparison". <https://x.com/LinusEkenstam/status/1692602911518343502>.
- [33] Ronen Eldan and Mark Russinovich. 2023. Who's Harry Potter? Approximate Unlearning in LLMs. *arXiv:2310.02238* [cs.CL]
- [34] Tao Fan, Yan Kang, Guoqiang Ma, Weijing Chen, Wenbin Wei, Lixin Fan, and Qiang Yang. 2023. Fate-llm: A industrial grade federated learning framework for large language models. *arXiv preprint arXiv:2310.10049* (2023).
- [35] Richard Fang, Rohan Bindu, Akul Gupta, and Daniel Kang. 2024. LLM Agents can Autonomously Exploit One-day Vulnerabilities. *arXiv preprint arXiv:2404.08144* (2024).
- [36] Vitaly Feldman. 2020. Does learning require memorization? a short tale about a long tail. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, 954–959.
- [37] Wenjie Fu, Huandong Wang, Chen Gao, Guanghua Liu, Yong Li, and Tao Jiang. 2023. Practical membership inference attacks against fine-tuned large language models via self-prompt calibration. *arXiv preprint arXiv:2311.06062* (2023).
- [38] Neil Zhenqiang Gong and Bin Liu. 2018. Attribute inference attacks in online social networks. *ACM Transactions on Privacy and Security (TOPS)* 21, 1 (2018), 1–30.
- [39] Xiangming Gu, Xiaosen Zheng, Tianyu Pang, Chao Du, Qian Liu, Ye Wang, Jing Jiang, and Min Lin. 2024. Agent Smith: A Single Image Can Jailbreak One Million Multimodal LLM Agents Exponentially Fast. *arXiv preprint arXiv:2402.08567* (2024).
- [40] Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A Smith. 2020. Don't stop pretraining: Adapt language models to domains and tasks. *arXiv preprint arXiv:2004.10964* (2020).
- [41] David Haber. 2024. Introducing Lakera Guard – Bringing Enterprise-Grade Security to LLMs with One Line of Code | Lakera – Protecting AI teams that disrupt the world. – lakera.ai. <https://www.lakera.ai/blog/lakera-guard-overview>. [Accessed 15-05-2024].
- [42] Meng Hao, Hongwei Li, Hanxiao Chen, Pengzhi Xing, Guowen Xu, and Tianwei Zhang. 2022. Iron: Private inference on transformers. *Advances in neural information processing systems* 35 (2022), 15718–15731.
- [43] Florian Hartmann, Duc-Hieu Tran, Peter Kairouz, Victor Cărbune, et al. 2024. Can LLMs get help from other LLMs without revealing private information? *arXiv preprint arXiv:2404.01041* (2024).
- [44] Wenyue Hua, Xianjun Yang, Zelong Li, Cheng Wei, and Yongfeng Zhang. 2024. TrustAgent: Towards Safe and Trustworthy LLM-based Agents through Agent Constitution. *arXiv preprint arXiv:2402.01586* (2024).
- [45] Daphne Ippolito, Florian Tramèr, Milad Nasr, Chiyuan Zhang, Matthew Jagielski, Katherine Lee, Christopher A Choquette-Choo, and Nicholas Carlini. 2022. Preventing verbatim memorization in language models gives a false sense of privacy. *arXiv preprint arXiv:2210.17546* (2022).
- [46] Abhyuday Jagannatha, Bhanu Pratap Singh Rawat, and Hong Yu. 2021. Membership inference attack susceptibility of clinical language models. *arXiv preprint arXiv:2104.08305* (2021).
- [47] Matthew Jagielski, Om Thakkar, Florian Tramèr, Daphne Ippolito, Katherine Lee, Nicholas Carlini, Eric Wallace, Shuang Song, Abhradeep Thakurta, Nicolas Papernot, et al. 2022. Measuring forgetting of memorized training examples. *arXiv preprint arXiv:2207.00099* (2022).
- [48] Tanqu Jiang, Zian Wang, Jiacheng Liang, Changjiang Li, Yuhui Wang, and Ting Wang. 2025. RobustKV: Defending Large Language Models against Jailbreak Attacks via KV Eviction. In *The Thirteenth International Conference on Learning Representations*. <https://openreview.net/forum?id=L5godAOC2z>
- [49] Junfeng Jiao, Saleh Afroogh, Yiming Xu, and Connor Phillips. 2024. Navigating llm ethics: Advancements, challenges, and future directions. *arXiv preprint arXiv:2406.18841* (2024).
- [50] Nikhil Kandpal, Krishna Pillutla, Alina Oprea, Peter Kairouz, Christopher A Choquette-Choo, and Zheng Xu. 2023. User inference attacks on large language models. *arXiv preprint arXiv:2310.09266* (2023).
- [51] Nikhil Kandpal, Eric Wallace, and Colin Raffel. 2022. Deduplicating training data mitigates privacy risks in language models. In *International Conference on Machine Learning*. PMLR, 10697–10707.
