Privacy Preserving Information Retrieval: Defining Privacy Research Pillars for a Future Research Agenda

Francesco Luigi De Faveri*
University of Padova
Padova, Italy
francescoluigi.defaveri@phd.unipd.it

Guglielmo Faggioli
University of Padova
Padova, Italy
guglielmo.faggioli@unipd.it

Asia J. Biega
Max Planck Institute for Security and Privacy
Bochum, Germany
asia.biega@mpi-sp.org

Nicola Ferro
University of Padova
Padova, Italy
nicola.ferro@unipd.it

Abstract

Advancements in computer science are raising concerns and preoccupations about the privacy of users' data submitted to and used by Information Retrieval (IR) Systems. IR Systems, such as search engines, integrate new generative information access pipelines that implement effective and efficient document retrieval and answer generation. However, critical challenges arise in Privacy-Preserving IR (PPIR): Are such data leaking personal information or being used to train generative systems? Are current privacy solutions sufficient to guarantee user privacy and limit such information leakage? Are users' queries and retrieved documents protected throughout the entire retrieval process? How has the privacy threats landscape changed, and in which directions should the IR and Privacy research community investigate to address such new risks?

In this perspective paper, we provide initial answers to these questions, analysing state-of-the-art solutions for protecting user privacy when accessing information and highlighting areas of concern. We propose a new PPIR research agenda to address the unsolved problem of private data use and access. The agenda includes novel privacy research pillars aimed at addressing objectives grounded in literature gaps, user survey findings, and structured interviews with experts from research, industry, and regulatory bodies. By defining these privacy research pillars, we advise the IR community to pursue research toward a more resilient privacy direction that can address future challenges stemming from rapidly advancing technology eager for user data.

CCS Concepts

• **Security and privacy → Privacy protections**; **Usability in security and privacy**; • **Information systems → Users and interactive retrieval**; **Retrieval tasks and goals**.

*Corresponding Author.

SIGIR '26, Melbourne, Australia

ACM ISBN 978-1-4503-XXXX-X/2018/06
https://doi.org/XXXXXXX.XXXXXXX

Keywords

Privacy Preserving Information Retrieval, Perspectives in Information Retrieval, Information Security, User Safety

ACM Reference Format:
Francesco Luigi De Faveri, Guglielmo Faggioli, Asia J. Biega, and Nicola Ferro. 2026. Privacy Preserving Information Retrieval: Defining Privacy Research Pillars for a Future Research Agenda. In *Proceedings of The 49th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '26).* ACM, New York, NY, USA, 12 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

Thanks to their pervasiveness, Information Retrieval (IR) systems, especially search engines, handle increasingly large volumes of data, including sensitive ones. Examples of such data include queries issued by the users, documents clicked, and the search history. Indeed, users employ IR systems to find critical documents relevant to their personal sphere, such as health and finance [18, 21, 30, 44]. Another common behaviour is the *ego-surfing*, which consists of searching for one's own name, email, or phone number on the web, leaving such identity information available to the IR system [56]. Collecting, processing, and storing this data, both identifying and highly sensitive, raises pressing privacy concerns. In fact, upon theft or data exfiltration, search logs might expose users' private information and reveal their personal details [24]. Similarly, a malicious employee might be tempted to access query logs to discover private information and blackmail the individual who generated the queries. Finally, based on their digital footprint, users might be profiled against their will. The problem is further aggravated by Large Language Models (LLMs) becoming an integral part of modern search engines. If private data, including query logs, is used to train such models, there is a high risk of leakage and exposure [15, 26].

To address these privacy challenges, Privacy-Preserving Information Retrieval (PPIR) focuses on limiting the disclosure of personal information to both internal and external threats in the retrieval pipeline [70, 81, 82]. PPIR combines Information Security and IR to minimise the risk of information leakage from legitimate users' and systems' data. To achieve this, PPIR often incurs non-negligible costs in terms of effectiveness and efficiency. Multiple strategies have been adopted to address the PPIR task. One of the most common solutions is query obfuscation, which perturbs a query to ensure that it does not disclose the original information need, while

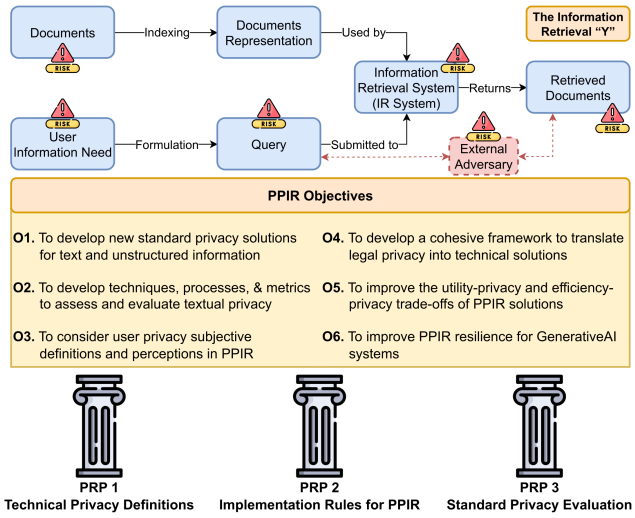

Figure 1: A schematic high-level view of the IR "Y" and potential risks and threats for the users' privacy, its relation with our proposed PRPs to achieve the privacy objectives.

still retrieving relevant documents. Traditionally, query obfuscation has been implemented via heuristic approaches [8, 39] that rely on hypernyms and words co-occurrence, or ε-Differential Privacy (DP)-based solutions [28, 37, 53, 79, 80]. While query obfuscation has a (relatively) limited impact on the IR system efficiency, it often achieves privacy at the expenses of effectiveness. Another alternative is based on homomorphic encryption protocols [3], in which the data, both documents and queries, remain encrypted throughout the entire retrieval pipeline, including during the matching phase. Most of the homomorphic encryption solutions preserve the retrieval effectiveness, but come with huge efficiency costs, making it a viable strategy only for extremely sensitive domains (e.g., military scenarios and international security documents) [50, 66]. As a final remark, we stress that, at the moment, the research regarding the evaluation of the privacy induced by different PPIR techniques is limited and often based on the formal parameters of the privacy protection strategy (e.g., the key size in homomorphic encryption, level of words' generalisation, or the privacy budget ε in DP).

To address these limitations, we argue that the PPIR research community need well-defined and common reference framework for the development of novel PPIR techniques and for guiding their evaluation and comparison. In this paper, we contribute to the definition of such a reference framework by putting into perspective the existing literature on PPIR, identifying potential critical limitations of current privacy-preserving solutions, and outlining future research steps to overcome them. More in detail, to help us identify the limits of the current privacy-oriented research landscape, outline the most urgent needs, and the desired outcomes of the PPIR research, we present here the results from a user study investigating the current perceptions of both experts and general users regarding PPIR techniques. Guided by the results of such a study, we define the research objectives that PPIR should strive to achieve. Critical to reaching the identified objectives, we finally delineate the following **Privacy Research Pillars (PRPs)**, summarised and linked with the IR "Y" in Figure 1, to shape the future research paths and address the gaps identified:

PRP1 Technical Privacy Definitions: To realise this pillar, we envision the development of new PPIR techniques that better align "formal" guarantees provided by the major privacy frameworks, e.g., DP, and the users' perception.

PRP2 Implementation Rules for PPIR: This pillar can be achieved through new implementation rules for IR pipelines, with a *privacy-by-desing* perspective. In particular, we call on the community to identify effective strategies to guide the implementation of legal privacy requirements into the components of actual IR systems ready for a user-safe deployment.

PRP3 Standard Privacy Evaluation: To uphold this pillar, the community should investigate the development of a unified evaluation framework to compare and improve different PPIR solutions.

To promote the reproducibility of our findings, the code used in this paper and the results are available in the online repository[1].

The paper is structured as follows: in Section 2, we discuss the background and frame current PPIR techniques. In Section 3, we report the results of the interviews with privacy experts from industry, research, and legislative backgrounds, aimed at identifying the research objectives for the research agenda. Section 4 discusses the identified privacy research objectives supporting the discussion with users' perspectives and experimental results. Finally, Section 5 defines the future PPIR research agenda and its pillars.

[1] https://anonymous.4open.science/r/PrivacyResearchAgenda4PPIR/README.md

2 Background

In this Section, we define the core privacy concepts from the Information Security and legal domains, providing also a comprehensive background on the definition and methodologies for PPIR.

2.1 Privacy Definitions

Different definitions of privacy have been articulated through different legal and technical frameworks. In the cybersecurity context, the National Institute of Standards and Technology Special Pubblication, i.e., NIST SP 800-130 [58], defines "Privacy" as the property of a system to guarantee the "[. . .] *assurance that the confidentiality of, and access to, certain information about an entity is protected.*". The publication emphasises privacy as an operational characteristic of secure and reliable information systems, focusing on mechanisms necessary to safeguard information. Furthermore, the NIST Privacy Framework[2] is a tool developed to assist organisations in identifying and managing privacy risks effectively, while promoting the deployment of new products and services.

On a more technical level, ε-Differential Privacy (DP), introduced by Dwork et al. [33], provides the *gold standard* mathematical framework for quantifying privacy guarantees when querying and analysing data. Formally, the framework states that, considering a randomisation mechanism $\mathcal{M}$, i.e., an algorithm that takes an input and produces an output, is said to be ε-DP if and only if, for a privacy budget $\varepsilon > 0$ and two neighbouring datasets[3] D, D', it holds

[2] https://www.nist.gov/privacy-framework
[3] Datasets that diverge for at most one entry.

the inequality $\Pr[\mathcal{M}(D) \in S] \leq e^{\varepsilon} \cdot \Pr[\mathcal{M}(D') \in S]$ for all possible outputs $S \subseteq \text{Range}(\mathcal{M})$. The privacy budget ε controls the privacy loss: a smaller ε corresponds to stronger privacy guarantees, and a higher ε provides weaker privacy guarantees. The DP definition was also extended and relaxed to different application scenarios, including Machine and Deep Learning [2] and Natural Language Processing (NLP) [17, 40], with the latter used to formalise the concept of *metric*-DP. Such extension says that considering every triplets of points $x, x', \hat{x}$ in a vector space $\mathcal{W}$ with distance function $d : \mathcal{W} \times \mathcal{W} \rightarrow \mathbb{R}$ and a privacy budget $\varepsilon > 0$, a mechanism $\mathcal{M}$ is ε-DP with respect to the metric d if it holds the inequality $\Pr[\mathcal{M}(x) = \hat{x}] \leq e^{\varepsilon \cdot d(x,x')} \cdot \Pr[\mathcal{M}(x') = \hat{x}]$.

