

PRIVACY IN OPEN SEARCH: A REVIEW OF CHALLENGES AND SOLUTIONS

Samuel Sousa*¹, Roman Kern¹, Know-Center GmbH, Graz, Austria

¹also at ISDS, Graz University of Technology, Graz, Austria

Christian Guetl, ISDS, Graz University of Technology, Graz, Austria

Abstract

Privacy is of worldwide concern regarding activities and processes that include sensitive data. For this reason, many countries and territories have been recently approving regulations controlling the extent to which organizations may exploit data provided by people. Artificial intelligence areas, such as machine learning and natural language processing, have already successfully employed privacy-preserving mechanisms in order to safeguard data privacy in a vast number of applications. Information retrieval (IR) is likewise prone to privacy threats, such as attacks and unintended disclosures of documents and search history, which may cripple the security of users and be penalized by data protection laws. This work aims at highlighting and discussing open challenges for privacy in the recent literature of IR, focusing on tasks featuring user-generated text data. Our contribution is threefold: firstly, we present an overview of privacy threats to IR tasks; secondly, we discuss applicable privacy-preserving mechanisms which may be employed in solutions to restrain privacy hazards; finally, we bring insights on the tradeoffs between privacy preservation and utility performance for IR tasks.

Index terms— Privacy, information retrieval, personal information, open search.

INTRODUCTION

Data is the cornerstone of many research fields, as well as the source of information used by professionals, such as journalists, statisticians, and police makers [1]. At the pace of the Internet popularization, the number of data publishers rose to the thousands [2], encompassing university library files, government open data, Wikipedia articles, social media platforms, commercial data providers, digital data markets, scientific data repositories, among many others. Accessing these data sources is crucial for solving the reproducibility issues of scientific results and improving the means for data journalists to obtain reliable information [1]. Furthermore, the idea of open science relies on sharing research data and materials for third-parties to reuse and reproduce experiments, assuring the trustworthiness of the research process [2]. Some companies also need to have their products openly visible on Internet for their business to keep going, e.g., e-commerce platforms. However, privacy issues may keep some kinds of data from being released in order to avoid the exposure of confidential information. As an example, user-generated data, such as user behavior, search

interests, and search profiles, demands anonymization steps prior to its release. Otherwise, it should remain unrevealed. Other privacy issue regards the level playing field for search services since those which monitor the intents of users can eventually obtain benefits out of it.

Privacy is a concept related to limiting the extend of information an individual is willing to share [3]. Many countries consider privacy as a right protected by law. For instance, the General Data Protection Regulation (GDPR) [4], which entered into force in May 2018 in the European Union (EU), draws guidelines for the collection, transfer, storage, management, and deletion of personal data within the member states of the economic bloc. GDPR grants the residents of the EU the control over their personal data. Therefore, any processing activities over personal data must comply with the regulation and provide protection measures. In case of data breaches or noncompliance with the legal principles, penalties and fines are applicable. Since 2018, data protection bills have also been passed in several countries worldwide.

In the recent past years, the preservation of privacy has been gaining attention in the field of information retrieval (IR) [5–8]. Several privacy-preserving mechanisms have been developed to safeguard personal data from threats, as attacks, disclosures, and unintended usages. Some of these mechanisms, such as encryption [9, 10], differential privacy (DP) [11], multi-party computation (MPC) [12], and federated learning (FL) [13, 14], are implemented alongside models for enabling applications to safeguard data privacy. There are also some attacking methods that aim at retrieving data samples used to train models, mostly neural networks, which can be seen as a threat or a safety checker, depending on the attacker’s intention. Furthermore, some privacy-preserving tools may present computational overheads or performance reductions, which call attention to privacy-utility trade-offs.

Searches on personal data often present privacy risks from both data provider and model sides. There is a need for open data, due to scientific, governmental, and press reasons [1, 2], however data generators and IR systems users cannot have their identities and personal data exposed. This work aims, therefore, at pointing out open challenges with regards to privacy in IR tasks, as well as reviewing appropriate privacy-preserving methods to safeguard personal data or model. We focus primarily on tasks featuring text data since the private content of data in written format may be presented explicitly, as a person’s name or an ID number, or implicitly like a the gender of a person which can be inferred from the text, based on gendered words like profession terms.