- [52] Antonia Karamolegkou, Jiaang Li, Li Zhou, and Anders Søgaard. 2023. Copyright violations and large language models. *arXiv preprint arXiv:2310.13771* (2023).
- [53] Aly Kassem, Omar Mahmoud, and Sherif Saad. 2023. Preserving Privacy Through Dememorization: An Unlearning Technique For Mitigating Memorization Risks In Language Models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 4360–4379. <https://doi.org/10.18653/v1/2023.emnlp-main.265>
- [54] Yigitcan Kaya, Sanghyun Hong, and Tudor Dumitras. 2020. On the Effectiveness of Regularization Against Membership Inference Attacks. *arXiv:2006.05336* [cs.LG] <https://arxiv.org/abs/2006.05336>
- [55] Sunder Ali Khowaja, Parus Khuwaja, and Kapal Dev. 2023. Chatgpt needs spade (sustainability, privacy, digital divide, and ethics) evaluation: A review. *arXiv preprint arXiv:2305.03123* (2023).
- [56] Siwon Kim, Sangdoon Yun, Hwaran Lee, Martin Gubri, Sungroh Yoon, and Seong Joon Oh. 2024. Propile: Probing privacy leakage in large language models. *Advances in Neural Information Processing Systems* 36 (2024).
- [57] Jacob Leon Kröger, Leon Gellrich, Sebastian Pape, Saba Rebecca Brause, and Stefan Ullrich. 2022. Personal information inference from voice recordings: User awareness and privacy concerns. *Proc. Priv. Enhancing Technol.* 2022, 1 (2022), 6–27.
- [58] Weirui Kuang, Bingchen Qian, Zitao Li, Daoyuan Chen, Dawei Gao, Xuchen Pan, Yuexiang Xie, Yaliang Li, Bolin Ding, and Jingren Zhou. 2023. Federatedscope-llm: A comprehensive package for fine-tuning large language models in federated learning. *arXiv preprint arXiv:2309.00363* (2023).
- [59] Harrison Lee, Samrat Phatale, Hassan Mansoor, Kellie Lu, Thomas Mesnard, Colton Bishop, Victor Carbune, and Abhinav Rastogi. 2023. Rlaif: Scaling reinforcement learning from human feedback with ai feedback. *arXiv preprint arXiv:2309.00267* (2023).
- [60] Katherine Lee, Daphne Ippolito, Andrew Nystrom, Chiyuan Zhang, Douglas Eck, Chris Callison-Burch, and Nicholas Carlini. 2021. Deduplicating training data makes language models better. *arXiv preprint arXiv:2107.06499* (2021).
- [61] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems* 33 (2020), 9459–9474.
- [62] Haoran Li, Dadi Guo, Wei Fan, Mingshi Xu, Jie Huang, Fanpu Meng, and Yangqiu Song. 2023. Multi-step jailbreaking privacy attacks on chatgpt. *arXiv preprint arXiv:2304.05197* (2023).
- [63] Kunhong Li, Ruwei Huang, and Bo Yang. 2025. Privacy-Preserving Text Classification on Deep Neural Network. *Neural Processing Letters* 57, 2 (2025), 29.
- [64] Linyang Li, Demin Song, and Xipeng Qiu. 2022. Text adversarial purification as defense against adversarial attacks. *arXiv preprint arXiv:2203.14207* (2022).
- [65] Tianshi Li, Sauvik Das, Hao-Ping Lee, Dakuo Wang, Bingsheng Yao, and Zhiping Zhang. 2024. Human-Centered Privacy Research in the Age of Large Language Models. *arXiv preprint arXiv:2402.01994* (2024).
- [66] Xuechen Li, Florian Tramèr, Percy Liang, and Tatsunori Hashimoto. 2021. Large language models can be strong differentially private learners. *arXiv preprint arXiv:2110.05679* (2021).
- [67] Yuanchun Li, Hao Wen, Weijun Wang, Xiangyu Li, Yizhen Yuan, Guohong Liu, Jiacheng Liu, Wenxing Xu, Xiang Wang, Yi Sun, et al. 2024. Personal llm agents: Insights and survey about the capability, efficiency and security. *arXiv preprint arXiv:2401.05459* (2024).
- [68] Guo Lin, Wenyue Hua, and Yongfeng Zhang. 2024. PromptCrypt: Prompt Encryption for Secure Communication with Large Language Models. *arXiv preprint arXiv:2402.05868* (2024).
- [69] Tong Liu, Yingjie Zhang, Zhe Zhao, Yinpeng Dong, Guozhu Meng, and Kai Chen. 2024. Making Them Ask and Answer: Jailbreaking Large Language Models in Few Queries via Disguise and Reconstruction. *arXiv preprint arXiv:2402.18104* (2024).
- [70] Yi Liu, Gelei Deng, Zhengzi Xu, Yuekang Li, Yaowen Zheng, Ying Zhang, Lida Zhao, Tianwei Zhang, and Yang Liu. 2023. Jailbreaking chatgpt via prompt engineering: An empirical study. *arXiv preprint arXiv:2305.13860* (2023).