From a human rights perspective, the Universal Declaration of Human Rights [1] establishes privacy as a fundamental right in Article 12, which states that *"No one shall be subjected to arbitrary interference with his privacy, family, home or correspondence, nor to attacks upon his honour and reputation. Everyone has the right to the protection of the law against such interference or attacks."*. In addition, different technology regulations, such as the European Union's General Data Protection Regulation General Data Protection Regulation (GDPR) [34], mandate different principles for the handling and processing of personal data. Among others, the GDPR establishes the principle of Data Minimisation [11, 38, 69]: personal data must be *"[. . .] adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed."*. This principle requires data controllers, i.e., the entities that determine the purposes and means of processing personal data, to critically evaluate and justify each data element they collect, ensuring specified purpose, adequacy of the data, and limited data volume, quality, and retention periods. In this broad context of privacy definitions, we now show how privacy is studied and investigated in IR.

2.2 What is PPIR?

To clearly define Privacy-Preserving Information Retrieval (PPIR), we first must distinguish it from Private Information Retrieval (PIR). A PIR protocol [22, 23] is a process that enables the extraction of a specific data point from a server, given its exact index in it. The implementation of PIR requires using cryptographic primitives and specialised secure communication protocols [65]. In contrast, PPIR involves a user who typically conveys a vaguely defined information need in the form of a plain natural language query, which is then altered to produce an obfuscated version. Specifically, the server needs to estimate document relevance based on the query text received. This distinction holds two significant implications to diversify between PPIR and PIR:

i) Cryptography cannot protect the contents of the query from the IR system, as the system must evaluate the relevance of the query and the documents. A non-privacy cooperative IR system is the intended recipient of the query and would require decryption if the query were encrypted. Indeed, matching encrypted queries to documents poses significant challenges (cf., *Homomorphic Encryption* paragraph).
ii) The "index" of relevant documents is indeterminate prior to the search; thus, relevance remains latent and often uncertain for the user until they review the list of retrieved documents to resolve their information need.

Query Obfuscation Techniques. Query Obfuscation [13, 36, 71, 78] is a client-side privacy-preserving approach in IR in which the user's original query is transformed into one (or multiple) obfuscated variants before submitting it to the server. This technique enables privacy-conscious users to formulate natural language queries while hiding their actual information needs through textual obfuscation. In such scenarios, the IR system operates under the *honest-but-curious* adversarial model [29], where the server correctly returns ranked document results in response to user queries while potentially collecting and aggregating such queries to construct user profiles based on their search behaviour. An obfuscation mechanism is an algorithm that takes the original query text and produces one or more modified versions of it. In the literature, there are two main categories of mechanisms: heuristics and ε-DP based algorithms [33]. In the heuristic family, the terms in the query can be changed with synonyms and hyponyms, like in the Arampatzis Et Al. (AEA) [8] mechanism or with co-occurring words from a local corpus, such as in the Fröbe Et Al. (FEA) algorithm [39]. On the other hand, ε-DP-based obfuscation mechanisms can be divided into *noisy sampling* [16, 19, 83] and *noisy embeddings* [28, 37, 79, 80]. Noisy term embedding strategies, including Cumulative Multivariate Perturbation Mechanism (CMP) [37], the Mahalanobis (Mhl) mechanism [79], and Vickrey's auction-based modifications [80], changes non-contextual embeddings, e.g., GloVe [59] 300-d vectors, by adding ε-calibrated statistical noise to each query term vector, selecting the obfuscated term nearest to the perturbed representation. DP-COMET [28] modifies the CMP and Mhl algorithms by adding noise to the contextual representations of the entire query text and sampling obfuscated queries from a publicly available query log. Conversely, noisy sampling mechanisms, such as SanText [83], CusText [19], and the TEM [16], select obfuscation terms via distance-based probability sampling, with the parameter ε and the embedding proximity used to compute these probabilities.

Homomorphic Encryption. On the other hand, Homomorphic Encryption frameworks [20, 42, 63] use a cryptographic approach to IR that enables computationally secure operations to be performed directly on encrypted data, i.e., queries and documents, without requiring decryption. These protocols allow users to submit encrypted queries to the IR system, which can then compute relevance scores and retrieve documents while operating exclusively on ciphertext, thereby preserving the query confidentiality throughout the whole retrieval process. Similarly, Secure Inner Product protocols [4, 9] provide privacy-preserving mechanisms for computing similarity measures between query and document embedding vectors in encrypted spaces. In such cases, the IR system operates under a strong privacy model in which neither the server nor potential malicious external adversaries can access plaintext queries or infer user information needs from encrypted operations. These present two principal challenges: computing the score, which is due to the encrypted vectors and thus reduces effectiveness, and the efficiency constraints imposed by the encryption operations.

Other Privacy Techniques. The Query Obfuscation and Homomorphic Encryption frameworks are not the only ones used to provide privacy. Federated learning [52] is a machine learning technique that has been employed as a privacy technique, especially for Learning to Rank and Recommender Systems tasks [5–7, 77].

In a federated setting, neither users' interactions nor their queries are shared directly with anyone: the central server serves as an aggregation point for local parameters, which are then used to update the local nodes. In addition to federated learning, synthetic data generation has been employed in recommendation systems to protect user privacy [12, 49]. The idea is to create synthetic datasets using generative models that preserve key statistical properties of the original data. Therefore, by employing "fake" data that resembles the originals, the trade-off between privacy and utility can limit the exfiltration of private information while minimising performance loss. However, not using real data during data processing can prevent generalisation to real-world scenarios.

3 Expert-Identified PPIR Challenges and Directions

To identify the main objectives for future PPIR research, we draw on two sources: the perspectives of experts who deal with privacy-related issues, and the views of the subjects whose privacy should be protected. We collect the former through structured interviews, and the latter via an anonymous survey. In this section, we describe the research challenges and future directions, as outlined by the privacy experts during their interviews. Figure 2 displays boxes that summarise the main findings from such interviews.

We developed a structured expert interview protocol comprising open-ended questions to capture the perspectives of four specialists across three domains of privacy implementation: Technical Research, Industry Applications, and Policy and Regulation. Two of the participants were computer science researchers: one focused on privacy in NLP and the other on privacy in IR, both of whom had published extensively on privacy at major computer science conferences. Further, two participants included an industry Cybersecurity expert and a Legal Expert. We refer to these experts as follows: Privacy Researcher (PR), Information Retrieval Researcher (IRR), Cybersecurity Expert (CE), Legal Expert (LE). The interview protocol is organised into the following key topical sections: i) Main advantages and bottlenecks of current PPIR solutions, ii) Threats and Risks associated with artificial intelligence and privacy when retrieving information, and iii) Future challenges regarding PPIR. The complete interview protocol, including the questions asked, is publicly available in the repository[4]. Each of the four expert interviews was conducted independently by one author.

The first objective that PPIR research should aim at achieving is highlighted by the PR, who focused on the limitations of DP. In particular, they observed that although ε-DP is insufficient to protect the privacy of textual data such as queries and documents, it is an initial defence strategy for protecting user privacy when dealing with unstructured data. PR also stressed that one of the main challenges presented by obfuscation mechanisms based on DP is that the formal privacy budget ε used can vary widely across mechanisms, and higher or lower formal parameters do not affect the actual privacy provided to the texts. This observation also aligns with previous literature on the topic [29, 45, 76]. In practice, most existing privacy frameworks are designed for structured data and provide strong guarantees only in such a usecase. Therefore, as an actionable objective derived from the interview with PR, we mention the following statement:

[4]https://anonymous.4open.science/r/PrivacyResearchAgenda4PPIR/README.md

***Objective O1**: "To develop new standard privacy solutions for text and unstructured information."*

PR emphasised the importance of developing a unified framework to evaluate the privacy of textual data used in different steps of IR, creating a formal benchmark able to produce a hierarchy among privacy mechanisms, developing *ad-hoc* test collections, privacy metrics, taking into account also the user perspective on what privacy in such unstructured context is. Furthermore, as a methodological note, PR and IRR agreed on the importance of developing open-source algorithms to help the research validate, reproduce, and further improve state-of-the-art mechanisms. This leads to the second objective to achieve effective PPIR:

***Objective O2**: "To develop techniques, processes, and metrics to assess and evaluate textual privacy."*

IRR stated that the privacy issue necessitates a more thorough explanation to users to enhance their awareness and education, so that they do not settle for the default privacy settings provided by private companies, but seek more suitable and tailored alternatives. Both PR and IRR underscored the importance of understanding what users perceive as privacy for textual data and the need to develop new standard definitions that are strictly dependent on user awareness and the application context. Therefore, to develop effective PPIR, it is necessary to achieve the following objective:

***Objective O3**: "To consider User Privacy subjective definitions and perceptions in PPIR."*

A different point of view on the concept of privacy is provided from a legal and policy perspective. LE emphasised that the existing technical and legal frameworks do not provide a unified definition for translating legal privacy guarantees into technical privacy solutions. Hence, this absence of a clear definition complicates the *a priori* definition of which contexts should be classified as private. Furthermore, this disparity deepens the range between privacy as an area of research in computer science and law, and the perceptions users hold about what privacy is and what they believe it should represent. As a future research objective, LE identified the need for a cohesive framework that encourages and enables the translation of technical definitions of privacy into their legal counterparts. This can be translated into a practical objective:

