



Robust Information Retrieval

Yu-An Liu

Ruqing Zhang

CAS Key Lab of Network Data
Science and Technology, ICT, CAS
University of Chinese Academy of
Sciences
Beijing, China
{liuyuan21b,zhangruqing}@ict.ac.cn

Jiafeng Guo

CAS Key Lab of Network Data
Science and Technology, ICT, CAS
University of Chinese Academy of
Sciences
Beijing, China
guojiafeng@ict.ac.cn

Maarten de Rijke

University of Amsterdam
Amsterdam, The Netherlands
m.derijke@uva.nl

Abstract

Beyond effectiveness, the robustness of an information retrieval (IR) system is increasingly attracting attention. When deployed, a critical technology such as IR should not only deliver strong performance on average but also have the ability to handle a variety of exceptional situations. In recent years, research into the robustness of IR has seen significant growth, with numerous researchers offering extensive analyses and proposing myriad strategies to address robustness challenges. In this tutorial, we first provide background information covering the basics and a taxonomy of robustness in IR. Then, we examine adversarial robustness and out-of-distribution (OOD) robustness within IR-specific contexts, extensively reviewing recent progress in methods to enhance robustness. The tutorial concludes with a discussion on the robustness of IR in the context of large language models (LLMs), highlighting ongoing challenges and promising directions for future research. This tutorial aims to generate broader attention to robustness issues in IR, facilitate an understanding of the relevant literature, and lower the barrier to entry for interested researchers and practitioners.

CCS Concepts

• Information systems → Information retrieval.

Keywords

Robustness in IR models, Adversarial robustness, OOD robustness

ACM Reference Format:

Yu-An Liu, Ruqing Zhang, Jiafeng Guo, and Maarten de Rijke. 2025. Robust Information Retrieval. In *Proceedings of the Eighteenth ACM International Conference on Web Search and Data Mining (WSDM '25)*, March 10–14, 2025, Hannover, Germany. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3701551.3703476>

1 General Information

This tutorial targets every researcher and practitioner who is interested in, or troubled by, the robustness of information retrieval systems. The material will be accessible to anyone who has a basic knowledge of IR and NLP. This is intended as a half-day (three hours plus breaks) tutorial.



WSDM '25, March 10–14, 2025, Hannover, Germany
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1329-3/25/03
<https://doi.org/10.1145/3701551.3703476>

2 Presenters

Yu-An Liu is a Ph.D. student at the Institute of Computing Technology, Chinese Academy of Sciences. He obtained his B.Eng. from SDU. His research centers on information retrieval, with a particular focus on adversarial and out-of-distribution robustness of IR systems. He has conducted robust IR workshops at SIGIR'24 [22] as a presenter with existing personnel. He is the first author of several full papers on adversarial robustness in IR, presented at SIGIR'23 [24], CIKM'23 [23], AAAI'24 [27], and SIGIR'24 [25], as well as a paper on OOD robustness in IR at Gen-IR@SIGIR'23 [21].

Ruqing Zhang is an Associate Researcher at the Institute of Computing Technology, Chinese Academy of Sciences. Her recent research focuses on information retrieval, with a particular emphasis on the robustness of information retrieval systems, trustworthy retrieval through the lens of causality, and generative information retrieval. She has authored several papers in the field of robust information retrieval [21, 23–25, 27, 37–39]. Additionally, she was a co-organizer of tutorials and workshops at SIGIR, WWW, ECIR, and SIGIR-AP, e.g., Gen-IR workshops at SIGIR'23 and SIGIR'24, and Gen-IR tutorials at SIGIR-AP'23/WWW'24/ECIR'24.

Jiafeng Guo is a Researcher at the Institute of Computing Technology, Chinese Academy of Sciences (CAS) and a Professor at the University of Chinese Academy of Sciences. He is the director of the CAS key lab of network data science and technology. He has worked on a number of topics related to Web search and data mining, with a current focus on neural models for information retrieval and natural language understanding. He has received multiple best paper (runner-up) awards at leading conferences (CIKM'11, SIGIR'12, CIKM'17, WSDM'22). He has been (co)chair for many conferences, e.g., reproducibility track co-chair of SIGIR'23, workshop co-chair of SIGIR'21, and short paper co-chair of SIGIR'20. He served as an associate editor for ACM Transactions on Information Systems and Information Retrieval Journal. Jiafeng has previously taught tutorials at many IR-related conferences.