- [71] Zechun Liu, Changsheng Zhao, Forrest Iandola, Chen Lai, Yuandong Tian, Igor Fedorov, Yunyang Xiong, Ernie Chang, Yangyang Shi, Raghuraman Krishnamoorthi, Liangzhen Lai, and Vikas Chandra. 2024. MobileLLM: Optimizing Sub-billion Parameter Language Models for On-Device Use Cases. *arXiv:2402.14905* [cs.LG]. <https://arxiv.org/abs/2402.14905>
- [72] Scott M Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. *Advances in neural information processing systems* 30 (2017).
- [73] Aatish Mandelecha. 2024. How to Secure Sensitive Data in LLM Prompts? – strac.io. <https://www.strac.io/blog/secure-sensitive-data-in-llm-prompts>. [Accessed 14-05-2024].
- [74] E McGowan. 2024. Is ChatGPT's use of people's data even legal? – blog.avast.com. <https://blog.avast.com/chatgpt-data-use-legal>. [Accessed 04-04-2024].
- [75] Fatemehsadat Mirehghallah, Archit Uniyal, Tianhao Wang, David K Evans, and Taylor Berg-Kirkpatrick. 2022. An empirical analysis of memorization in fine-tuned autoregressive language models. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. 1816–1826.
- [76] Sumeet Ramesh Motwani, Mikhail Baranchuk, Martin Strohmeier, Vijay Bolina, Philip HS Torr, Lewis Hammond, and Christian Schroeder de Witt. 2024. Secret Collusion Among Generative AI Agents. *arXiv preprint arXiv:2402.07510* (2024).
- [77] Silen Naihini, David Atkinson, Marc Green, Merwane Hamadi, Craig Swift, Douglas Schonholtz, Adam Tauman Kalai, and David Bau. 2023. Testing language model agents safely in the wild. *arXiv preprint arXiv:2311.10538* (2023).
- [78] Reichihiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. 2021. Webgpt: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332* (2021).
- [79] Milad Nasr, Nicholas Carlini, Jonathan Hayase, Matthew Jagielski, A Feder Cooper, Daphne Ippolito, Christopher A Choquette-Choo, Eric Wallace, Florian Tramèr, and Katherine Lee. 2023. Scalable extraction of training data from (production) language models. *arXiv preprint arXiv:2311.17035* (2023).
- [80] Seth Neel and Peter Chang. 2023. Privacy issues in large language models: A survey. *arXiv preprint arXiv:2312.06717* (2023).
- [81] Helen Nissenbaum. 2004. Privacy as contextual integrity. *Wash. L. Rev.* 79 (2004), 119.
- [82] Chris Olah, Arvind Satyanarayan, Ian Johnson, Shan Carter, Ludwig Schubert, Katherine Ye, and Alexander Mordvintsev. 2018. The building blocks of interpretability. *Distill* 3, 3 (2018), e10.
- [83] Matthew J Page, Joanne E McKenzie, Patrick M Bossuyt, Isabelle Boutron, Tammy C Hoffmann, Cynthia D Mulrow, Larissa Shamseer, Jennifer M Tetzlaff, Elie A Akl, Sue E Brennan, et al. 2021. The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. *bmj* 372 (2021).
- [84] Xianghe Pang, Shuo Tang, Rui Ye, Yuxin Xiong, Bolun Zhang, Yanfeng Wang, and Siheng Chen. 2024. Self-Alignment of Large Language Models via Monopolylogue-based Social Scene Simulation. *arXiv preprint arXiv:2402.05699* (2024).
- [85] Ildikó Pílán, Pierre Lison, Lilja Øvrelid, Anthi Papadopoulou, David Sánchez, and Montserrat Batet. 2022. The text anonymization benchmark (tab): A dedicated corpus and evaluation framework for text anonymization. *Computational Linguistics* 48, 4 (2022), 1053–1101.
- [86] Aman Priyanshu, Supriti Vijay, Ayush Kumar, Rakshit Naidu, and Fatemehsadat Mirehghallah. 2023. Are chatbots ready for privacy-sensitive applications? an investigation into input regurgitation and prompt-induced sanitization. *arXiv preprint arXiv:2305.15008* (2023).
- [87] Md Abdur Rahman, Lamyaa Alqahtani, Amna Alboooq, and Alaa Ainousah. 2024. A Survey on Security and Privacy of Large Multimodal Deep Learning Models: Teaching and Learning Perspective. In *2024 21st Learning and Technology Conference (L&T)*. IEEE, 13–18.
- [88] Abhinav Rao, Sachin Vashistha, Atharva Naik, Somak Aditya, and Monojit Choudhury. 2023. Tricking llms into disobedience: Formalizing, analyzing, and detecting jailbreaks. *arXiv preprint arXiv:2305.14965* (2023).
- [89] Protection Regulation. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council. *Regulation (eu)* 679 (2016), 2016.
- [90] Protection Regulation. 2024. AI Act, Regulation (EU) 2024/1689 of the European Parliament and of the Council. *Regulation (eu)* 1689 (2024).
- [91] Jingjing Ren, Ashwin Rao, Martina Lindorfer, Arnaud Legout, and David Choffnes. 2016. Recon: Revealing and controlling pii leaks in mobile network traffic. In *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services*. 361–374.
- [92] Donghwan Rho, Taeseong Kim, Minje Park, Jung Woo Kim, Hyunsik Chae, Ernest K Ryu, and Jung Hee Cheon. 2024. Encryption-friendly LLM architecture. *arXiv preprint arXiv:2410.02486* (2024).
- [93] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "Why should i trust you?" Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. 1135–1144.
- [94] Toran Bruce Richards. 2023. Auto-gpt: Autonomous artificial intelligence software agent. <https://github.com/Significant-Gravitas/AutoGPT>. [Accessed 08-04-2024].