* ssousa@know-center.at

Our contributions comprise a summary of research directions on privacy for IR tasks alongside adequate privacy-preserving methods for the privacy risks these tasks may present. Problems that put privacy at risk are also discussed. Moreover, we discuss how privacy-preserving methods can influence the results of IR tasks. Finally, we provide the readers with essential understanding of privacy-related issues and privacy-preserving methods for IR.

This paper is organized as follows. Firstly, we review recent works in the literature of privacy in section *Related Work*. Secondly, we describe open challenges and solutions for privacy preservation in open search tasks in section *Privacy Challenges and Solutions*. Further, we discuss our results in the section *Discussion*. Finally, our contributions and upcoming works are brought into context in the section *Conclusion and Future Works*.

RELATED WORK

Data privacy is a large concern within IR [6–8] since privacy threats and risks often arise from searches on private data [6] and tasks which involve the search behavior or intent of users [7, 8, 15]. Examples of such issues include private data publishing [6], string searches [16], user’s context [8], information sharing of web search logs [17], interactive search [15], and the information selection behavior of users [15]. Additional privacy threats include revealing private data from IR systems that compute distributed information, performing face recognition without the consent of the people whose faces are captured, and so on. Therefore, a large number of real-world IR systems can be prone to breaches of personal data, unintended disclosures of private information, and penalties established by data protection regulations.

Nowadays, many organizations collect and process personal data, such as governments, companies, and search service providers. As a consequence, the volume of collected data has increased alongside the risks related to breaches of sensitive information from these datasets. Zhu et al. [6] survey DP applications for data publishing and data analysis. The authors define data publishing as publicly sharing data itself or the result of queries, whereas data analysis consists on releasing data models to the public. In both scenarios, DP can offer privacy guarantees resistant against attacks and mathematically provable. Riazi et al. [16] come up with a mathematically provable mechanism for privacy preservation in IR tasks. The authors implement a methodology for two-party string search based on the Yao’s Garbled Circuit protocol. For instance, two users can hold queries and data for string search simultaneously, whereas they both aim at keeping these search components private without relying on a third party like a trusted server. Therefore, the proposed protocol converts the search algorithm into a Boolean circuit that evaluates the private queries and texts. Tamine and Daoud [8] survey methodologies and metrics for context IR evaluation, specifying the impact of data privacy towards the evaluation design.

User search behavior can also be seen as private information since queries and search result’s selection can disclose demographic attributes, preferences, political views, etc. Orso et al. [15] investigate the role of user search behavior and information selection in order to understand which layers of social information, namely personal preferences, tags combined to personal preferences, or tags and social ratings combined with personal preferences, can enhance search efficiency. The authors found empirical evidence that personalized preferences and social ratings make it easier for users to select information without external sources. In this study, the publicly available Yelp dataset¹ was used. Therefore, privacy aspects that would arise from this use case in a real-world scenario, as the compliance with data protection regulations, were not approached. Sharing search log records is a process ruled by data protection laws, which demand anonymization techniques to be applied before the data leaves the servers it is stored on. Mivule et al. [17] introduce an heuristic for privacy preservation of individual web search log records based on swapping. Firstly, an individual has a set of logs \mathcal{A} . Secondly, the records in \mathcal{A} are switched within this set. Finally, the swapped files from \mathcal{A} are switched again using records of a set \mathcal{B} . This heuristic is efficient when it comes to preventing an attacker from tracing the issuer of a search query. Additional privacy-preserving IR solutions include homomorphic encryption (HE) [18], DP [19], and hashing [7].

Our work focuses on identifying privacy risks across IR tasks, alongside suggesting privacy-preserving mechanisms that have the potential to suit the needs for privacy protection. We briefly introduce the task description followed by the privacy risks associated with both data and IR model.