***Objective O4**: "To develop a cohesive framework to translate legal privacy concepts into technical solutions."*

According to CE, the main issue with the current PPIR techniques is related to the effectiveness of masking information before querying and accessing it. They highlight the negative impact of PPIR techniques on the trade-off between privacy and utility, often resulting in significant performance degradation. Both effectiveness and efficiency are severely harmed by PPIR. These shortcomings include retrieving non-relevant documents and extending the overall retrieval process time. In this regard, both PR and the IRR highlighted the current computational effort of PPIR solutions, including the resources required by cryptographic protocols and the embedding of textual representations. Such resource constraints may slow the widespread implementation of PPIR measures in practical scenarios. In such contexts, users may choose to lose privacy guarantees in

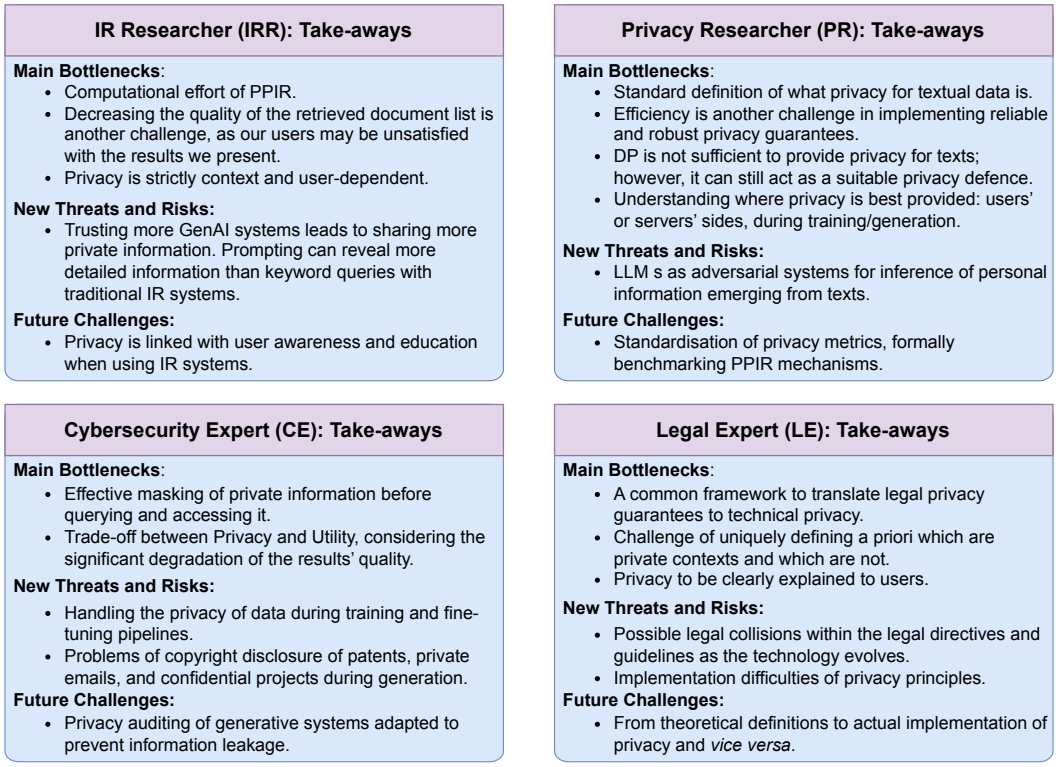

Figure 2: Main Take-aways from the interviews with experts. The points raised by the experts are organised considering the main topic of the interview protocol: bottlenecks of PPIR, new threats and risks, and future challenges.

favour of rapid access to relevant documents. This guides us to formalise the fifth objective of our PPIR research agenda:

***Objective O5**: "To improve the utility-privacy and efficiency-privacy trade-offs of PPIR solutions."*

A common topic discussed by all the experts regards the role that LLMs and GenAI have in novel IR paradigms and with respect to PPIR. According to CE, the risks and threats associated with the use of LLMs for IR include the secure handling of private and personal data throughout the training and fine-tuning processes of the generative systems used by company staff and production. These concerns include the potential for the leakage of sensitive information, such as private data, confidential emails, and corporate patents, particularly when employing RAG pipelines. Furthermore, as future challenges, CE emphasised the importance of researching new privacy-auditing procedures for these systems to effectively address them and mitigate the exposure of private information. Similarly, IRR highlighted a potential risk of self-disclosure of personal information within generative systems. In traditional IR systems, users typically construct queries consisting of a limited number of terms to represent their information needs. In contrast, RAG systems often yield prompts that span a broader range of terms and information, thereby enhancing the quality of generated responses. This more specific context may pose privacy concerns, as it increases the likelihood of disclosing more personal information. Both researchers recognised that another privacy issue arises from the inadvertent generation of personal information that may have been used during the training phase. On a slightly different line, PR suggested that LLMs cannot be used as black-box systems to ensure privacy, as they provide no formal guarantees. Finally, LE reported that complying with legal requirements can be challenging. Although privacy principles exist, e.g., Data Minimisation in the GDPR [34], it is often unclear how to implement them concretely in technical systems, especially in generative systems [55]. As a common objective for our community, this can be translated into:

***Objective O6**: "To improve PPIR resilience for GenAI Systems."*

4 Perspectives on Achieving the PPIR Objectives

This Section discusses the objectives defined in Section 3 for the IR pipeline. For each objective, we also detail its perceived criticality to the group most affected by privacy violations: the end users.

To understand IR users' perspectives and preferences regarding privacy in information access systems, we developed a multiple-choice question survey. The survey was promoted using online social networks and survey platforms[5]. The study was available in English to people aged 18 years and above. We did not use any additional inclusion or exclusion criteria for the survey. Table 1 reports the demographics of the 63 users who completed the survey.

[5]We published the user study on research-related online platforms, i.e., SurveyCircle (https://www.surveycircle.com/en/) and SurveySwap (https://surveyswap.io/)

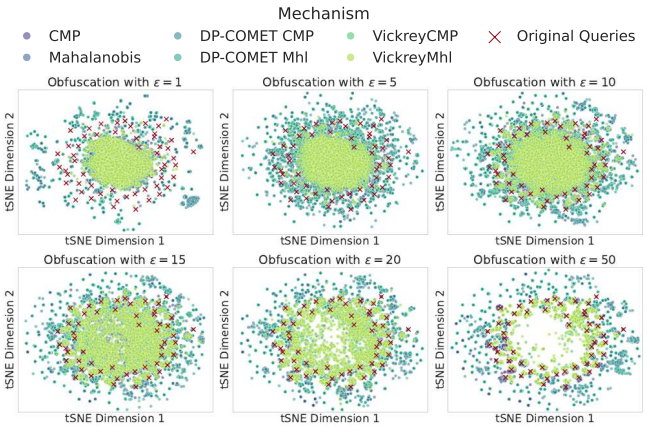

Figure 3: Bidimensional t-SNE of original and obfuscated versions of the TREC Deep Learning 2020 queries [25]. The × are the embeddings of the original queries, while the coloured circles are the ones produced by different obfuscations.

Data was collected anonymously. The first page of the survey outlined the scope of the questions and informed participants that no personal information would be collected or processed. We release the survey items and the results in the paper repository.

Table 1: Users' demographic observed in the survey.

Variable	Category	Users	%	Variable	Category	Users	%
Age	18-25	20	31.7	*Location*	Europe	49	77.7
	26-35	29	46.0		North America	10	15.9
	36-45	6	9.5		Asia	2	3.2
	46-55	3	4.8		Other	2	3.2
	56-65	3	4.8	*Occupation*	Students (BSc, MSc)	19	30.1
	>65	2	3.2		Researchers	18	28.6
Gender	Male	40	63.5		Industry Employees	17	27.0
	Female	23	36.5		Others	9	14.3

O1: New Privacy Solutions. Specifically highlighted in the experts' interviews, the first objective for a more robust PPIR is the need for new privacy-preserving definitions. Current state-of-the-art techniques in PPIR rely mainly on heuristics, formal privacy guarantees, or cryptographic frameworks. However, IR and privacy practitioners should focus on the meaning of privacy for unstructured content, such as text. Unlike structured data, like tables containing *identifiers*, *quasi-identifiers* and *sensitive information* [31], whose privacy can be adequately protected using solutions such as k-anonymity [72], ℓ-diversity [51] and t-closeness [47] or the traditional definition of ε-DP [33], unstructured data might be vague and contain private information beyond those categories that cannot be protected in the same manner. An important consideration currently missing in the application of formal frameworks for textual privacy is a complete understanding of the goals for safeguarding words in user queries and documents. As noted by the PR expert, DP is a possible solution to enhance privacy guarantees of obfuscated textual data, as it is formally grounded in a mathematical framework. However, to ensure privacy concretely, obfuscation mechanisms should go beyond formal formulations and explore this problem through a contextual lens [10, 57]. In non-structured texts, the concept of privacy does not depend only on the presence of specific terms, and it is especially influenced by the context in which a word appears. For example, the sentences *"The patient shows symptoms of mercury present in the blood"* and *"The child said that Mercury is the smallest planet in the solar system"* illustrate how context can reveal privacy concerns. Despite both sentences containing the term *"mercury"*, they require a distinct approach to safeguard privacy: the first sentence reveals an individual's medical diagnosis and thus concerns sensitive and confidential information, whereas the second sentence addresses a widely recognised scientific fact. Thus, it is important to handle these cases with privacy strategies tailored to the specific context, and only when necessary.

O2: Standard Textual Privacy Evaluation. Once privacy is provided for unstructured content, an important challenge is to assess the level of privacy for such data. In particular, we want to ensure that the obfuscated text does not leak information about the original one. This is particularly challenging when it comes to encoding text in dense embedding spaces, used by most modern IR systems, as it has been observed that sentence embeddings can leak the overall sentence, enabling inference attacks [46, 74].