- [95] Jae Hun Ro, Srinadh Bhojanapalli, Zheng Xu, Yanxiang Zhang, and Ananda Theertha Suresh. 2024. Efficient Language Model Architectures for Differentially Private Federated Learning. *arXiv preprint arXiv:2403.08100* (2024).
- [96] Alexander Robey, Eric Wong, Hamed Hassani, and George J Pappas. 2023. Smoothllm: Defending large language models against jailbreaking attacks. *arXiv preprint arXiv:2310.03684* (2023).
- [97] Oscar J Romero, John Zimmerman, Aaron Steinfeld, and Anthony Tomic. 2023. Synergistic integration of large language models and cognitive architectures for robust ai: An exploratory analysis. In *Proceedings of the AAAI Symposium Series*, Vol. 2. 396–405.
- [98] Lorenzo Rovida and Alberto Leporati. 2024. Transformer-based language models and homomorphic encryption: An intersection with bert-tiny. In *Proceedings of the 10th ACM International Workshop on Security and Privacy Analytics*. 3–13.
- [99] Yangjun Ruan, Honghua Dong, Andrew Wang, Silviu Pitis, Yongchao Zhou, Jimmy Ba, Yann Dubois, Chris J Maddison, and Tatsunori Hashimoto. 2023. Identifying the risks of llm agents with an llm-emulated sandbox. *arXiv preprint arXiv:2309.15817* (2023).
- [100] Zhang Ruoyan, Zheng Zhongxiang, and Bao Wankang. 2025. Practical Secure Inference Algorithm for Fine-tuned Large Language Model Based on Fully Homomorphic Encryption. *arXiv preprint arXiv:2501.01672* (2025).
- [101] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108* (2019).
- [102] Weiyang Shi, Aiqi Cui, Evan Li, Ruoxi Jia, and Zhou Yu. 2021. Selective differential privacy for language modeling. *arXiv preprint arXiv:2108.12944* (2021).
- [103] Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. 2020. Autoprompt: Eliciting knowledge from language models with automatically generated prompts. *arXiv preprint arXiv:2010.15980* (2020).
- [104] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. 2017. Membership inference attacks against machine learning models. In *2017 IEEE symposium on security and privacy (SP)*. IEEE, 3–18.
- [105] Tanmay Singh, Harshvardhan Aditya, Vijay K Madiseti, and Arshdeep Bahga. 2024. Whispered Tuning: Data Privacy Preservation in Fine-Tuning LLMs through Differential Privacy. *Journal of Software Engineering and Applications* 17, 1 (2024), 1–22.
- [106] Victoria Smith, Ali Shahin Shamsabadi, Carolyn Ashurst, and Adrian Weller. 2023. Identifying and mitigating privacy risks stemming from language models: A survey. *arXiv preprint arXiv:2310.01424* (2023).
- [107] Robin Staab, Mark Vero, Mislav Balunovic, and Martin Vechev. 2024. Beyond Memorization: Violating Privacy via Inference with Large Language Models. In *The Twelfth International Conference on Learning Representations*. <https://openreview.net/forum?id=kmn0BhQk7p>
- [108] Robin Staab, Mark Vero, Mislav Balunović, and Martin Vechev. 2024. Large Language Models are Advanced Anonymizers. *arXiv preprint arXiv:2402.13846* (2024).
- [109] Lukas Struppek, Minh Hieu Le, Dominik Hintersdorf, and Kristian Kersting. 2024. Exploring the Adversarial Capabilities of Large Language Models. *arXiv preprint arXiv:2402.09132* (2024).
- [110] Nishant Subramani, Sasha Luccioni, Jesse Dodge, and Margaret Mitchell. 2023. Detecting personal information in training corpora: an analysis. In *Proceedings of the 3rd Workshop on Trustworthy Natural Language Processing (TrustNLP 2023)*. 208–220.
- [111] Zhiqing Sun, Yikang Shen, Qinlong Zhou, Hongxin Zhang, and et al Chen. 2024. Principle-driven self-alignment of language models from scratch with minimal human supervision. *Advances in Neural Information Processing Systems* 36 (2024).
- [112] Jasper Tan, Blake Mason, Hamid Javadi, and Richard G. Baraniuk. 2022. Parameters or Privacy: A Provable Tradeoff Between Overparameterization and Membership Inference. *arXiv:2202.01243* [stat.ML]. <https://arxiv.org/abs/2202.01243>
- [113] Xinyu Tang, Richard Shin, Huseyin A Inan, Andre Manoel, Fatemehsadat Mirehghallah, Zinan Lin, Sivakanth Gopi, Janardhan Kulkarni, and Robert Sim. 2023. Privacy-preserving in-context learning with differentially private few-shot generation. *arXiv preprint arXiv:2309.11765* (2023).
- [114] Yu Tian, Xiao Yang, Jingyuan Zhang, Yinpeng Dong, and Hang Su. 2023. Evil geniuses: Delving into the safety of llm-based agents. *arXiv preprint arXiv:2311.11855* (2023).
- [115] Minh N Vu, Truc Nguyen, Tre'R Jeter, and My T Thai. 2024. Analysis of Privacy Leakage in Federated Large Language Models. *arXiv preprint arXiv:2403.04784* (2024).
- [116] Jeffrey G Wang, Jason Wang, Marvin Li, and Seth Neel. 2024. Pandora's White-Box: Increased Training Data Leakage in Open LLMs. *arXiv preprint arXiv:2402.17012* (2024).
- [117] Alexander Wei, Nika Haghtalab, and Jacob Steinhardt. 2024. Jailbroken: How does llm safety training fail? *Advances in Neural Information Processing Systems* 36 (2024).
- [118] Jules White, Quchen Fu, Sam Hays, Michael Sandborn, Carlos Olea, Henry Gilbert, Ashraf Elnashar, Jesse Spencer-Smith, and Douglas C Schmidt. 2023.