PRIVACY CHALLENGES AND SOLUTIONS

This section firstly introduces open challenges with regards to privacy in IR tasks. We survey recently published works for ten IR tasks, highlighting their privacy issues. Afterwards, we overview privacy-preserving methods that are suitable options to address these problems.

Privacy Challenges

Search tasks

Ad-hoc search. Modern search engines often rely on bag-of-words models to represent documents and search queries. Consequently, accurate quantification of context-specific term importance in documents is a tricky problem since term’s context is often not captured by these models. When it comes to data privacy, bag-of-words models pose risks related to data re-identification. For instance, some terms in the vocabulary of the model may refer to personal identities or private attributes, such as age, gender, location, and demographic information, which allow a de-identified document to be re-identified. Dai and Callan [20] propose a novel

¹ <https://www.yelp.com/dataset>.

document term weighting framework which preserves word context, using BERT [21] embeddings to extract document representations. Although keeping document’s contextual information, BERT embeddings can also encode private information and suffer attacks as model inversion, reverse engineering, and membership inference.

Query expansion. Query expansion (QE) helps users of IR applications to find more relevant information by expanding the search queries, hence increasing recall. For instance, synonyms and hypernyms of terms in the questions for question answering (QA) are used to rise the likelihood of matching sentences stating the most appropriate candidate answer [22,23]. However, in some QA scenarios the number of retrieved sentences may be small, and then mismatch the intents of the user who queries the QA system [23]. Common privacy threats in applications which use QE regard disclosing query content or private information in the documents in which the sentences are extracted, such as names and locations of people.

Feature extraction for ranking. In the context of ad-hoc search, document ranking consists on returning a ranked list of documents from a large collection based on the assumed information need expressed by a search query, maximizing some ranking metric like average precision [24]. Ranking is a cornerstone for IR systems and search engines, which can also be performed by machine learning (ML) classification models [24,25]. Therefore, in order to reduce computation time, boost learning results, and prevent overfitting for those models, feature extraction techniques can be implemented to better represent documents. Pandey et al. [25] come up with a method for representing documents as matrices with reduced dimensions when compared to the original document representations. Such representations are useful for improving the results of ranking algorithms. However, feature extraction models can have its original training samples disclosed by attacks of model inversion, membership inference, or reverse engineering.

Online learning for ranking. A critical drawback of ranking models regards the hardness to obtain labelled data for model training. Thus, Zhuang and Zucco [26] propose a counterfactual learning to rank method based on logs of user interaction collected from the ranking model in production. In a real-world setting, users would be able to confirm the effectiveness of the ranking model by clicking on the results. In the experimental evaluation, the authors use publicly available datasets and generate user clicks automatically with a cascade click model [27]. This scenario can pose privacy threats to personal data collection if the user clicks are collected from actual system users. Therefore, regulations for personal data collection would be applicable alongside the need for privacy-preserving methods.

Query composition. ML algorithms can be used to predict query properties like answer size, run-time, and error

class [28]. These algorithms can therefore be prone to unintended memorization of query content alongside the attacks which aim at disclosing training document samples. Another privacy threat regards the location of the ML model, e.g., on a cloud server, since computation parties sometimes may not be trusted, or the communication channels for transferring data or model updates may allow eavesdropping attacks to take place.

Healthcare tasks

Healthcare data tasks. Electronic medical records or electronic health records are digital documents, in which medical staff inputs patient data, including personal information, health condition, disease diagnosis, medication, etc. [29]. This type of document has the advantages of easy storage, transfer, sharing, and deletion. However, healthcare data is inherently private due to the sensitivity of its content. Therefore data protection regulations, as the EU’s GDPR [4], prevents publicly releasing such data for research activities and public searches, without the explicit consent of data owners and the application of data anonymization methods. In the context of IR, medical applications have to safeguard medical data from queries by malicious users or computation parties, e.g., corrupted servers.