To assess such a risk, we use `all-MiniLM` model [61] to embed the original TREC Deep Learning 2020 [25] queries and their obfuscated versions into a 768-dimensional embedding space. To generate the obfuscated queries, we use the pyPANTERA framework [27]. Such obfuscations are constructed using different DP mechanisms, namely *Noisy Embedding* and *Sampling*, cf. Section 2. Traditional privacy evaluation measures compute similarity between dense representation vectors to assess the semantics of the underlying texts. Formally, this means computing the cosine of the angle between two vectors in an Euclidean space. To visually show what this means, we use t-distributed Stochastic Neighbour Embedding (t-SNE) dimensionality reduction to obtain a two-dimensional representation of the queries and their obfuscations. The t-SNE algorithm [75] maps high-dimensional embeddings to a two-dimensional plane, i.e., $\mathbb{R}^2$. This technique ensures that vectors that are similar in the higher-dimensional plane are represented as points close together in $\mathbb{R}^2$, while dissimilar vectors are mapped as points farther apart, thereby preserving the original similarity relation within the vectors. Figure 3 shows the t-SNE of the DP query obfuscations at different levels of privacy budget ε. From the results we obtained for the query obfuscation protocol, we showed how the embeddings of the obfuscations distribute in the vector space relative to the originals. As the ε increases, we report a clear trend for almost every DP mechanism: the embeddings of the obfuscated queries converge significantly towards those of the original ones. This geometric proximity implies that the obfuscated variants retain a semantic representation close to the sensitive original one, which is intended to be masked. Thus, the cosine similarity computed using the traditional privacy metric for high ε indicates that the obfuscation mechanisms fail to properly assess the privacy protection for the users' intents. In a real-world scenario, an adversary, or a malicious IR system, could exploit this high semantic similarity to easily invert the obfuscation and infer the original users' query [29].

O3: Users' Privacy Perspectives. To develop more reliable and resilient privacy solutions, PPIR researchers should also consider

Table 2: Average nDCG@10 achieved on the TREC Deep Learning 2020 [25] submitting 20 obfuscated variants (per query in the collection) to the TAS-B retrieval model [41], reranking the retrieved list of documents with Dragon+ [48].

Kind of Queries submitted		ε - Privacy Budget					No DP
		1	5	10	15	20	
Original Queries							0.672
Obfuscation Strategy	**Mechanism**						
Heuristics	AEA [8]						0.590
	FEA [39]						0.153
Noisy Embedding	CMP [37]	0.002	0.024	0.581	0.668	0.666	
	Mahalanobis [79]	0.000	0.002	0.426	0.657	0.667	
	VickreyCMP [80]	0.002	0.001	0.480	0.653	0.667	
	VickreyMhl [80]	0.010	0.016	0.302	0.630	0.657	
	DP-COMET CMP [28]	0.050	0.687	0.697	0.697	0.697	
	DP-COMET Mhl [28]	0.024	0.687	0.694	0.695	0.696	
Noisy Sampling	TEM [16]	0.000	0.671	0.671	0.671	0.671	
	SanText [83]	0.003	0.666	0.665	0.666	0.666	

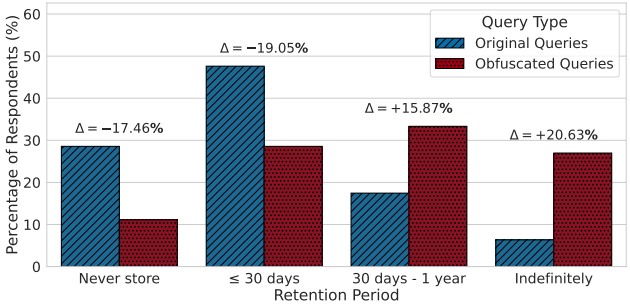

Figure 4: *"If a search engine stores your original/obfuscated queries to improve services, what do you consider an acceptable retention time period for them?"*. The Δ represents the percentage difference between users' preferences for the acceptable retention time of obfuscated *vs.* original queries.

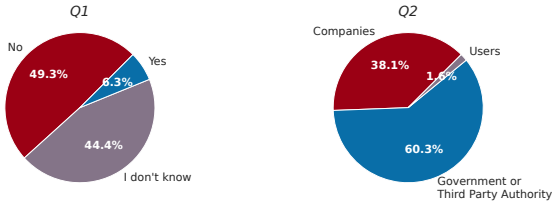

Q1: Do you feel that current Privacy Policies provided by search engines are written in a way that you can easily understand?

Q2: Who do you think should be responsible for protecting the user privacy during the search?

Figure 5: User results on the perspectives about privacy policies and responsibilities during search tasks.

users' perspectives on what is and how intended privacy is achieved in such scenarios. Regarding the use of cryptography in interactions with IR systems, a significant majority of survey respondents (54%) indicated that current measures are inadequate for ensuring their privacy. Conversely, only 7.9% of participants believed that the current level of cryptographic protection is sufficient to achieve strong privacy in this process[6]. On the other hand, in Figure 4, we report the aggregated findings from two questions about the retention times of queries submitted to a search engine. When asked for how long their original queries should be retained by the IR system, the majority of respondents were divided between the two options: "Never store" and "Less than 30 days". Conversely, for the obfuscated queries, the retention time shifted towards an elapsed period, i.e., "From 30 days to 1 year" and "Indefinitely". This distribution shift reveals a significant difference in users' acceptable retention times between the two query types. Users generally prefer shorter retention periods for original queries, whereas they favour longer storage times for obfuscated queries. This preference indicates a difference of 15.87% to 20.63% in the retention times of obfuscated and original queries, suggesting a tendency for obfuscated queries to be retained for more than 30 days and to persist indefinitely.

[6]Pie charts with the complete survey statistics are available in the repository.

O4: Legal and Technical Privacy. An important aspect to consider when developing PPIR solutions is the compliance with the law, thus being able to translate a legal guideline or directive into a technical framework [55, 62]. In our survey, we examined users' perceptions of search engines' privacy policies, specifically their transparency and clarity. Figure 5 reports the users' perspectives on such matters, also analysing them from a legal point of view when using IR systems, requiring who should be accountable for privacy in such scenarios. The findings indicate that only 6.3% of users feel they can easily comprehend these policies. In contrast, 49.3% of respondents believe the policies are not clearly written, hindering their effective understanding. The remaining users, i.e., 44.4% of participants, reported that they are either unaware of the policies or do not read them. On the other hand, when asked who should be held accountable for privacy, the general public primarily falls into two categories: a negligible minority of 1.6% of the users believed that the responsibility should lie with the users themselves. Conversely, 60.3% of the interviewed were in favour of government or third-party authority intervention, while 38.1% expect search engine providers to be responsible for the privacy of their systems. This signals strong demand for solutions that enforce a *"privacy by default"* and for regulatory oversight rather than active user configuration. These results emphasise a major concern also pointed out during the IRR interview: users' awareness and behaviour regarding privacy solutions. Figure 5 shows a scenario where, since the users do not understand the dynamics of privacy, they do not want to be accountable for it, rejecting the concept of individual liability, and thus viewing PPIR not as a strictly personal task, but as an external obligation directed at them.

***O5: Utility* vs. *Privacy*.** The Utility *vs.* Privacy trade-off is at the core of PPIR. This remains one of the most critical issues, especially when retrieving relevant documents from an untrusted system that lacks any mechanism to protect users' privacy. For example, in query obfuscation protocols [13, 27, 36], the user query is altered before being submitted to the IR system using obfuscation mechanisms, cf. Section 2. To empirically test this utility-privacy trade-off, Table 2 reports the results of a query obfuscation protocol, using the same approaches used to build Figure 3. In addition, we used heuristic-based query obfuscation employing the AEA and FEA mechanisms. The utility, measured as normalized Discounted Cumulative Gain (nDCG) [43] at the first 10 documents (nDCG@10), is achieved by submitting 20 obfuscated queries, for each of the

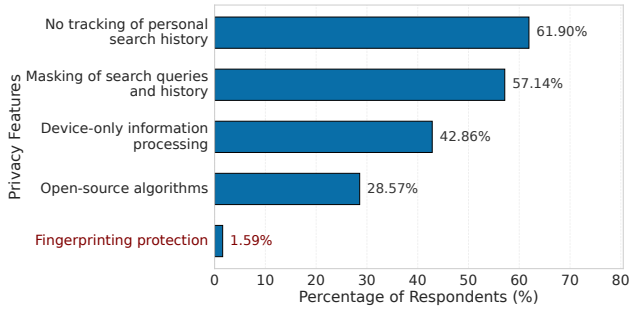

Figure 7: *"If a new search engine were to be released, which privacy feature would be a "must-have" for you to switch?"*. Privacy Features preferences noted by users, the red item shows a feature added by one of the respondents.

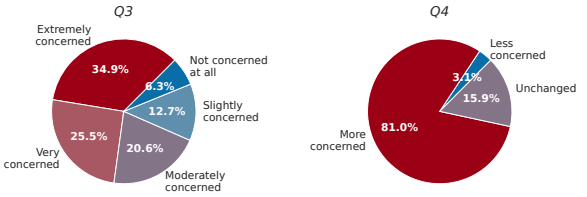

Q3: How concerned are you that the queries you prompt to an AI system might be used to "train" the model and the data you provided could potentially reappear in future answers given?

Q4: With the rise of AI-powered search (e.g., Google Gemini or MS Copilot), how does your concern regarding privacy change?

Figure 6: Users' concerns about personal texts being used by Generative AI models integrated in IR systems.

original queries in the TREC Deep Learning 2020 collection [25] to TAS-B [41]. The original queries were ultimately used to rerank the retrieved document list employing Dragon+ [48] as a local system.