- A prompt pattern catalog to enhance prompt engineering with chatgpt. *arXiv preprint arXiv:2302.11382* (2023).
- [119] Tong Wu, Ashwinee Panda, Jiachen T Wang, and Prateek Mittal. 2023. Privacy-preserving in-context learning for large language models. In *The Twelfth International Conference on Learning Representations*.
- [120] Xiaodong Wu, Ran Duan, and Jianbing Ni. 2023. Unveiling security, privacy, and ethical concerns of chatgpt. *Journal of Information and Intelligence* (2023).
- [121] Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe Wang, Senjie Jin, Enyu Zhou, et al. 2023. The rise and potential of large language model based agents: A survey. *arXiv preprint arXiv:2309.07864* (2023).
- [122] Daliang Xu, Wangsong Yin, Xin Jin, Ying Zhang, Shiyun Wei, Mengwei Xu, and Xuanzhe Liu. 2023. Llmcad: Fast and scalable on-device large language model inference. *arXiv preprint arXiv:2309.04255* (2023).
- [123] Fuzhao Xue, Yao Fu, Wangchunshu Zhou, Zangwei Zheng, and Yang You. 2024. To repeat or not to repeat: Insights from scaling llm under token-crisis. *Advances in Neural Information Processing Systems* 36 (2024).
- [124] Biwei Yan, Kun Li, Minghui Xu, Yueyan Dong, Yue Zhang, Zhaochun Ren, and Xiuzheng Cheng. 2024. On Protecting the Data Privacy of Large Language Models (LLMs): A Survey. *arXiv preprint arXiv:2403.05156* (2024).
- [125] Haomiao Yang, Kunlan Xiang, Mengyu Ge, Hongwei Li, Rongxing Lu, and Shui Yu. 2024. A comprehensive overview of backdoor attacks in large language models within communication networks. *IEEE Network* (2024).
- [126] Linyao Yang, Hongyang Chen, Zhao Li, Xiao Ding, and Xindong Wu. 2024. Give us the facts: Enhancing large language models with knowledge graphs for fact-aware language modeling. *IEEE Transactions on Knowledge and Data Engineering* (2024).
- [127] Zhou Yang, Zhipeng Zhao, Chenyu Wang, Jieke Shi, Dongsun Kim, DongGyun Han, and David Lo. 2023. What do code models memorize? an empirical study on large language models of code. *arXiv preprint arXiv:2308.09932* (2023).
- [128] 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.
- [129] Yixiang Yao, Fei Wang, Srivatsan Ravi, and Muhao Chen. 2024. Privacy-Preserving Language Model Inference with Instance Obfuscation. *arXiv preprint arXiv:2402.08227* (2024).
- [130] Yuanshun Yao, Xiaojun Xu, and Yang Liu. 2023. Large language model unlearning. *arXiv preprint arXiv:2310.10683* (2023).
- [131] Rui Ye, Wenhao Wang, Jingyi Chai, Dihan Li, Zexi Li, Yinda Xu, Yaxin Du, Yanfeng Wang, and Siheng Chen. 2024. OpenFedLLM: Training Large Language Models on Decentralized Private Data via Federated Learning. *arXiv preprint arXiv:2402.06954* (2024).
- [132] Rongjie Yi, Liwei Guo, Shiyun Wei, Ao Zhou, Shangguang Wang, and Mengwei Xu. 2025. EdgeMoE: Empowering Sparse Large Language Models on Mobile Devices. *IEEE Transactions on Mobile Computing* (2025).
- [133] Keun Soo Yim. 2023. Privacy-friendly Personalization of LLM Responses Using Hashed Entity Injection. *Technical Disclosure Commons* (2023).
- [134] Jin Yong Yoo and Yanjun Qi. 2021. Towards improving adversarial training of NLP models. *arXiv preprint arXiv:2109.00544* (2021).
- [135] Tongxin Yuan, Zhiwei He, Lingzhong Dong, Yiming Wang, Ruijie Zhao, Tian Xia, Lizhen Xu, Binglin Zhou, Fangqi Li, Zhuosheng Zhang, et al. 2024. R-Judge: Benchmarking Safety Risk Awareness for LLM Agents. *arXiv preprint arXiv:2401.10019* (2024).
- [136] Ahtsham Zafar, Venkatesh Balavadhani Parthasarathy, Chan Le Van, Saad Shahid, Arsalan Shahid, et al. 2023. Building trust in conversational ai: A comprehensive review and solution architecture for explainable, privacy-aware systems using llms and knowledge graph. *arXiv preprint arXiv:2308.13534* (2023).
- [137] Santiago Zanella-Béguelin, Lukas Wutschitz, Shruti Tople, Victor Rühle, Andrew Paverd, Olga Ohrimenko, Boris Köpf, and Marc Brockschmidt. 2020. Analyzing information leakage of updates to natural language models. In *Proceedings of the 2020 ACM SIGSAC conference on computer and communications security*, 363–375.
- [138] Yifan Zeng, Yiran Wu, Xiao Zhang, Huazheng Wang, and Qingyun Wu. 2024. AutoDefense: Multi-Agent LLM Defense against Jailbreak Attacks. *arXiv preprint arXiv:2403.04783* (2024).
- [139] Jiawen Zhang, Kejia Chen, Zunlei Feng, Jian Lou, Mingli Song, Jian Liu, and Xiaohu Yang. 2024. SecPE: Secure Prompt Ensembling for Private and Robust Large Language Models. In *ECAI*.
- [140] Jinghan Zhang, Junteng Liu, Junxian He, et al. 2023. Composing parameter-efficient modules with arithmetic operation. *Advances in Neural Information Processing Systems* 36 (2023), 12589–12610.
- [141] Kaiyan Zhang, Jianyu Wang, Ermo Hua, Biqing Qi, Ning Ding, and Bowen Zhou. 2024. CoGenesis: A Framework Collaborating Large and Small Language Models for Secure Context-Aware Instruction Following. *arXiv preprint arXiv:2403.03129* (2024).
- [142] Lexin Zhang, Changxiang Li, Qi Hu, Jingjing Lang, Sirui Huang, Linyue Hu, Jingwen Leng, Qiuhuan Chen, and Chunli Lv. 2023. Enhancing Privacy in Large Language Model with Homomorphic Encryption and Sparse Attention. *Applied Sciences* 13, 24 (2023). <https://doi.org/10.3390/app132413146>
- [143] Ruiqi Zhang, Seira Hidano, and Farinaz Koushanfar. 2022. Text revealer: Private text reconstruction via model inversion attacks against transformers. *arXiv preprint arXiv:2209.10505* (2022).
- [144] Xinyu Zhang, Huiyu Xu, Zhongjie Ba, Zhibo Wang, Yuan Hong, Jian Liu, Zhan Qin, and Kui Ren. 2024. PrivacyAsst: Safeguarding User Privacy in Tool-Using Large Language Model Agents. *IEEE Transactions on Dependable and Secure Computing* (2024).
- [145] Yao Zhang, Zijian Ma, Yunpu Ma, Zhen Han, Yu Wu, and Volker Tresp. 2024. WebPilot: A Versatile and Autonomous Multi-Agent System for Web Task Execution with Strategic Exploration. *arXiv:2408.15978 [cs.AI]* <https://arxiv.org/abs/2408.15978>
- [146] Zhiping Zhang, Michelle Jia, Bingsheng Yao, Sauvik Das, Ada Lerner, Dakuo Wang, Tianshi Li, et al. 2023. "It's a Fair Game", or Is It? Examining How Users Navigate Disclosure Risks and Benefits When Using LLM-Based Conversational Agents. *arXiv preprint arXiv:2309.11653* (2023).
- [147] Zhixin Zhang, Jiaxin Wen, and Minlie Huang. 2023. Ethicist: Targeted training data extraction through loss smoothed soft prompting and calibrated confidence estimation. *arXiv preprint arXiv:2307.04401* (2023).
- [148] Zhixin Zhang, Junxiao Yang, Pei Ke, Fei Mi, Hongning Wang, and Minlie Huang. 2024. Defending Large Language Models Against Jailbreaking Attacks Through Goal Prioritization. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Bangkok, Thailand, 8865–8887. <https://doi.org/10.18653/v1/2024.acl-long.481>
- [149] Haiyan Zhao, Hanjie Chen, Fan Yang, Ninghao Liu, Huiqi Deng, Hengyi Cai, Shuaiqiang Wang, Dawei Yin, and Mengnan Du. 2024. Explainability for large language models: A survey. *ACM Transactions on Intelligent Systems and Technology* 15, 2 (2024), 1–38.
- [150] Xuandong Zhao, Lei Li, and Yu-Xiang Wang. 2022. Provably confidential language modelling. *arXiv preprint arXiv:2205.01863* (2022).
- [151] Xin Zhou, Yi Lu, Ruotian Ma, Tao Gui, Yuran Wang, Yong Ding, Yibo Zhang, Qi Zhang, and Xuan-Jing Huang. 2023. Textobfuscator: Making pre-trained language model a privacy protector via obfuscating word representations. In *Findings of the Association for Computational Linguistics: ACL 2023*, 5459–5473.
- [152] Ligeng Zhu, Zhijian Liu, and Song Han. 2019. Deep leakage from gradients. *Advances in neural information processing systems* 32 (2019).