Social media tasks

Opinion mining. Many online data sources contain opinions, which can be classified with regards to their polarity. For instance, Nguyen and Nguyen [30] predict sentiment polarity on YouTube² comments in English and Italian languages, using a DL model based on a Bidirectional Long Short-Term Memory architecture. Tasks performed over user-generated data pose privacy risks of unintended data memorization by the neural network, as well as sensitivity to attacks that aim at retrieving the model training samples. Therefore, there is a room for privacy-preserving mechanisms that prevent such privacy issues.

Advisor for hashtag sharing. On social media platforms, users often put their privacy at risk unconsciously by releasing details of their personal lives publicly, revealing their exact location, and posting political or societal views. Furthermore, some features of such platforms like hashtags may induce privacy threats, mainly related to location, since attacking models can easily predict precise user location from the hashtags they post online [5]. This situation can be prevented by privacy-preserving mechanisms of data obfuscation or de-identification, which get rid of privacy-sensitive information before any sharing step by the user. Moreover, social media data is frequently protected by data protection regulations, then, ensuring data privacy to be safeguarded.

² <http://youtube.com>.

Social media profile linking. User identity linking is the task of connecting accounts owned by the same user across different social media platforms [31]. For instance, a user can have profile simultaneously on TikTok³, Twitter⁴, Instagram⁵, and so on. However, more relevant advertisement can be suggested to this user by linking these different profiles. This scenario, therefore, presents privacy threats related to monitoring user behavior online besides linking profiles where the user uses a pseudonym to remain anonymous.

Recommendation tasks

Recommender systems. Recommender systems aim at predicting the preferences of users based on their interest, making more effective use of information [32]. However, user interest and their search history should be preserved from leakages, hence these are pieces of private information. A recent application of recommender systems is news recommendation. Qi et al. [33] propose a framework to recommend news to users in distributed computation scenarios, e.g., smartphones running a mobile application. This framework computes updates locally on each distributed device separately and, then, sends these local updates to a centralized server which updates the global model by aggregating the local updates. Finally, the updated global model is distributed over the distributed devices to news recommendation and successive updates. This is an outstanding application of FL, which prevents threats of data leakage from the distributed devices. However, this model still can suffer from unintended memorization of user’s behavior, preferences, and search history.

Solutions

Many privacy-preserving methods have been proposed for the sake of preventing attacks and unintended data breaches. In this subsection, we review computational techniques for privacy preservation, which can be conveniently integrated into IR tasks. We group these techniques according to their technical aspects into cryptographic approaches and ML-based applications.

Cryptographic approaches Cryptographic protocols have been extensively used to protect private data when sharing activities are not advised. In a nutshell, encryption consists on the application of a function over data in raw format, which is referred as ‘*plaintext*’. This function results on an output called ‘*cyphertext*’, which inhibits the identification of its original format or content.

Encryption. Homomorphic encryption (HE) is a form of encryption that allows arithmetic operations to be computed over cyphertexts without the need for decryption [9]. For instance, ML models can be used for inference on encrypted data, whereas the results are still consistent. Encryption

schemes that implement HE are particularly useful for scenarios in which data transfers and centralized storage occur, e.g., cloud servers, alongside the lack of trust. However, it is worth to mention the computational overheads that FHE often leads to, hence its use becomes prohibitive for devices with small memory capacity.

IR models can also be privacy-preserving by the use of searchable encryption (SE). This technique encrypts document collections enabling the data owner to delegate search capabilities, whereas the server or a service provider, like a search engine, does not demand decryption [10]. Therefore, the so-called ‘honest-but-curious’ server can provide searches, whereas the content of the stored data and the input queries is preserved [34]. On the other hand, semantic relations between words and documents may be lost in the encrypted forms, so that decaying search results [35]. SE usually encompasses algorithms for the steps of key generation, encryption, token generation, and search [34]. Finally, the main threat this encryption scheme faces is related to keyword inference attacks that aim at recovering the content of encrypted keywords.