Maarten de Rijke is a Distinguished University Professor of Artificial Intelligence and Information Retrieval at the University of Amsterdam. His research is focused on designing and evaluating trustworthy technology to connect people to information, particularly search engines, recommender systems, and conversational assistants. He is the scientific director of the Innovation Center for Artificial Intelligence and a former editor-in-chief of ACM Transactions on Information Systems and of Foundations and Trends in Information Retrieval, and a current co-editor-in-chief of Springer's

Information Retrieval book series, (associate) editor for various journals and book series. He has been general (co)chair or program (co)chair for CIKM, ECIR, ICTIR, SIGIR, WSDM, WWW, and has previously taught tutorials at these same venues and AAAI.

3 Motivation

Information retrieval (IR) systems are an important way for people to access information. In recent years, with the development of deep learning, deep neural networks have begun to be applied in IR systems [15, 20], achieving remarkable effectiveness. However, beyond their effectiveness, these neural IR models also inherit the inherent *robustness flaws* of neural networks [36, 38]. This poses a hindrance to their widespread application in the real world.

In the past few years, the issue of the robustness of IR has received wide attention, e.g., Wu et al. [39] analyzed the robustness of neural ranking models (NRMs), and a perspective paper on competitive search [18] discussed adversarial environments in search engines. Since then, there has been a lot of work that focuses on different robustness aspects in IR, such as adversarial robustness [19, 23, 25, 38], out-of-distribution (OOD) robustness [11, 36], performance variance [39] and on the corresponding improvement options. Today, the research community can effectively scrutinize IR models leading to more robust and reliable IR systems.

To ensure the quality of the tutorial, we will focus on the two most widely studied types of robustness issues, namely *adversarial robustness* and *OOD robustness*. There are many analyses and suggestions for improvement around these two robustness issues, but it has not yet been systematically organized. Through this tutorial, we aim to summarize and review the progress of robust IR to attract attention and promote widespread development.

4 Objectives

1. Introduction. We start by reminding our audience of the required background and introducing the motivation and scope of the robustness issue in IR in our tutorial.

2. Preliminaries. In IR, robustness signifies an IR system's consistent performance and resilience against diverse unexpected situations. There is a large volume of work that covers many aspects of IR robustness, e.g., (i) *Performance variance* [39], which emphasizes the worst-case performance across different individual queries under the independent and identically distributed (IID) data; (ii) *OOD robustness* [36, 39], which measures the performance of an IR model on unseen queries and documents from different distributions of the training dataset; and (iii) *Adversarial robustness* [23, 38], which focuses on the ability of the IR model to defend against malicious adversarial attacks aimed at manipulating rankings.

In this tutorial, we focus on adversarial robustness and OOD robustness, which have received the most attention. Interest in adversarial robustness stems largely from the widespread practice of search engine optimization (SEO) [4]. Concerns about OOD robustness are primarily due to the need for adaptation across diverse and complex real-world scenarios. Moreover, as large language models (LLMs) are being integrated into IR, new robustness challenges emerge; LLMs also offer novel opportunities for enhancing the robustness of IR systems.

Building on these preliminaries, we will cover adversarial robustness, OOD robustness, and robust IR in the age of LLMs.

3. Adversarial robustness. The web is a competitive search environment, which can lead to the emergence of SEO, in turn causing a decline in the content quality of search engines [4, 18]. With the gradual rise of SEO, traditional web spamming [4] started to become an effective way to attack IR systems. However, this approach based on keyword stacking is easily detected by statistical-based spamming detection methods.

Adversarial attacks. In order to exploit the vulnerability of neural IR models, many research works have simulated real black-hat SEO scenarios and proposed a lot of adversarial attack methods. (i) First, we introduce the differences between attacks in IR and CV/NLP, including task scenarios and attack targets; (ii) Then, we present adversarial retrieval attacks [23, 28, 43] against the first-stage retrieval models, including the task definition and evaluation. Current retrieval attack methods mainly include corpus poison attacks [23, 43], backdoor attacks [28], and encoding attacks [1]; and (iii) Finally, we introduce adversarial ranking attacks [10, 19, 24, 25, 38] against NRMs with task definitions and evaluation setups. These include word substitution attacks [24, 38], trigger attacks [19, 24], prompt attacks [10], and multi-granular attacks [25].