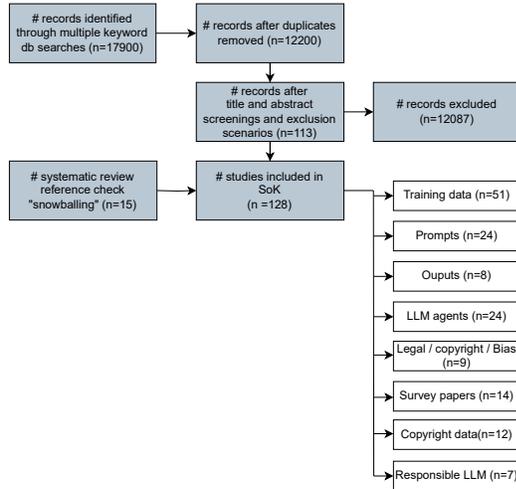
## 6 APPENDICES

### 6.1 Source Selection and Strategy

We conducted a literature review adhering to the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) guidelines [83]. Our search spanned databases such as ACM, IEEE Xplore, Springer, ScienceDirect, and Google Scholar, targeting research papers that explore various privacy challenges in LLM systems published after 2022. The initial screening employed broad keywords such as ‘privacy challenges in LLM OR ChatGPT’, ‘privacy in LLM’, ‘privacy attacks in LLM’, ‘responsible LLM’, and ‘security AND privacy challenges in LLM.’ We then refined our search strategy by categorizing the identified challenges (such as ‘privacy in LLM agents’, ‘privacy in LLM prompts’, ‘privacy in LLM output OR decision’, and ‘privacy in LLM training data’) via screening the identified challenges via initial search. Additionally, we used the mitigation methodologies found in the identified papers to search for further relevant studies. Survey papers on training data privacy challenges in LLMs were included as well.

Figure 5 shows the PRISMA and snowball approach of the paper search and selection process. Our search yielded a total of 17,900 records. We removed duplicates, screened titles and abstracts, and applied exclusion criteria, such as eliminating long-short repetitive papers, papers irrelevant to the focused area, security-related papers, those with no citations or published in non-A/A\*/B conferences, and technical papers focused on root techniques with modifications. After this filtering, we were left with 101 records. We also employed a snowballing technique by reviewing the bibliographies of the identified papers to capture additional relevant

studies. Ultimately, our study comprised 116 papers categorized as follows: Training data (n=51), Prompts (n=24), Outputs (n=8), LLM agents (n=24), Legal/copyright/bias (n=9), Survey papers (n=14), copyright data (n=12), and Responsible LLM (n=7).