Multi-party computation. Document collections may be stored in distributed search system hosted by other companies or even distributed across members of a broad community. Therefore, when the exchange of documents among the members of these computation scenario is not an advisable option, multi-party computation (MPC) can be successfully used. MPC is a cryptographic primitive that computes aggregated functions over multiple sources of data, which cannot be revealed [12]. Formally, MPC assumes a set of inputs $\{x_1, x_2, \dots, x_n\}$ so that each party P_i will store x_i and agree to compute $y = \hat{f}(x_1, x_2, \dots, x_n)$, in which y is the output information to be released, and \hat{f} is the agreed function on the entire input set [36]. The input set may be composed of keywords, documents, medical records, etc.

Differential privacy DP can be understood as a randomized function \hat{k} which is applied to document collections or query results prior to their public release [11]. Therefore, for all subsets S in the range of \hat{k} , and document collections D and D' differing on at most one element, \hat{k} provides ϵ -differential privacy if:

$$Pr[\hat{k}(D) \in S] \leq \exp(\epsilon) Pr[\hat{k}(D') \in S]. \quad (1)$$

Approaches for DP exploit mathematical formalism to neutralize de-anonymization attacks and keep a lookout for membership inference attacks that may disclose the original documents in the collection. The function \hat{k} adds random noise to any input query and, consequently, yields the response [11]. Every mechanism that satisfies ϵ -DP will mitigate risks of leakages of private information from any individual element since its inclusion or removal from the collection would not turn the output significantly more or less likely [11]. DP provides privacy guarantees usually at the cost of performance and computational overhead. However, one of the

³ <https://www.tiktok.com/>.

⁴ <https://twitter.com/>.

⁵ <https://www.instagram.com/>.

main advantages of this method regards managing the trade-off between privacy and utility, finding the ideal value of ϵ which preserves data privacy and affects the model results at a controlled extent.

ML-based approaches

Federated learning. FL is a methodology for training ML models in distributed computation scenarios proposed by Google [37]. This methodology prevents data from leaving its owner’s device during the computations and addresses privacy threats related to training over private data [14]. The federated training consists on, firstly, distributing copies of a global model with pre-defined parameters, computing local updates on the distributed devices, sending these updates to the server for aggregation, updating the global model parameters, and sending the updated parameters of the global model to each distributed device [13]. In formal terms, FL assumes a model \hat{m} with parameters $\Theta_{\hat{m}}$ which are stored in a matrix M . The model \hat{m} is thus shared with a subset T of η clients, which will update \hat{m} with their locally stored data at each training step $t \geq 0$ [13, 37]. Every client i will send its update $H_t^i := M_t^i - M_t$ to the central server, which is responsible for aggregating client-side updates for updating global model [13]. FL has advantages for enabling distributed computations over devices with restrained memory, bandwidth, and computation power.

DISCUSSION

A large number of IR systems are sensitive to privacy threats in scenarios which include personal information, search history, personal preferences, private documents, to name a few. Additionally, regulations as the EU’s GDPR establish the guidelines for protecting user generated data from breaches and non-consented usages. As a consequence, many algorithmic methods for protecting data privacy have been proposed and integrated into IR systems. However, protecting privacy using such methods can mean coping with utility performance decays and computational overheads. Therefore, the choice for a convenient privacy-preserving method for an IR scenario has to take into account the target of privacy protection and the computational resources available.

In Table 1, we bring a summary of IR applications for the privacy-preserving methods surveyed in the section *Solutions*. Cryptographic approaches, such as HE, SE, and MPC, are suitable to scenarios in which the original content of documents or datasets should not be revealed to unauthorized parties or some of these parties are not trusted, e.g., malicious servers. DP provides formal privacy guarantees which can be employed on a myriad of applications, such as data anonymization and protecting results of queries and ML models against inversion or reverse engineering. Finally, FL can be successfully implemented for scenarios with distributed devices and limited resources for data sharing or prohibitions of data exchange.

Table 1: Applications of IR alongside suitable privacy-preserving methods.