Adversarial defense. To cope with adversarial attacks, research has proposed a series of adversarial defense methods to enhance the robustness of IR models. (i) We introduce the objective and evaluation of IR defense tasks. Based on these defense principles, adversarial defense methods in IR can be classified as attack detection, empirical defense, and certified robustness; (ii) We turn to attack detection, which includes perplexity-based, linguistic-based, and learning-based detection [9]; (iii) We present empirical defenses, which encompass data augmentation [8], traditional adversarial training [29], and theory-guided adversarial training [27]; and (iv) We introduce the certified robustness in IR [37].

4. Out-of-distribution robustness. In real-world scenarios, search engines are in an ever-changing data environment, and new data are often not IID with the training data. Therefore, the ability to generalize to OOD data or not is the basis for the evaluation of IR systems in terms of OOD robustness [39].

OOD generalizability on unseen documents. In IR, the OOD robustness scenarios that have been examined can be categorized into unseen documents and unseen queries. The unseen documents scenario may be caused by adaptation to new corpus [36] or by incrementation of original corpus [2]. (i) Adaptation to new corpus usually refers to the phenomenon that the corpus on which an IR model is trained is not in the same domain as the corpus on which it is tested. Due to the overhead of retraining, the performance of the model on the new domain needs to be guaranteed under zero/few-shot scenario, which is usually solved by data augmentation [7, 32], domain modeling [41], architectural modification [16], scaling up the model capacity [30]; and (ii) Incrementation of original corpus refers to the scenario where new documents are continuously added to the corpus with potential distribution drift. In this situation, the IR model should effectively adapt to the evolving distribution with the unlabeled new-coming data, which is usually solved by continual learning [2].

OOD generalizability on unseen queries. Unseen queries concern query variations [45] and unseen query types [39]. (i) The query variations are usually different expressions of the same information need [21, 45] which may impact the effectiveness of IR models. Many noise-resistant approaches such as self-teaching methods [11, 44], contrastive learning methods [33, 45], hybrid methods [3, 35] have been proposed for neural IR models; and (ii) Unseen query types refer to the unfamiliar query type with new query intents [39]. Domain regularization [12] is effective for dealing with new query types.

5. Robust IR in the age of LLMs. (i) We first discuss the potential robustness challenges with applications of LLMs in IR, such as retrieval augmentation [17, 31], and LLMs for ranking [34, 42]; and (ii) Then, we will discuss how LLMs can be used to enhance the robustness of IR systems. These explorations will inspire many novel attempts in this area.

6. Conclusions and future directions. We conclude our tutorial by discussing several important questions and future directions, including (i) There is a diverse focus on the robustness of IR models from multiple perspectives. Establishing a unified benchmark of analysis to systematically analyze the robustness of all aspects of existing models. (ii) For adversarial robustness, existing work on adversarial attacks focuses on specific stages (first-stage retrieval or re-ranking) [23, 38] in IR systems. Customizing adversarial examples to make them effective for all stages is challenging. Therefore, one potential future direction is to explore how we can design a general unified attack method that can cater to every IR stage. (iii) For OOD robustness, the main limitation of existing work is the difficulty of seeing enough diverse domain data in advance, leading to insufficient transfer capabilities of the model. Using the generation capabilities of LLMs to synthesize corpora for adaptation domains seems to be a promising direction.

5 RELEVANCE TO THE IR COMMUNITY

In recent years, a considerable number of tutorials focusing on the topic of robustness have emerged across disciplines within computer science. In KDD'21 [13], CVPR'21 [6], and AAAI'22, there were tutorials on robustness for AI and computer vision. In EMNLP'21 [5] and EMNLP'23 [40], there were tutorials on robustness and security challenges in NLP. The focus of these tutorials was not on search tasks and models.

Web search is a core theme at WSDM, and the sub-themes of adversarial search, search benchmarking and evaluation, query analysis, etc. Robust information retrieval aligns well with these core themes. Recently, robust information retrieval has gained considerable attention as more and more work is now devoted to analyzing and improving the robustness of information retrieval systems [14, 19, 26, 36, 39]. Our tutorial will describe recent advances in robust information retrieval and shed light on future research directions. It would benefit the community and help to encourage further research into robust IR.

6 Previous Editions

We have presented this tutorial at SIGIR'24 on July 2024. At WSDM'25 we offer a new edition of the tutorial that has been revised based on the feedback received and incorporates coverage of new relevant

work. Especially during the half-year period, the cutting-edge work in robust IR and the new trends brought about by the development of LLMs will be the focus of discussion.