**Figure 5: PRISMA and Snowballing Approach for Paper Searching and Selection (# implies number of)**

## 6.2 Glossary

Table 1 lists the abbreviations used in the paper along with their descriptions.

## 6.3 Responsible Large Language Models Concerns and Privacy Implications

This section addresses responsible AI concerns in LLM systems, particularly from user interactions and LLM systems' inherent black-box nature. Responsible LLMs' concerns include ethical and legal issues due to LLMs tendency to copy proprietary material, transparency, reliability, explainability, accountability, fairness, and bias [52, 55].

The first challenge lies in the lack of transparency and interoperability of LLM, with service providers managing user data without clear mechanisms for users to understand its use. Ensuring stringent data protection is crucial due to the rising number of data breaches<sup>5</sup> and the absence of strict regulations. The black-box nature of LLMs complicates verifying data processes, heightening privacy concerns.

Another concern within LLM systems is the perpetuation of unfair discrimination and representational harm through reinforcing stereotypes and social biases [55]. The association of social identities with LLM decisions often results in excluding specific social communities from the LLM's outcomes. Furthermore, hate

speech generated by LLMs may incite violence and cause significant offense. Despite efforts by many LLM service providers to remove such content from their training corpora, bias and unfair discrimination remain persistent issues in LLMs.

Next, the potential for misinformation and disinformation is heightened by LLMs' ability to generate persuasive yet potentially false or misleading content. This capability can be exploited to spread misinformation, manipulate public opinion, or create fraudulent materials, posing significant risks to public health, democracy, and social harmony [10].

Accountability in the deployment and use of LLM systems presents a critical issue, as it involves determining responsibility for the actions and outputs of these models [49]. When LLMs generate harmful or unethical content, such as misinformation or biased responses, it can be challenging to pinpoint who is responsible—the developers, users, or the organizations deploying the models. This issue is compounded by the complexity and opacity of LLMs, which can obscure the decision-making processes of the models. Establishing precise accountability mechanisms, such as transparent reporting standards, ethical guidelines for development, and oversight committees, is essential to ensure that all parties involved are held responsible for the impacts of LLM technologies [11].

**6.3.1 Mitigation.** Various approaches have been proposed to address responsible AI concerns, but they are still nascent for LLM systems. Potential solutions include improving user-friendly privacy policies and ethical data principles, conducting model audits, and using explainable AI techniques like Local Interpretable Model-agnostic Explanations (LIME) [93] and SHapley Additive exPlanations (SHAP) [72] to make LLM processes transparent. Feature visualization and activation maximization offer insights into a model's learning process, though their application to LLMs with millions of parameters is challenging. Additionally, it is crucial to educate users about the privacy implications of LLMs, their data usage policies, and the importance of protecting their privacy. Implementing mechanisms to quantify or explain privacy risks can further inform users of potential consequences. However, further work is needed to evaluate the effectiveness of these solutions.

In conclusion, addressing these concerns requires a multi-faceted approach that encompasses technical strategies (such as algorithms for detecting and mitigating biases), policy development (including privacy protection regulations), ethical guidelines (like respecting intellectual property), and stakeholder engagement (with collaboration among industry, academia, and regulatory bodies).

## 6.4 Summary of Privacy Challenges and Solutions

Table 2 summarizes the privacy challenges and the details of existing mitigation techniques proposed in the literature.

<sup>5</sup><https://www.statista.com/statistics/290525/cyber-crime-biggest-online-data-breaches-worldwide/>

**Table 1: The Glossary Table**

Abbreviations	Description
Artificial Intelligence (AI)	The science of making machines that can think like humans
Machine Learning (ML)	The capability of a machine to imitate intelligent human behavior
Large Language Model (LLM)	Trained on vast datasets and characterized by extensive parameters to generate sensible responses
GDPR	General Data Protection Regulation
HIPAA	Health Insurance Portability and Accountability Act of 1996
LLM systems user prompts	Natural language instructions to describe the task and achieve the desired outcomes
LLM agents	A system endowed with intricate reasoning capabilities, enabling it to interact with other systems and perform actions
Personally Identifiable Information (PII)	Information that can be used to identify individuals
Responsible AI (RAI)	An approach to developing and deploying AI in an ethical, trustworthy, safe, and legal way
Fine-tuning in LLM	The process of adjusting the parameters of a pre-trained large language model to a specific task or domain
In Context Learning (ICL)	A method to adapt the LLM to a specific task by incorporating demonstrations into the prompt
Retrieval Augmented Generation (RAG)	A technique where an LLM retrieves relevant information from external sources to enhance its generated responses
Differential Privacy (DP)	A formal mathematical framework for ensuring privacy in data analysis. It provides strong guarantees that the inclusion or exclusion of any individual’s data in a dataset does not significantly impact the output of a computation, thereby protecting individual privacy. It achieves this by introducing controlled noise to statistical computations, making inferring private details about any single participant difficult. The strength of privacy protection is controlled by a parameter called epsilon ( $\epsilon$ )—a lower $\epsilon$ provides stronger privacy but may reduce accuracy.
Knowledge unlearning	Knowledge unlearning is the process of selectively removing specific information from a trained model while preserving overall performance, which is crucial for privacy protection and regulatory compliance. Common approaches include exact unlearning (retraining without specific data), influence-based methods, and fine-tuning to mitigate privacy risks and data poisoning.
Federated Learning (FL)	A machine learning approach where models are trained across multiple decentralized devices or servers, preserving data privacy by keeping data local
MIA	Membership Inference Attack: The objective of a MIA is to determine whether a specific data point $x$ (a sentence or document in LLM context) was included in the training dataset $\mathcal{D}$ of a model $M$ . This is achieved by computing a membership score $f(x; M)$ , which is then compared against a predefined threshold to decide whether the target sample was part of the training data by analyzing its output [104].
Attribute Inference Attack	In an attribute inference attack, the adversary aims to predict sensitive attributes of an individual whose data was used to train a target model, given that the adversary already knows other attributes of that individual [38].
Jailbreak attacks	A jailbreak attack is an adversarial technique used to bypass the safety restrictions and ethical guardrails of LLMs. These attacks manipulate the model into generating responses that it is designed to avoid, such as harmful, biased, or restricted content. Jailbreak attacks exploit techniques like prompt engineering (deceptive inputs), role-playing (bypassing safeguards through scenarios), encoding tricks (obfuscation or special characters), and multi-turn exploits (gradual manipulation) to evade AI restrictions.
RLHF	Reinforcement Learning from Human Feedback
Intellectual Property (IP)	It encompasses creations of the mind, including inventions and works of art, protected by law through patents, copyrights, trademarks, and trade secrets
API	Application Programming Interface
NER	Named Entity Recognition: This involves identifying and classifying named entities in a text into predefined categories, such as the names of persons, organizations, locations, dates, times, and other entities.