The results illustrate the impact of varying privacy levels, as indicated by changes in the privacy budget ε. These changes yield distinct values of nDCG@10, therefore showing one major concern raised by PR during their interview. Specifically, stricter privacy regimes tend to result in greater user effectiveness loss, while less private configurations are associated with higher nDCG. Conversely, heuristic-based mechanisms are more effective for users than hard DP privacy mechanisms, such as those with $\varepsilon = 1, 5$. We emphasise that this heuristic level of privacy is not strictly comparable to the formal assurances of ε-DP, as the former mechanisms are not defined with respect to any ε, and the privacy parameters are used in different ways during obfuscation. Nevertheless, these results offer a promising outlook for PPIR: the obfuscated queries generated by the DP mechanisms achieve the same, or even outperform, e.g., DP-COMET, the original queries' effectiveness. This shows that it is possible to maintain a full task utility while providing a layer of privacy for the users' requests.

O6: Privacy and Generative Systems. One result of our user study, Figure 6, reports users' attitudes toward integrating generative AI systems into commercial search engines. Specifically, participants of the study were asked about their concerns regarding personal search queries and prompts submitted to such systems that might be used for the model training, potentially resulting in the "data regurgitation" [14] during future system-generated responses. The results show that 34.9% of respondents express extreme concern about this vulnerability, while 6.3% and 12.7% report being not at all or slightly concerned, respectively. The remaining 46% of users indicated either very high concern (25.5%) or moderate concern (20.6%). Furthermore, we interviewed participants about their concerns regarding the rise of AI-powered search engines, e.g., MS Copilot and Google Gemini [32, 73]. A marked 81% majority reported significant concerns about these systems, with only 3.1% expressing less concern about their privacy and 15.9% reporting unchanged concerns toward potential privacy risks.

5 Proposing a Future PPIR Research Agenda

In this Section, we present our proposed PPIR Research Agenda composed of three concepts to shape future research trajectories in IR. Considering the results of the interviews with the experts about the technical feasibility and regulatory frameworks, and the experimental evaluation, to address the objectives defined in Section 4, we present the Privacy Research Pillars (PRPs), principles that we advocate to be investigated by the research community in IR: 1) Technical Privacy Definitions, 2) Implementation Rules for PPIR, and 3) Standard Privacy Evaluation.

5.1 PRP1: Technical Privacy Definitions

In the first Privacy Research Pillar (PRP 1), we propose a redesign of the theoretical principles for technical definitions of privacy. Sensitive information can be exploited for malicious purposes, leading individuals to be unwilling to share it with IR systems. In our survey, we asked participants to express their level of concern about sharing personal information, i.e., medical conditions, personal income, and sexual preferences, with commercial search engines when searching online. A combined 50.80% of respondents indicated that they are either extremely concerned (23.8%) or very concerned (27.0%) about this issue. These findings confirm what was reported by Sharma et al. [69] and raise an important open challenge: the current privacy definitions applied to textual queries and documents are insufficient for appropriate user privacy protection and do not account for which data are sensitive and thus should be more privatised than non-sensitive data. We suggest formulating new privacy definitions that consider the sensitivity of texts as a key aspect when privatising queries and documents in IR pipelines, going beyond formal privacy guarantees, and offering different levels of privacy based on. For example, using a risk-based approach, as outlined in the AI Act in the EU [35, 64], can represent an initial step towards more comprehensive privacy definitions.

PRP 1 aims to achieve the objectives ***O1, O3, O4*** and ***O5***. The development of new privacy approaches depends on formal mathematical definitions of privacy, either based on computational approaches such as cryptography or ε-DP. These definitions set the milestones upon which privacy and IR practitioners can develop new systems that have privacy in their core, not only from a formal perspective but also from a practical one, thereby achieving ***O1***. The new technical definitions should align with each user's personal understanding of privacy. This shift will transform the existing formal privacy model into a more user-centric one. Such a system

will enable users to privately retrieve information while allowing companies to process and analyse the data, ultimately achieving their business objectives in a trusted, user-safe fashion (***O3*** and ***O4***). Finally, the privacy-utility trade-off, maximising the utility of data used and the privacy guarantees provided to such data, represent another guideline for developing these new techniques: finding the optimal degree of trusted information released, how to evaluate such degree, and, at the same time, measuring the utility of the privatised data, can help practitioners to produce adaptive privacy mechanisms that can be tuned to specific scenarios and users (***O5***).

5.2 PRP2: Implementation Rules for PPIR

An essential aspect of privacy also concerns the rules for implementing PPIR solutions into IR systems, which we propose as the second Privacy Research Pillar (PRP 2) of the future research agenda.

Figure 7 shows the users' preferences regarding privacy features for a potential new search engine. The findings reveal that 61.90% of respondents would prefer the absence of tracking techniques for personal queries, while 57.14% favoured the masking of search queries and personal history. In the second place, participants expressed a desire for increased device-only information processing and the implementation of open-source algorithms in such a search engine, with 42.86% and 28.57% of participants expressing these preferences, respectively. These results suggest a hierarchical model of user privacy expectations: the strong preference for non-tracking [54] and for query obfuscation techniques [36] imply that users view anonymity as a fundamental requirement, whereas local processing and algorithmic transparency represent secondary choices. For the development of PPIR frameworks, this implies that, while architectural transparency is valued by users, the success of a privacy-resilient IR systems depends on protecting personal data during the interactions with the deployed system.

Therefore, PRP 2 aims to achieve the objectives ***O3, O4,*** and ***O5***. Such new implementation guidelines can codify users' privacy perspectives more closely, leading to a solution deployment that better reflects users' privacy priorities rather than technical and formal privacy constraints. For example, by organising feedback channels within IR systems, taking into account the heterogeneity of the user population and the different views on privacy and contextual concerns, a new privacy definition and frameworks can be grounded through a user validation process (***O3***). Finally, establishing strict implementation rules for PPIR can help to address the trade-off between privacy and utility, setting new frameworks that quantify and optimise both dimensions. Such implementation guidelines would enable privacy practitioners and researchers to make more informed decisions about acceptable losses in order to achieve specific performance or privacy goals, such as threshold results quality, indistinguishability between original and privatised results, protection against external eavesdroppers, and against inference and membership attacks. Privacy practitioners can develop frameworks that adjust privacy based on query sensitivity and user context, enabling dynamic optimisation of the trade-off rather than a static one. Moreover, new rules for implementing PPIR can bridge the gap between technical privacy mechanisms, user awareness, and policy guidelines, mandating transparency and fairness requirements, generalised consent protocols, and user-designed privacy paradigms. By setting up clear and more understandable privacy policies, cf. Figure 5, and architectural choice, e.g., integrating federated learning or ε-DP into IR pipelines [60, 77], user privacy models can be properly addressed, while maintaining a robustness towards the personalisation and effectiveness of the systems (***O4*** and ***O5***).

5.3 PRP3: Standard Privacy Evaluation

The final Privacy Research Pillar (PRP 3) that we want to define aims to support the importance of assessing the levels of privacy provided to unstructured content. As explained in Section 4, current privacy evaluation frameworks for IR tasks, such as query obfuscation, rely either on the ε-DP definition or on *ad hoc* heuristics. This pillar addresses the central question: *How can we develop rigorous, standardised evaluation frameworks that quantify privacy protection for IR related tasks by accounting for semantic inference, linguistic context, and adversarial reasoning capabilities?*

PRP 3 aims to achieve the objectives ***O2, O3,*** and ***O6***. We envision research integrating information measures of semantic similarity, adversarial protocols, and inference risk frameworks, using state-of-the-art language models as informed attackers, and compositional privacy metrics that capture both direct information leakage and indirect inference through semantic relationships. Such frameworks would serve not only as a ground-truth benchmark for comparing redaction and perturbation-based obfuscation techniques, but also as a means to provide interpretable privacy guarantees that account for what adversaries can realistically infer from protected text (***O2*** and ***O6***). We also advocate a transition to more empirical privacy auditing, shifting from formal assumptions to practical evaluations beyond formal privacy parameter analysis, which may not provide users with a proper assessment of textual privacy. To achieve this standard evaluation of textual privacy, we encourage IR communities, e.g., TREC[7] and CLEF[8], to promote initiatives that aim to create not only *ad hoc* collections for evaluating privacy mechanisms, but also support challenges aimed to formalise new measures for PPIR tasks, as previously accomplished in other studies [67, 68]. Such new collections and challenges can raise awareness in the research community, thus improving the development and deployment of private IR systems that prioritise privacy (***O3***).

6 Conclusion

In this perspective paper, we proposed a novel research agenda for Privacy-Preserving Information Retrieval (PPIR). The main contributions of this paper are the identification of the objectives to achieve a sustainable and effective PPIR and the definition of new Privacy Research Pillars (PRPs), i.e., novel research trajectories for unsolved privacy problems. These pillars aim to open new research paths to develop solutions that address the objectives identified and discussed to mitigate threats and emerging privacy risks in IR. The study considered current state-of-the-art techniques for protecting user and system privacy against adversaries seeking to disclose personal information used in IR systems. The PRPs' objectives are supported by privacy experts from industry, research, and legislative backgrounds who were interviewed, by users' perspectives collected via an online survey, and by experimental evidence obtained employing state-of-the-art protocols. The results show

[7] https://trec.nist.gov/
[8] https://www.clef-initiative.eu/

that such technical solutions provide a preliminary answer to the privacy issue in IR, but are not yet fully satisfactory. Thus, the proposed PRPs outline future work to pursue in new directions, offering valuable insights for the research community. Despite the limited users' sample size, i.e., 4 privacy experts and 63 survey respondents, this study significantly contributes to the definition of a future research agenda for the IR community, and we claim that these findings provide a robust starting point.

GenAI Usage Disclosure

Following ACM's guidelines on the use of generative AI tools (Grammarly Pro), we disclose that generative AI technologies were used solely to assist in code debugging and grammar checking during the preparation of this paper. All research ideas, experiments, analysis, and writing were conducted and critically reviewed by the authors. No part of the scientific content or creative reasoning has been generated or substantially rewritten by generative AI tools alone.

References

[1] 1948. Universal Declaration of Human Rights.