7 FORMAT AND DETAILED SCHEDULE

A detailed schedule for our proposed half-day tutorial is as follows:

1. **Introduction** (15 minutes)
 - Introduction to robust IR: motivation and scope
2. **Preliminaries** (20 minutes)
 - Definition & taxonomy of robustness in IR
3. **Adversarial Robustness** (50 minutes)
 - Traditional Web spamming
 - Adversarial attacks
 - Comparison: IR attacks vs. CV/NLP attacks
 - Retrieval attacks: definition, evaluation, method, etc.
 - Ranking attacks: definition, evaluation, method, etc.
 - Adversarial defense
 - IR defense tasks: objective & evaluation
 - Empirical defense: adversarial training, detection, etc.
 - Theoretical defense: certified defense, etc.
4. **Out-of-distribution Robustness** (45 minutes)
 - OOD generalizability in IR
 - OOD generalizability on unforeseen corpus
 - Definition & evaluation
 - Adaptation to new corpus
 - Incrementation of original corpus
 - OOD generalizability on unforeseen queries
 - Definition & evaluation
 - Query variation
 - Unseen query type
5. **Robust IR in the Age of LLMs** (20 minutes)
 - New challenges to IR robustness from LLMs
 - New solutions for IR robustness via LLMs
6. **Challenges and Future Directions** (20 minutes)

8 TUTORIAL MATERIALS

We plan to make all teaching materials available online for attendees, including: (i) Slides: The slides will be made publicly available. (ii) Annotated bibliography: This compilation will contain references listing all works discussed in the tutorial, serving as a valuable resource for further study. (iii) Reading list: We will provide a reading list with a compendium of existing work, open-source code libraries, and datasets relevant to the work discussed in the tutorial. We intend to ensure that all instructional materials are available online. Moreover, we grant permission to include slides and video recordings in the ACM anthology.

Acknowledgments

This work was funded by the National Natural Science Foundation of China (NSFC) under Grants No. 62472408, the Strategic Priority Research Program of the CAS under Grants No. XDB0680102, the National Key Research and Development Program of China under Grants No. 2023YFA1011602, the Lenovo-CAS Joint Lab Youth Scientist Project, and the project under Grants No. JCKY2022130C039. This work was also (partially) funded by the Dutch Research Council (NWO), under project numbers 024.004.022, NWA.1389.20.183,

and KICH3.LTP.20.006, and the European Union’s Horizon Europe program under grant agreement No. 101070212.