**Table 2: Summary of Privacy Challenges in LLM Systems and Solutions. Experimented on ‘Open/closed source LLM system’: ● = Closed-source LLM systems, ○ = Open source LLM systems, ◐ = Both. ‘Threat model’: ○ = Trusting LLM service provider and consider threats from external parties only, ● = Threat from both service provider and external parties. (This table continues on the next page.)**

Privacy Challenges	Specific Issues	Solution	Technical description	Tested on Open/closed source LLM	Threat model	
● In LLM Training data	Data memorization and personal data leakage	Data deduplication	Fuzzy deduplication techniques [13]	○	○	
		Detection and Filtering	Bloom filter [45], Filtering with restrictive terms [91], PII scrubbing filters and NER [23], Output filtering [105]	○	○	
		Data cleaning	Correcting errors and inconsistencies, implementing anonymization, data minimization, and secure practices to protect sensitive information [13, 124]	○	○	
		Differential privacy	Adding noise to data when training[12, 15, 66, 105, 137, 150]	○	○	
		Knowledge unlearning	Force models to forget specific knowledge without requiring full retraining [33, 130, 140]	○	○	
	Privacy attacks	Model stacking / Dropout	Model stacking: combines multiple base models[27], Dropout: regularization technique used in neural networks to prevent overfitting [123]	○	○	
		Test time defense and instruction processing	Filters malicious inputs, detects abnormal queries, and post-processes LLM-generated output. [64, 96]	◐	○	
		Adversarial regularizations	LLM training process employs robust optimization methods like adversarial training [134] and robust fine-tuning [26] to prevent malicious text attacks	○	○	
		Prompt-level approaches	Filtering adversarial prompts using rule-based detection or classifier-based approaches [148]	◐	●	
		Model-level approaches	Enhancing LLMs through safety training, RLHF, and adversarial fine-tuning [6, 48]	○	○	
● In Prompts	Direct leakage	Input validation and sanitization	NER and encryption [68]	●	●	
			NER and Substitution/masking [17]	●	●	
			NER and obfuscation [73]	◐	●	
		Local small LM and remote LLM combination	Use NER to remove sensitive information before sending data to remote LLM, and add again locally[43]	●	●	
			Text obfuscation	Obscure original word information while retaining original word functionality [151]	○	●
			Cryptographic approach: Fully harmonic encryption	Encrypting data in such a way that computations can be performed on the encrypted data without needing to decrypt it first [16, 42, 100, 139]	○	●
	Inference of sensitive	Local small LM and remote LLM combination	Adversarial feedback-guided approach using LLM [108]	◐	●	
	Contextual information leakage	Hash operation	Transmitting only hashed values of sensitive data and service providers can revert back [133]	-	○	
			Prompt ensembling	Using multiple different prompts to hide actual prompt of the user[28]	○	●
Local small LM and remote LLM combination			Remote LLMs provide high-level guidance and enhance the result in local LM with contextual information [141]	◐	●	

● In LLM-generated Outputs	Revealing Sensitive Information	Text obfuscation	Appending an another random text to the original prompt [129]	-	●
		Differentially private aggregation	Aggregate and release responses without relying on any individual outputs [119]	●	○
			Synthetic few-shot demonstrations [113]	●	●
	Transparency in decision making	Explainable AI techniques	Feature importance analysis[72], interactive exploration tools [82], and human-AI interpretation	-	-
● In LLM Agents	Ambiguities in automated task execution	LLM sandbox	LLM powered emulator and safety evaluator [77, 99], Complex multi-turn interactions [135]	●	○
		Injecting safety knowledge in different stages	Safety knowledge injection in pre-planning, in-planning, and post-planning [44]	●	○
		Self-alignment	Self-alignment of LLMs through multi-agent role-playing [84]	●	○
	Adversarial interactions of agents	Multi-agent debate mechanism	Self-evaluate through discussion and feedback	○	○
		Mitigation measures based on AI and security principles	Monitoring agents' communication content [76]	●	●
	Exposure of Sensitive Information to External Tools	Encryption and shuffling-based solutions	Operate on encrypted inputs and attribute-based forgery generative models with shuffling mechanism [144]	●	●
Access restriction methods		Limits access to necessary data based on user privacy preferences [5]	●	●	