[2] Martín Abadi, Andy Chu, Ian J. Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. 2016. Deep Learning with Differential Privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016*, Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi (Eds.). ACM, 308–318. doi:10.1145/2976749.2978318

[3] Abbas Acar, Hidayet Aksu, A. Selcuk Uluagac, and Mauro Conti. 2018. A Survey on Homomorphic Encryption Schemes: Theory and Implementation. *ACM Comput. Surv.* 51, 4, Article 79 (July 2018), 35 pages. doi:10.1145/3214303

[4] Shweta Agrawal, Benoît Libert, and Damien Stehlé. 2016. Fully Secure Functional Encryption for Inner Products, from Standard Assumptions. In *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part III (Lecture Notes in Computer Science, Vol. 9816)*, Matthew Robshaw and Jonathan Katz (Eds.). Springer, 333–362. doi:10.1007/978-3-662-53015-3_12

[5] Waqar Ali, Xiangmin Zhou, and Jie Shao. 2025. Privacy-preserved and Responsible Recommenders: From Conventional Defense to Federated Learning and Blockchain. *ACM Comput. Surv.* 57, 5 (2025), 119:1–119:35. doi:10.1145/3708982

[6] Vito Walter Anelli, Luca Belli, Yashar Deldjoo, Tommaso Di Noia, Antonio Ferrara, Fedelucio Narducci, and Claudio Pomo. 2021. Pursuing Privacy in Recommender Systems: the View of Users and Researchers from Regulations to Applications. In *RecSys '21: Fifteenth ACM Conference on Recommender Systems, Amsterdam, The Netherlands, 27 September 2021 - 1 October 2021*, Humberto Jesús Corona Pampín, Martha A. Larson, Martijn C. Willemsen, Joseph A. Konstan, Julian J. McAuley, Jean Garcia-Gathright, Bouke Huurnink, and Even Oldridge (Eds.). ACM, 838–841. doi:10.1145/3460231.3473326

[7] Vito Walter Anelli, Yashar Deldjoo, Tommaso Di Noia, Antonio Ferrara, and Fedelucio Narducci. 2021. How to put users in control of their data in federated top-N recommendation with learning to rank. In *SAC '21: The 36th ACM/SIGAPP Symposium on Applied Computing, Virtual Event, Republic of Korea, March 22-26, 2021*, Chih-Cheng Hung, Jiman Hong, Alessio Bechini, and Eunjee Song (Eds.). ACM, 1359–1362. doi:10.1145/3412841.3442010

[8] Avi Arampatzis, Pavlos S. Efraimidis, and George Drosatos. 2011. Enhancing Deniability against Query-Logs. In *Advances in Information Retrieval - 33rd European Conference on IR Research, ECIR 2011, Dublin, Ireland, April 18-21, 2011. Proceedings (Lecture Notes in Computer Science, Vol. 6611)*, Paul D. Clough, Colum Foley, Cathal Gurrin, Gareth J. F. Jones, Wessel Kraaij, Hyowon Lee, and Vanessa Murdock (Eds.). Springer, 117–128. doi:10.1007/978-3-642-20161-5_13

[9] Michael Backes, Martin Gagné, and Sri Aravinda Krishnan Thyagarajan. 2015. Fully Secure Inner-Product Proxy Re-Encryption with Constant Size Ciphertext. In *Proceedings of the 3rd International Workshop on Security in Cloud Computing* (Singapore, Republic of Singapore) *(SCC '15)*. Association for Computing Machinery, New York, NY, USA, 31–40. doi:10.1145/2732516.2732525

[10] Adam Barth, Anupam Datta, John C. Mitchell, and Helen Nissenbaum. 2006. Privacy and Contextual Integrity: Framework and Applications. In *Proceedings of the 2006 IEEE Symposium on Security and Privacy (SP '06)*. IEEE Computer Society, USA, 184–198. doi:10.1109/SP.2006.32

[11] Asia J. Biega, Peter Potash, Hal Daumé, Fernando Diaz, and Michèle Finck. 2020. Operationalizing the Legal Principle of Data Minimization for Personalization. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Virtual Event, China) *(SIGIR '20)*. Association for Computing Machinery, New York, NY, USA, 399–408. doi:10.1145/3397271.3401034

[12] Vincent Bindschaedler, Reza Shokri, and Carl A. Gunter. 2017. Plausible Deniability for Privacy-Preserving Data Synthesis. *Proc. VLDB Endow.* 10, 5 (2017), 481–492. doi:10.14778/3055540.3055542

[13] Danushka Bollegala, Tomoya Machide, and Ken-ichi Kawarabayashi. 2022. Query Obfuscation by Semantic Decomposition. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference, LREC 2022, Marseille, France, 20-25 June 2022*, Nicoletta Calzolari, Frédéric Béchet, Philippe Blache, Khalid Choukri, Christopher Cieri, Thierry Declerck, Sara Goggi, Hitoshi Isahara, Bente Maegaard, Joseph Mariani, Hélène Mazo, Jan Odijk, and Stelios Piperidis (Eds.). European Language Resources Association, 6200–6211. https://aclanthology.org/2022.lrec-1.667

[14] Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramèr, and Chiyuan Zhang. 2023. Quantifying Memorization Across Neural Language Models. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net. https://openreview.net/forum?id=TatRHT_1cK

[15] Nicholas Carlini, Florian Tramèr, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom B. Brown, Dawn Song, Úlfar Erlingsson, Alina Oprea, and Colin Raffel. 2021. Extracting Training Data from Large Language Models. In *30th USENIX Security Symposium, USENIX Security 2021, August 11-13, 2021*, Michael D. Bailey and Rachel Greenstadt (Eds.). USENIX Association, 2633–2650. https://www.usenix.org/conference/usenixsecurity21/presentation/carlini-extracting

[16] Ricardo Silva Carvalho, Theodore Vasiloudis, Oluwaseyi Feyisetan, and Ke Wang. 2023. TEM: High Utility Metric Differential Privacy on Text. In *Proceedings of the 2023 SIAM International Conference on Data Mining, SDM 2023, Minneapolis-St. Paul Twin Cities, MN, USA, April 27-29, 2023*, Shashi Shekhar, Zhi-Hua Zhou, Yao-Yi Chiang, and Gregor Stiglic (Eds.). SIAM, 883–890. doi:10.1137/1.9781611977653.CH99

[17] Konstantinos Chatzikokolakis, Miguel E. Andrés, Nicolás Emilio Bordenabe, and Catuscia Palamidessi. 2013. Broadening the Scope of Differential Privacy Using Metrics. In *Privacy Enhancing Technologies - 13th International Symposium, PETS 2013, Bloomington, IN, USA, July 10-12, 2013. Proceedings (Lecture Notes in Computer Science, Vol. 7981)*, Emiliano De Cristofaro and Matthew K. Wright (Eds.). Springer, 82–102. doi:10.1007/978-3-642-39077-7_5

[18] Chung-Chi Chen, Yongjae Lee, Alejandro Lopez-Lira, Chanyeol Choi, Richard McCreadie, and Javier Sanz-Cruzado. 2025. Information Retrieval in Finance: Industry and Academic Perspectives on Innovation. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2025, Padua, Italy, July 13-18, 2025*, Nicola Ferro, Maria Maistro, Gabriella Pasi, Omar Alonso, Andrew Trotman, and Suzan Verberne (Eds.). ACM, 4090–4093. doi:10.1145/3726302.3731685

[19] Sai Chen, Fengran Mo, Yanhao Wang, Cen Chen, Jian-Yun Nie, Chengyu Wang, and Jamie Cui. 2023. A Customized Text Sanitization Mechanism with Differential Privacy. In *Findings of the Association for Computational Linguistics: ACL 2023*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 5747–5758. doi:10.18653/v1/2023.findings-acl.355

[20] Jung Hee Cheon, Andrey Kim, Miran Kim, and Yong Soo Song. 2017. Homomorphic Encryption for Arithmetic of Approximate Numbers. In *Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 10624)*, Tsuyoshi Takagi and Thomas Peyrin (Eds.). Springer, 409–437. doi:10.1007/978-3-319-70694-8_15

[21] Chanyeol Choi, Jihoon Kwon, Jaeseon Ha, Hojun Choi, Chaewoon Kim, Yongjae Lee, Jy-yong Sohn, and Alejandro Lopez-Lira. 2025. FinDER: Financial Dataset for Question Answering and Evaluating Retrieval-Augmented Generation. *CoRR* abs/2504.15800 (2025). arXiv:2504.15800 doi:10.48550/ARXIV.2504.15800

[22] Benny Chor, Oded Goldreich, Eyal Kushilevitz, and Madhu Sudan. 1995. Private Information Retrieval. In *36th Annual Symposium on Foundations of Computer Science, FOCS 1995, Milwaukee, Wisconsin, USA, 23-25 October 1995*. IEEE Computer Society, 41–50. doi:10.1109/SFCS.1995.492461

[23] Benny Chor, Eyal Kushilevitz, Oded Goldreich, and Madhu Sudan. 1998. Private Information Retrieval. *J. ACM* 45, 6 (1998), 965–981. doi:10.1145/293347.293350

[24] Jonathan Cohn. 2016. My TiVo Thinks I'm Gay: Algorithmic Culture and Its Discontents. *Television & New Media* 17, 8 (2016), 675–690. arXiv:https://doi.org/10.1177/1527476416644978 doi:10.1177/1527476416644978

[25] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, and Daniel Campos. 2021. Overview of the TREC 2020 deep learning track. *CoRR* abs/2102.07662 (2021). arXiv:2102.07662 https://arxiv.org/abs/2102.07662

[26] Badhan Chandra Das, M. Hadi Amini, and Yanzhao Wu. 2025. Security and Privacy Challenges of Large Language Models: A Survey. *ACM Comput. Surv.* 57, 6 (2025), 152:1–152:39. doi:10.1145/3712001

[27] Francesco Luigi De Faveri, Guglielmo Faggioli, and Nicola Ferro. 2024. pyPANTERA: A Python PAckage for Natural language obfuscaTion Enforcing pRivacy & Anonymization. In *Proceedings of the 33rd ACM International Conference on*