IR application	PP methods
Ad-hoc search	SE, MPC
Query expansion	HE, SE
Feature extraction for ranking	DP
Online learning for ranking	DP, HE
Query composition	SE
Healthcare data tasks	DP, HE, SE
Opinion mining	DP
Advisor for hashtag sharing	SE
Social media profile linking	HE, SE
Recommender systems	FE, MPC

CONCLUSION AND FUTURE WORKS

Privacy is a critical point for the development of IR systems which deal with personal, user generated, or sensitive data. In this work we overview recent developments in the IR literature, pointing out privacy issues and suggesting suitable privacy-preserving methods. Data types, IR tasks, and privacy-preserving method drawbacks are taken into account to provide the reader with essential understating of this research field. As future works, we aim to address the aforementioned challenges for Open Search use cases, as well as studying and discussing compliance with legal requirements, such as those of the EU’s GDPR.

REFERENCES

- [1] D. Brickley, M. Burgess, and N. Noy, “Google dataset search: Building a search engine for datasets in an open web ecosystem,” in *The World Wide Web Conference*, 2019, pp. 1365–1375.
- [2] D.-L. Magazine, “The landscape of research data repositories in 2015: A re3data analysis,” *D-Lib Magazine*, vol. 23, no. 3/4, 2017.
- [3] A. F. Westin, “Privacy and freedom,” *Washington and Lee Law Review*, vol. 25, no. 1, p. 166, 1968.
- [4] E. Commission, “2018 reform of eu data protection rules,” https://ec.europa.eu/commission/sites/beta-political/files/data-protection-factsheet-changes_en.pdf, 2018, date: 2018-05-25, URL Date: 2019-06-17.
- [5] Y. Zhang, M. Humbert, T. Rahman, C.-T. Li, J. Pang, and M. Backes, “Tagvisor: A privacy advisor for sharing hashtags,” in *Proceedings of the 2018 World Wide Web Conference*, 2018, pp. 287–296.
- [6] T. Zhu, G. Li, W. Zhou, and S. Y. Philip, “Differentially private data publishing and analysis: A survey,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 29, no. 8, pp. 1619–1638, 2017.
- [7] L. Weng, L. Amsaleg, A. Morton, and S. Marchand-Maillet, “A privacy-preserving framework for large-scale content-based information retrieval,” *IEEE Transactions on Information Forensics and Security*, vol. 10, no. 1, pp. 152–167, 2014.