All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

References

- [1] Nicholas Boucher, Luca Pajola, Ilia Shumailov, Ross Anderson, and Mauro Conti. 2023. Boosting Big Brother: Attacking Search Engines with Encodings. *arXiv preprint arXiv:2304.14031* (2023).
- [2] Yinqiong Cai, Keping Bi, Yixing Fan, Jiafeng Guo, Wei Chen, and Xueqi Cheng. 2023. L2R: Lifelong Learning for First-stage Retrieval with Backward-Compatible Representations. In *CIKM*. 183–192.
- [3] Daniel Campos, ChengXiang Zhai, and Alessandro Magnani. 2023. Noise-robust Dense Retrieval via Contrastive Alignment Post Training. *arXiv e-prints* (2023), arXiv–2304.
- [4] Carlos Castillo and Brian D. Davison. 2011. Adversarial Web Search. *Foundations and Trends in Information Retrieval* 4, 5 (2011), 377–486.
- [5] Kai-Wei Chang, He He, Robin Jia, and Sameer Singh. 2021. Robustness and Adversarial Examples in Natural Language Processing. In *EMNLP*. 22–26.
- [6] Pin-Yu Chen and Sayak Paul. 2021. Practical Adversarial Robustness in Deep Learning: Problems and Solutions. In *CVPR*.
- [7] Xu Chen, Zida Cheng, Shuai Xiao, Xiaoyi Zeng, and Weilin Huang. 2023. Cross-Domain Augmentation Networks for Click-Through Rate Prediction. *ArXiv abs/2305.03953* (2023). <https://doi.org/10.48550/arXiv.2305.03953>
- [8] Xuanang Chen, Ben He, Kai Hui, Le Sun, and Yingfei Sun. 2023. Dealing with Textual Noise for Robust and Effective BERT Re-ranking. *IPM* 60, 1 (2023), 103135.
- [9] Xuanang Chen, Ben He, Le Sun, and Yingfei Sun. 2023. Defense of Adversarial Ranking Attack in Text Retrieval: Benchmark and Baseline via Detection. *arXiv preprint arXiv:2307.16816* (2023).
- [10] Xuanang Chen, Ben He, Zheng Ye, Le Sun, and Yingfei Sun. 2023. Towards Imperceptible Document Manipulations against Neural Ranking Models. In *ACL*. 6648–6664.
- [11] Xuanang Chen, Jian Luo, Ben He, Le Sun, and Yingfei Sun. 2022. Towards Robust Dense Retrieval via Local Ranking Alignment. In *IJCAL*. 1980–1986.
- [12] Daniel Cohen, Bhaskar Mitra, Katja Hofmann, and W Bruce Croft. 2018. Cross Domain Regularization for Neural Ranking Models Using Adversarial Learning. In *SIGIR*. 1025–1028.
- [13] Anupam Datta, Matt Fredrikson, Klas Leino, Kaiji Lu, Shayak Sen, and Zifan Wang. 2021. Machine Learning Explainability and Robustness: Connected at the Hip. In *SIGKDD*. 4035–4036.
- [14] Gregory Goren, Oren Kurland, Moshe Tennenholtz, and Fiana Raiber. 2020. Ranking-incentivized Quality Preserving Content Modification. In *SIGIR*. 259–268.
- [15] Jiafeng Guo, Yixing Fan, Qingyao Ai, and W Bruce Croft. 2016. A Deep Relevance Matching Model for Ad-hoc Retrieval. In *CIKM*. 55–64.
- [16] Michelle Chen Huebscher, Christian Buck, Massimiliano Ciaramita, and Sascha Rothe. 2022. Zero-Shot Retrieval With Search Agents and Hybrid Environments. *arXiv preprint arXiv:2209.15469* (2022).
- [17] Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. 2023. Atlas: Few-shot Learning with Retrieval Augmented Language Models. *Journal of Machine Learning Research* 24, 251 (2023), 1–43.
- [18] Oren Kurland and Moshe Tennenholtz. 2022. Competitive Search. In *SIGIR*.
- [19] Jiawei Liu, Yangyang Kang, Di Tang, Kaisong Song, Changlong Sun, Xiaofeng Wang, Wei Lu, and Xiaozhong Liu. 2022. Order-Disorder: Imitation Adversarial Attacks for Black-box Neural Ranking Models. In *CCS*. 2025–2039.
- [20] Shichen Liu, Fei Xiao, Wenwu Ou, and Luo Si. 2017. Cascade Ranking for Operational E-commerce Search. In *SIGKDD*. 1557–1565.
- [21] Yu-An Liu, Ruqing Zhang, Jiafeng Guo, Wei Chen, and Xueqi Cheng. 2023. On the Robustness of Generative Retrieval Models: An Out-of-Distribution Perspective. In *Gen-IR@SIGIR*.
- [22] Yu-An Liu, Ruqing Zhang, Jiafeng Guo, and Maarten de Rijke. 2024. Robust Information Retrieval. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 3009–3012.
- [23] Yu-An Liu, Ruqing Zhang, Jiafeng Guo, Maarten de Rijke, Wei Chen, Yixing Fan, and Xueqi Cheng. 2023. Black-Box Adversarial Attacks against Dense Retrieval Models: A Multi-View Contrastive Learning Method. In *CIKM*. 1647–1656.
- [24] Yu-An Liu, Ruqing Zhang, Jiafeng Guo, Maarten de Rijke, Wei Chen, Yixing Fan, and Xueqi Cheng. 2023. Topic-Oriented Adversarial Attacks against Black-Box Neural Ranking Models. In *SIGIR*. 1700–1709.