Information and Knowledge Management, CIKM 2024, Boise, ID, USA, October 21-25, 2024, Edoardo Serra and Francesca Spezzano (Eds.). ACM, 5348–5353. doi:10.1145/3627673.3679173

[28] Francesco Luigi De Faveri, Guglielmo Faggioli, and Nicola Ferro. 2025. DP-COMET: A Differential Privacy Contextual Obfuscation MEchanism for Texts in Natural Language Processing. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management, CIKM 2025, Seoul, Republic of Korea, November 10-14, 2025*, Meeyoung Cha, Chanyoung Park, Noseong Park, Carl Yang, Senjuti Basu Roy, Jessie Li, Jaap Kamps, Kijung Shin, Bryan Hooi, and Lifang He (Eds.). ACM, 4700–4705. doi:10.1145/3746252.3760888

[29] Francesco Luigi De Faveri, Guglielmo Faggioli, and Nicola Ferro. 2025. Measuring Actual Privacy of Obfuscated Queries in Information Retrieval. In *Advances in Information Retrieval - 47th European Conference on Information Retrieval, ECIR 2025, Lucca, Italy, April 6-10, 2025, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 15572)*, Claudia Hauff, Craig Macdonald, Dietmar Jannach, Gabriella Kazai, Franco Maria Nardini, Fabio Pinelli, Fabrizio Silvestri, and Nicola Tonellotto (Eds.). Springer, 49–66. doi:10.1007/978-3-031-88708-6_4

[30] Cinzia Di Novi, Matija Kovacic, and Cristina Elisa Orso. 2024. Online health information seeking behavior, healthcare access, and health status during exceptional times. *Journal of Economic Behavior & Organization* 220 (2024), 675–690. doi:10.1016/j.jebo.2024.02.032

[31] Sabrina De Capitani di Vimercati, Sara Foresti, Giovanni Livraga, and Pierangela Samarati. 2012. Data Privacy: Definitions and Techniques. *Int. J. Uncertain. Fuzziness Knowl. Based Syst.* 20 (2012), 793–818. https://api.semanticscholar.org/CorpusID:13161915

[32] Eleanor Wiske Dillon, Sonia Jaffe, Sida Peng, and Alexia Cambon. 2025. Early Impacts of M365 Copilot. *CoRR* abs/2504.11443 (2025). arXiv:2504.11443 doi:10.48550/ARXIV.2504.11443

[33] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *Theory of Cryptography*, Shai Halevi and Tal Rabin (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 265–284.

[34] European Parliament and Council of the European Union. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council. https://data.europa.eu/eli/reg/2016/679/oj

[35] European Parliament and Council of the European Union. 2024. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act). http://data.europa.eu

[36] Guglielmo Faggioli and Nicola Ferro. 2024. Query Obfuscation for Information Retrieval Through Differential Privacy. In *Advances in Information Retrieval - 46th European Conference on Information Retrieval, ECIR 2024, Glasgow, UK, March 24-28, 2024, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 14608)*, Nazli Goharian, Nicola Tonellotto, Yulan He, Aldo Lipani, Graham McDonald, Craig Macdonald, and Iadh Ounis (Eds.). Springer, 278–294. doi:10.1007/978-3-031-56027-9_17

[37] Oluwaseyi Feyisetan, Borja Balle, Thomas Drake, and Tom Diethe. 2020. Privacy- and Utility-Preserving Textual Analysis via Calibrated Multivariate Perturbations. In *Proceedings of the 13th International Conference on Web Search and Data Mining*, James Caverlee, Xia (Ben) Hu, Mounia Lalmas, and Wei Wang (Eds.). ACM, 178–186. doi:10.1145/3336191.3371856

[38] Michèle Finck and Asia Biega. 2021. Reviving Purpose Limitation and Data Minimisation in Personalisation, Profiling and Decision-Making Systems. *CoRR* abs/2101.06203 (2021). arXiv:2101.06203 https://arxiv.org/abs/2101.06203

[39] Maik Fröbe, Eric Oliver Schmidt, and Matthias Hagen. 2021. Efficient Query Obfuscation with Keyqueries. In *WI-IAT '21: IEEE/WIC/ACM International Conference on Web Intelligence, Melbourne VIC Australia, December 14 - 17, 2021*, Jing He, Rainer Unland, Eugene Santos Jr., Xiaohui Tao, Hemant Purohit, Willem-Jan van den Heuvel, John Yearwood, and Jie Cao (Eds.). ACM, 154–161. doi:10.1145/3486622.3493950

[40] Ivan Habernal. 2021. When differential privacy meets NLP: The devil is in the detail. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (Eds.). Association for Computational Linguistics, 1522–1528. doi:10.18653/V1/2021.EMNLP-MAIN.114

[41] Sebastian Hofstätter, Sheng-Chieh Lin, Jheng-Hong Yang, Jimmy Lin, and Allan Hanbury. 2021. Efficiently Teaching an Effective Dense Retriever with Balanced Topic Aware Sampling. In *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11-15, 2021*, Fernando Diaz, Chirag Shah, Torsten Suel, Pablo Castells, Rosie Jones, and Tetsuya Sakai (Eds.). ACM, 113–122. doi:10.1145/3404835.3462891

[42] Yuval Ishai, Eyal Kushnir, and Ron D. Rothblum. 2023. Combinatorially Homomorphic Encryption. In *Theory of Cryptography - 21st International Conference, TCC 2023, Taipei, Taiwan, November 29 - December 2, 2023, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 14370)*, Guy N. Rothblum and Hoeteck Wee (Eds.). Springer, 251–278. doi:10.1007/978-3-031-48618-0_9

[43] Kalervo Järvelin and Jaana Kekäläinen. 2002. Cumulated gain-based evaluation of IR techniques. *ACM Trans. Inf. Syst.* 20, 4 (2002), 422–446. doi:10.1145/582415.582418

[44] Brent Kitchens, Christopher A. Harle, and Shengli Li. 2014. Quality of health-related online search results. *Decis. Support Syst.* 57 (2014), 454–462. doi:10.1016/J.DSS.2012.10.050

[45] Oleksandra Klymenko, Stephen Meisenbacher, and Florian Matthes. 2022. Differential Privacy in Natural Language Processing: The Story So Far. *CoRR* abs/2208.08140 (2022). arXiv:2208.08140 doi:10.48550/ARXIV.2208.08140

[46] Haoran Li, Mingshi Xu, and Yangqiu Song. 2023. Sentence Embedding Leaks More Information than You Expect: Generative Embedding Inversion Attack to Recover the Whole Sentence. In *Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023 (Findings of ACL, Vol. ACL 2023)*, Anna Rogers, Jordan L. Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, 14022–14040. doi:10.18653/V1/2023.FINDINGS-ACL.881

[47] Ninghui Li, Tiancheng Li, and Suresh Venkatasubramanian. 2007. t-Closeness: Privacy Beyond k-Anonymity and l-Diversity. In *Proceedings of the 23rd International Conference on Data Engineering, ICDE 2007, The Marmara Hotel, Istanbul, Turkey, April 15-20, 2007*, Rada Chirkova, Asuman Dogac, M. Tamer Özsu, and Timos K. Sellis (Eds.). IEEE Computer Society, 106–115. doi:10.1109/ICDE.2007.367856

[48] Sheng-Chieh Lin, Akari Asai, Minghan Li, Barlas Oguz, Jimmy Lin, Yashar Mehdad, Wen-tau Yih, and Xilun Chen. 2023. How to Train Your Dragon: Diverse Augmentation Towards Generalizable Dense Retrieval. In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023 (Findings of ACL, Vol. EMNLP 2023)*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, 6385–6400. doi:10.18653/V1/2023.FINDINGS-EMNLP.423

[49] Fan Liu, Zhiyong Cheng, Huilin Chen, Yinwei Wei, Liqiang Nie, and Mohan S. Kankanhalli. 2022. Privacy-Preserving Synthetic Data Generation for Recommendation Systems. In *SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, Madrid, Spain, July 11 - 15, 2022*, Enrique Amigó, Pablo Castells, Julio Gonzalo, Ben Carterette, J. Shane Culpepper, and Gabriella Kazai (Eds.). ACM, 1379–1389. doi:10.1145/3477495.3532044

[50] Allan Luedeman, Nicholas Baum, Andrew Quijano, and Kemal Akkaya. 2024. Privacy-Preserving Drone Navigation Through Homomorphic Encryption for Collision Avoidance. In *2024 IEEE 49th Conference on Local Computer Networks (LCN)*. 1–6. doi:10.1109/LCN60385.2024.10639770

[51] Ashwin Machanavajjhala, Daniel Kifer, Johannes Gehrke, and Muthuramakrishnan Venkitasubramaniam. 2007. L-diversity: Privacy beyond k-anonymity. *ACM Trans. Knowl. Discov. Data* 1, 1 (March 2007), 3–es. doi:10.1145/1217299.1217302

[52] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, AISTATS 2017, 20-22 April 2017, Fort Lauderdale, FL, USA (Proceedings of Machine Learning Research, Vol. 54)*, Aarti Singh and Xiaojin (Jerry) Zhu (Eds.). PMLR, 1273–1282. http://proceedings.mlr.press/v54/mcmahan17a.html

[53] Stephen Meisenbacher, Nihildev Nandakumar, Alexandra Klymenko, and Florian Matthes. 2024. A Comparative Analysis of Word-Level Metric Differential Privacy: Benchmarking the Privacy-Utility Trade-off. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, Nicoletta Calzolari, Min-Yen Kan, Veronique Hoste, Alessandro Lenci, Sakriani Sakti, and Nianwen Xue (Eds.). ELRA and ICCL, Torino, Italia, 174–185. https://aclanthology.org/2024.lrec-main.16

[54] William Melicher, Mahmood Sharif, Joshua Tan, Lujo Bauer, Mihai Christodorescu, and Pedro Giovanni Leon. 2016. (Do Not) Track Me Sometimes: Users' Contextual Preferences for Web Tracking. *Proc. Priv. Enhancing Technol.* 2016, 2 (2016), 135–154. doi:10.1515/POPETS-2016-0009

[55] Janos Meszaros, Davy Preuveneers, Elisabetta Biasin, Enzo Marquet, Irmak Erdogan, Isabela Maria Rosal, Koen Vranckaert, Lydia Belkadi, and Natalia Menendez. 2024. ChatGPT, Are You Lawfully Processing My Personal Data? GDPR Compliance and Legal Basis for Processing Personal Data by OpenAI. (2024). doi:10.21552/aire/2024/2/10

[56] Arvind Narayanan and Vitaly Shmatikov. 2008. Robust De-anonymization of Large Sparse Datasets. In *2008 IEEE Symposium on Security and Privacy (SP 2008), 18-21 May 2008, Oakland, California, USA*. IEEE Computer Society, 111–125. doi:10.1109/SP.2008.33

[57] Helen Nissenbaum. 2004. Privacy as contextual integrity. *Washington Law Review* 79 (2004), 119–157. https://api.semanticscholar.org/CorpusID:150528892

[58] National Institute of Standards and Technology. 2017. *NIST SP 800-130 Framework for Designing Cryptographic Key Management Systems: NIST SP 800-130 Aug 2013.* CreateSpace Independent Publishing Platform, North Charleston, SC, USA.