- [8] L. Tamine and M. Daoud, "Evaluation in contextual information retrieval: foundations and recent advances within the challenges of context dynamicity and data privacy," *ACM Computing Surveys (CSUR)*, vol. 51, no. 4, pp. 1–36, 2018.
- [9] C. Gentry, "Fully homomorphic encryption using ideal lattices," in *Proceedings of the forty-first annual ACM symposium on Theory of computing*, 2009, pp. 169–178.
- [10] D. Cash, P. Grubbs, J. Perry, and T. Ristenpart, "Leakage-abuse attacks against searchable encryption," in *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*, 2015, pp. 668–679.
- [11] C. Dwork, "Differential privacy: A survey of results," in *International conference on theory and applications of models of computation*. Springer, 2008, pp. 1–19.
- [12] Q. Feng, D. He, Z. Liu, H. Wang, and K.-K. R. Choo, "SecureNLP: A system for multi-party privacy-preserving natural language processing," *IEEE Transactions on Information Forensics and Security*, 2020.
- [13] J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, "Federated learning: Strategies for improving communication efficiency," in *NIPS Workshop on Private Multi-Party Machine Learning*, 2016.
- [14] M. Chen, A. T. Suresh, R. Mathews, A. Wong, C. Allauzen, F. Beaufays, and M. Riley, "Federated learning of n-gram language models," pp. 121–130, 2019.
- [15] V. Orso, T. Ruotsalo, J. Leino, L. Gamberini, and G. Jacucci, "Overlaying social information: The effects on users' search and information-selection behavior," *Information Processing & Management*, vol. 53, no. 6, pp. 1269–1286, 2017.
- [16] M. S. Riazi, E. M. Songhori, and F. Koushanfar, "Prisearch: Efficient search on private data," in *2017 54th ACM/EDAC/IEEE Design Automation Conference (DAC)*. IEEE, 2017, pp. 1–6.
- [17] K. Mivule, "Data swapping for private information sharing of web search logs," *Procedia computer science*, vol. 114, pp. 149–158, 2017.
- [18] A. El-Ansari, A. Beni-Hssane, M. Saadi, and M. El Fissaoui, "Papir: privacy-aware personalized information retrieval," *Journal of Ambient Intelligence and Humanized Computing*, pp. 1–17, 2021.
- [19] G. H. Yang and S. Zhang, "Differential privacy for information retrieval," in *Proceedings of the ACM SIGIR International Conference on Theory of Information Retrieval*, 2017, pp. 325–326.
- [20] Z. Dai and J. Callan, "Context-aware document term weighting for ad-hoc search," in *Proceedings of The Web Conference 2020*, 2020, pp. 1897–1907.
- [21] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 2019, pp. 4171–4186.
- [22] M. Esposito, E. Damiano, A. Minutolo, G. De Pietro, and H. Fujita, "Hybrid query expansion using lexical resources and word embeddings for sentence retrieval in question answering," *Information Sciences*, vol. 514, pp. 88–105, 2020.
- [23] C. Carpineto and G. Romano, "A survey of automatic query expansion in information retrieval," *Acm Computing Surveys (CSUR)*, vol. 44, no. 1, pp. 1–50, 2012.
- [24] R. Nogueira, Z. Jiang, R. Pradeep, and J. Lin, "Document ranking with a pretrained sequence-to-sequence model," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, 2020, pp. 708–718.
- [25] G. Pandey, Z. Ren, S. Wang, J. Vejjalainen, and M. de Rijke, "Linear feature extraction for ranking," *Information Retrieval Journal*, vol. 21, no. 6, pp. 481–506, 2018.
- [26] S. Zhuang and G. Zuccon, "Counterfactual online learning to rank," in *European Conference on Information Retrieval*. Springer, 2020, pp. 415–430.
- [27] A. Chuklin, I. Markov, and M. d. Rijke, "Click models for web search," *Synthesis lectures on information concepts, retrieval, and services*, vol. 7, no. 3, pp. 1–115, 2015.
- [28] Z. Zolaktaf, M. Milani, and R. Pottinger, "Facilitating sql query composition and analysis," in *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, 2020, pp. 209–224.
- [29] W. Sun, Z. Cai, Y. Li, F. Liu, S. Fang, and G. Wang, "Data processing and text mining technologies on electronic medical records: a review," *Journal of healthcare engineering*, vol. 2018, 2018.
- [30] H. T. Nguyen and M. Le Nguyen, "Multilingual opinion mining on youtube—a convolutional n-gram bilstm word embedding," *Information Processing & Management*, vol. 54, no. 3, pp. 451–462, 2018.
- [31] A. T. Hadgu and J. K. R. Gundam, "Learn2link: Linking the social and academic profiles of researchers," in *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 14, 2020, pp. 240–249.
- [32] H. Wang, N. Wang, and D.-Y. Yeung, "Collaborative deep learning for recommender systems," in *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, 2015, pp. 1235–1244.
- [33] T. Qi, F. Wu, C. Wu, Y. Huang, and X. Xie, "Privacy-preserving news recommendation model learning," in *Findings of the Association for Computational Linguistics: EMNLP 2020*. Online: Association for Computational Linguistics, Nov. 2020, pp. 1423–1432. [Online]. Available: <https://www.aclweb.org/anthology/2020.findings-emnlp.128>
- [34] H. Liu and B. Wang, "Mitigating file-injection attacks with natural language processing," in *Proceedings of the Sixth International Workshop on Security and Privacy Analytics*, 2020, pp. 3–13.
- [35] X. Dai, H. Dai, G. Yang, X. Yi, and H. Huang, "An efficient and dynamic semantic-aware multikeyword ranked search scheme over encrypted cloud data," *IEEE Access*, vol. 7, pp. 142 855–142 865, 2019.
- [36] R. Cramer, I. B. Damgård, and J. B. Nielsen, *Secure Multi-party Computation and Secret Sharing*. Cambridge University Press, 2015.
- [37] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Artificial Intelligence and Statistics*. PMLR, 2017, pp. 1273–1282.