- [25] Yu-An Liu, Ruqing Zhang, Jiafeng Guo, Maarten de Rijke, Yixing Fan, and Xueqi Cheng. 2024. Multi-granular Adversarial Attacks against Black-box Neural Ranking Models. In *SIGIR*.
- [26] Yu-An Liu, Ruqing Zhang, Jiafeng Guo, Maarten de Rijke, Yixing Fan, and Xueqi Cheng. 2024. Robust neural information retrieval: An adversarial and out-of-distribution perspective. *arXiv preprint arXiv:2407.06992* (2024).
- [27] Yu-An Liu, Ruqing Zhang, Mingkun Zhang, Wei Chen, Maarten de Rijke, Jiafeng Guo, and Xueqi Cheng. 2024. Perturbation-Invariant Adversarial Training for Neural Ranking Models: Improving the Effectiveness-Robustness Trade-Off. In *AAAI*, Vol. 38.
- [28] Quanyu Long, Yue Deng, LeiLei Gan, Wenya Wang, and Sinno Jialin Pan. 2024. Backdoor Attacks on Dense Passage Retrievers for Disseminating Misinformation. *arXiv preprint arXiv:2402.13532* (2024).
- [29] Simon Lupart and Stéphane Clinchant. 2023. A Study on FGSM Adversarial Training for Neural Retrieval. In *ECIR*. Springer, 484–492.
- [30] Jianmo Ni, Chen Qu, Jing Lu, Zhuyun Dai, Gustavo Hernandez Abrego, Ji Ma, Vincent Zhao, Yi Luan, Keith Hall, Ming-Wei Chang, et al. 2022. Large Dual Encoders Are Generalizable Retrievers. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. 9844–9855.
- [31] Shiyu Ni, Keping Bi, Jiafeng Guo, and Xueqi Cheng. 2024. When Do LLMs Need Retrieval Augmentation? Mitigating LLMs’ Overconfidence Helps Retrieval Augmentation. *arXiv preprint arXiv:2402.11457* (2024).
- [32] Seokjin Oh, Su ah Lee, and Woohwan Jung. 2023. Data Augmentation for Neural Machine Translation using Generative Language Model. *ArXiv abs/2307.16833* (2023). <https://doi.org/10.48550/arXiv.2307.16833>
- [33] Georgios Sidiropoulos and Evangelos Kanoulas. 2024. Improving the Robustness of Dense Retrievers Against Typos via Multi-Positive Contrastive Learning. In *European Conference on Information Retrieval*. Springer, 297–305.
- [34] Weiwei Sun, Lingyong Yan, Xinyu Ma, Shuaiqiang Wang, Pengjie Ren, Zhumin Chen, Dawei Yin, and Zhaochun Ren. 2023. Is ChatGPT Good at Search? Investigating Large Language Models as Re-Ranking Agents. In *EMNLP*. 14918–14937.
- [35] Panuthe Tasawong, Wuttikorn Ponwittayarat, Peerat Limkonchotiwat, Can Udomcharoenchaikit, Ekapol Chuangsuwanich, and Sarana Nutanong. 2023. Typo-Robust Representation Learning for Dense Retrieval. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. 1106–1115.
- [36] Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. BEIR: A Heterogeneous Benchmark for Zero-shot Evaluation of Information Retrieval Models. In *NIPS*.
- [37] Chen Wu, Ruqing Zhang, Jiafeng Guo, Wei Chen, Yixing Fan, Maarten de Rijke, and Xueqi Cheng. 2022. Certified Robustness to Word Substitution Ranking Attack for Neural Ranking Models. In *CIKM*. 2128–2137.
- [38] Chen Wu, Ruqing Zhang, Jiafeng Guo, Maarten de Rijke, Yixing Fan, and Xueqi Cheng. 2023. PRADA: Practical Black-Box Adversarial Attacks against Neural Ranking Models. *TOIS* 41, 4 (2023), Article 89.
- [39] Chen Wu, Ruqing Zhang, Jiafeng Guo, Yixing Fan, and Xueqi Cheng. 2022. Are Neural Ranking Models Robust? *TOIS* 41, 2 (2022), 1–36.
- [40] Qionghai Xu and Xuanli He. 2023. Security Challenges in Natural Language Processing Models. In *EMNLP*. 7–12.
- [41] Jingtao Zhan, Qingyao Ai, Yiqun Liu, Jiaxin Mao, Xiaohui Xie, Min Zhang, and Shaoping Ma. 2022. Disentangled Modeling of Domain and Relevance for Adaptable Dense Retrieval. *arXiv preprint arXiv:2208.05753* (2022).
- [42] Hengran Zhang, Ruqing Zhang, Jiafeng Guo, Maarten de Rijke, Yixing Fan, and Xueqi Cheng. 2024. Are Large Language Models Good at Utility Judgments?. In *SIGIR*.
- [43] Zexuan Zhong, Ziqing Huang, Alexander Wettig, and Danqi Chen. 2023. Poisoning Retrieval Corpora by Injecting Adversarial Passages. In *EMNLP*.
- [44] Shengyao Zhuang, Linjun Shou, Jian Pei, Ming Gong, Houxing Ren, Guido Zuccon, and Daxin Jiang. 2023. Typos-Aware Bottlenecked Pre-Training for Robust Dense Retrieval. In *Proceedings of the Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*. 212–222.
- [45] Shengyao Zhuang and Guido Zuccon. 2021. Dealing with Typos for BERT-based Passage Retrieval and Ranking. In *EMNLP*. 2836–2842.