[59] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global Vectors for Word Representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*,

Alessandro Moschitti, Bo Pang, and Walter Daelemans (Eds.). ACL, 1532–1543. doi:10.3115/v1/d14-1162

[60] Fabio Pinelli, Gabriele Tolomei, and Giovanni Trappolini. 2023. FLIRT: Federated Learning for Information Retrieval. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2023, Taipei, Taiwan, July 23-27, 2023*, Hsin-Hsi Chen, Wei-Jou (Edward) Duh, Hen-Hsen Huang, Makoto P. Kato, Josiane Mothe, and Barbara Poblete (Eds.). ACM, 3472–3475. doi:10.1145/3539618.3591926

[61] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics. https://arxiv.org/abs/1908.10084

[62] Mouna Rhahla, Sahar Allegue, and Takoua Abdellatif. 2021. Guidelines for GDPR compliance in Big Data systems. *J. Inf. Secur. Appl.* 61 (2021), 102896. doi:10.1016/J.JISA.2021.102896

[63] Michael L. Dertouzos Ronald L. Rivest, Len Adelman. 1978. ON DATA BANKS AND PRIVACY HOMOMORPHISMS. https://api.semanticscholar.org/CorpusID:6905087

[64] Nativi S and De Nigris S. 2021. AI Standardisation Landscape: state of play and link to the EC proposal for an AI regulatory framework. KJ-NA-30772-EN-N (online) (2021). doi:10.2760/376602(online)

[65] Bharath Kumar Samanthula, Wei Jiang, and Elisa Bertino. 2014. Privacy-Preserving Complex Query Evaluation over Semantically Secure Encrypted Data. In *Computer Security - ESORICS 2014 - 19th European Symposium on Research in Computer Security, Wroclaw, Poland, September 7-11, 2014. Proceedings, Part I (Lecture Notes in Computer Science, Vol. 8712)*, Miroslaw Kutylowski and Jaideep Vaidya (Eds.). Springer, 400–418. doi:10.1007/978-3-319-11203-9_23

[66] Naw Safrin Sattar, Muhammad Abdullah Adnan, and Maimuna Begum Kali. 2017. Secured aerial photography using Homomorphic Encryption. In *International Conference on Networking, Systems and Security, NSysS 2017, Dhaka, Bangladesh, January 5-8, 2017*, Mohammad Zulkernine and Mohammad Mahfuzul Islam (Eds.). IEEE, 107–114. doi:10.1109/NSYSS.2017.7885810

[67] Mahmoud F. Sayed, William Cox, Jonah Lynn Rivera, Caitlin Christian-Lamb, Modassir Iqbal, Douglas W. Oard, and Katie Shilton. 2020. A Test Collection for Relevance and Sensitivity. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, July 25-30, 2020*, Jimmy X. Huang, Yi Chang, Xueqi Cheng, Jaap Kamps, Vanessa Murdock, Ji-Rong Wen, and Yiqun Liu (Eds.). ACM, 1605–1608. doi:10.1145/3397271.3401284

[68] Mahmoud F. Sayed and Douglas W. Oard. 2019. Jointly Modeling Relevance and Sensitivity for Search Among Sensitive Content. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2019, Paris, France, July 21-25, 2019*, Benjamin Piwowarski, Max Chevalier, Éric Gaussier, Yoelle Maarek, Jian-Yun Nie, and Falk Scholer (Eds.). ACM, 615–624. doi:10.1145/3331184.3331256

[69] Tanusree Sharma, Lin Kyi, Yang Wang, and Asia J. Biega. 2024. "I'm not convinced that they don't collect more than is necessary": User-Controlled Data Minimization Design in Search Engines. In *33rd USENIX Security Symposium, USENIX Security 2024, Philadelphia, PA, USA, August 14-16, 2024*, Davide Balzarotti and Wenyuan Xu (Eds.). USENIX Association. https://www.usenix.org/conference/usenixsecurity24/presentation/sharma

[70] Luo Si and Hui Yang. 2014. Privacy-preserving IR: when information retrieval meets privacy and security. In *The 37th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '14, Gold Coast , QLD, Australia - July 06 - 11, 2014*, Shlomo Geva, Andrew Trotman, Peter Bruza, Charles L. A. Clarke, and Kalervo Järvelin (Eds.). ACM, 1295. doi:10.1145/2600428.2600737

[71] Samuel Sousa and Roman Kern. 2023. How to keep text private? A systematic review of deep learning methods for privacy-preserving natural language processing. *Artif. Intell. Rev.* 56, 2 (2023), 1427–1492. doi:10.1007/S10462-022-10204-6

[72] Latanya Sweeney. 2002. k-Anonymity: A Model for Protecting Privacy. *Int. J. Uncertain. Fuzziness Knowl. Based Syst.* 10, 5 (2002), 557–570. doi:10.1142/S0218488502001648

[73] Gemini Team. 2023. Gemini: A Family of Highly Capable Multimodal Models. *CoRR* abs/2312.11805 (2023). arXiv:2312.11805 doi:10.48550/ARXIV.2312.11805

[74] Antonios Tragoudaras, Theofanis Aslanidis, Emmanouil Georgios Lionis, Marina Orozco González, and Panagiotis Eustratiadis. 2025. Information Leakage of Sentence Embeddings via Generative Embedding Inversion Attacks. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Padua, Italy) *(SIGIR '25)*. Association for Computing Machinery, New York, NY, USA, 3234–3243. doi:10.1145/3726302.3730303

[75] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing Data using t-SNE. *Journal of Machine Learning Research* 9, 86 (2008), 2579–2605. http://jmlr.org/papers/v9/vandermaaten08a.html

[76] Isabel Wagner and David Eckhoff. 2018. Technical Privacy Metrics: A Systematic Survey. *ACM Comput. Surv.* 51, 3 (2018), 57:1–57:38. doi:10.1145/3168389

[77] Shuyi Wang, Bing Liu, Shengyao Zhuang, and Guido Zuccon. 2021. Effective and Privacy-preserving Federated Online Learning to Rank. In *ICTIR '21: The 2021 ACM SIGIR International Conference on the Theory of Information Retrieval, Virtual Event, Canada, July 11, 2021*, Faegheh Hasibi, Yi Fang, and Akiko Aizawa (Eds.). ACM, 3–12. doi:10.1145/3471158.3472236

[78] Benjamin Weggenmann and Florian Kerschbaum. 2018. SynTF: Synthetic and Differentially Private Term Frequency Vectors for Privacy-Preserving Text Mining. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval* (Ann Arbor, MI, USA) *(SIGIR '18)*. Association for Computing Machinery, New York, NY, USA, 305–314. doi:10.1145/3209978.3210008

[79] Zekun Xu, Abhinav Aggarwal, Oluwaseyi Feyisetan, and Nathanael Teissier. 2020. A Differentially Private Text Perturbation Method Using Regularized Mahalanobis Metric. In *Proceedings of the Second Workshop on Privacy in NLP*. Association for Computational Linguistics. doi:10.18653/v1/2020.privatenlp-1.2

[80] Zekun Xu, Abhinav Aggarwal, Oluwaseyi Feyisetan, and Nathanael Teissier. 2021. On a Utilitarian Approach to Privacy Preserving Text Generation. *CoRR* abs/2104.11838 (April 2021). arXiv:2104.11838 [cs.CL] doi:10.48550/ARXIV.2104.11838

[81] Grace Hui Yang, Ian Soboroff, Li Xiong, Charles L. A. Clarke, and Simson L. Garfinkel. 2016. Privacy-Preserving IR 2016: Differential Privacy, Search, and Social Media. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval, SIGIR 2016, Pisa, Italy, July 17-21, 2016*, Raffaele Perego, Fabrizio Sebastiani, Javed A. Aslam, Ian Ruthven, and Justin Zobel (Eds.). ACM, 1247–1248. doi:10.1145/2911451.2917763

[82] Hui Yang and Ian Soboroff. 2015. Privacy-Preserving IR 2015: When Information Retrieval Meets Privacy and Security. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, Santiago, Chile, August 9-13, 2015*, Ricardo Baeza-Yates, Mounia Lalmas, Alistair Moffat, and Berthier A. Ribeiro-Neto (Eds.). ACM, 1157–1158. doi:10.1145/2766462.2767857

[83] Xiang Yue, Minxin Du, Tianhao Wang, Yaliang Li, Huan Sun, and Sherman S. M. Chow. 2021. Differential Privacy for Text Analytics via Natural Text Sanitization. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (Eds.). Association for Computational Linguistics, Online, 3853–3866. doi:10.18653/v1/2021.findings-acl.337