

The Impact of Main Content Extraction on Near-Duplicate Detection

Maik Fröbe,^{*} Matthias Hagen,^{*} Martin-Luther-Universität Halle-Wittenberg, Germany
Janek Bevendorff,[†] Michael Völske,[†] Benno Stein,[†] Bauhaus-Universität Weimar, Germany
Christopher Schröder,[‡] Robby Wagner,[‡] Lukas Gienapp,[‡] Martin Potthast,[‡]
Leipzig University, Germany

Abstract

Commercial web search engines employ near-duplicate detection to ensure that users see each relevant result only once, albeit the underlying web crawls typically include (near-)duplicates of many web pages. We revisit the risks and potential of near-duplicates with an information retrieval focus, motivating that current efforts toward an open and independent European web search infrastructure should maintain metadata on duplicate and near-duplicate documents in its index.

Near-duplicate detection implemented in an open web search infrastructure should provide a suitable similarity threshold, a difficult choice since identical pages may substantially differ in parts of a page that are irrelevant to searchers (templates, advertisements, etc.). We study this problem by comparing the similarity of pages for five (main) content extraction methods in two studies on the ClueWeb crawls. We find that the full content of pages serves precision-oriented near-duplicate-detection, while main content extraction is more recall-oriented.

INTRODUCTION

Typical web crawls contain many pages with identical or very similar content and different URLs [10]. Search engines retrieving pages from such web crawls may encounter those near-duplicates in multiple stages of their pipeline. During indexing, omitting near-duplicates might reduce the index size. During retrieval, near-duplicates might occur in the search engine result pages, reducing the user experience because users gain nothing from viewing the same result twice or more on the search engine result pages [5]. Hence, identifying near-duplicates is a mandatory step in web search, with commercial search engines like Google showing only the “best” version from a set of near-duplicates for a query.^{*}

Widely available web crawls—most notably the ClueWebs[†] and the Common Crawl[‡]—contain the (near-)duplicate documents that the crawler encountered during the crawling process. While the inclusion of near-duplicates enables many applications (like research on text reuse [2]), it introduces problems for search engines (that we will discuss later in this paper). The CopyCat resource [12] addresses the problems introduced by near-duplicates in information retrieval experiments by providing a precision-oriented near-duplicate detection. CopyCat comes in two parts: (1) ready-to-use compilations of near-duplicate documents within and between selected web crawls, and (2) a software library to deduplicate arbitrary document sets, e.g., search engine result pages before they are shown to searchers.

^{*} <first-name>.<last-name>@informatik.uni-halle.de

[†] <first-name>.<last-name>@uni-weimar.de

[‡] martin.pothast@uni-leipzig.de

^{*} developers.google.com/search/docs/advanced/guidelines/duplicate-content

[†] lemurproject.org/clueweb09.php/ and lemurproject.org/clueweb12/

[‡] commoncrawl.org/

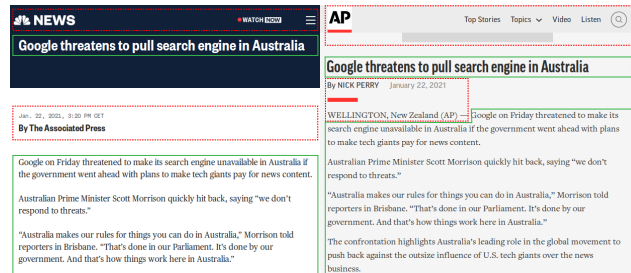


Figure 1: Pages with different URLs and the same article. Both pages have identical content (indicated by green boxes) but vary in parts irrelevant to searchers (indicated by red dashed boxes).

Removing near-duplicates from the search engine result pages with a framework like CopyCat comes with the inherent difficulty of balancing precision and recall. A low precision might reduce the effectiveness of a search engine because relevant and novel documents might be omitted. In contrast, low recall reduces the user experience because users see more near-duplicates. In the context of web search, tuning a similarity measure faces additionally the problem that some parts of the pages that are irrelevant for searchers can increase or decrease the similarity of pages. Figure 1 shows an example of two identical articles located at different URLs where the “noise,” i.e., the navigation bar, reduces the similarity, eventually having a negative impact on the recall.

Using only the “retrieval-relevant” part of documents for the near-duplicate detection might be a promising direction to improve the recall in cases as exemplified in Figure 1. However, the impact of main content extraction on near-duplicate detection is not studied so far. While a “perfect” main content extraction should positively affect precision and recall, existing implementations might even harm precision and recall. E.g., invisible changes in the HTML structure of a page might cause tag-based main content extraction approaches to extract different main contents from pages that have identical content from the users’ perspective.

We conduct experiments on the ClueWeb crawls to investigate to what extent main content extraction may improve near-duplicate detection in information retrieval. Therefore, we extract document pairs from the ClueWebs containing redundant content (according to their canonical URL) to draw a sharp line between “roughly similar” documents and near-duplicates. We calculate the syntactic similarities for all document pairs after extracting the text from the raw HTML with four main content extraction methods and one full content extraction method that does not discard the “noise” parts from the document. We evaluate the precision and recall of all five content extraction methods based on manual near-duplicate judgments in two case studies. First, we review document pairs uniformly sampled from the full similarity range for all content

extraction methods confirming that main content extraction increases the recall as exemplified in Figure 1. Secondly, we review 100 cases per main content extraction method in which main content extraction changes the similarity drastically, e.g., with identical documents having disjoint main content, or dissimilar pages having duplicate main content, finding inaccurate main content extraction in most of those cases.

RELATED WORK

This section reviews definitions for near-duplicates, their prevalence on the web, and approaches to near-duplicate detection implemented in the CopyCat framework.

Defining Near-Duplicates

The fact that there is no universal near-duplicate definition renders their detection and comparable analysis difficult. Restrictive near-duplicate definitions [15, 17] consider documents as near-duplicates if they differ only by their session or message IDs, timestamps, visitor counts, server names, invisible differences, URL parts, or if they are entry pages to the same site. Note that documents, even with minimal content changes, are often not considered near-duplicates under such a restrictive definition. Bernstein and Zobel [6] relax the near-duplicate definition by applying an information retrieval focus, allowing minimal changes in the content. They consider a document pair as near-duplicate if users get the same information from both documents for all “reasonable queries.” We adopt the near-duplicate definition of Bernstein and Zobel since it considers a pair of pages as duplicates if they are equivalent in terms of information provided, i.e., ignoring parts of pages irrelevant to searchers (templates, advertisements, etc.).

Studies on Near-Duplicates on the Web

According to previous studies by Fetterly et al. [10, 11], 30% of the pages on the web are near-duplicates. While web pages change regularly, consecutive versions of the same web page are usually highly similar [9]. Subsequent investigations [1, 10, 11, 18, 19] confirm this observation by tracking web pages between 5 weeks and one year. For example, Adar et al. [1] repeatedly crawl 55 000 URLs over 5 weeks finding two-thirds of the pages changed their content, observing that most of these changes were minimal. Ntoulas et al. [18] tracked 150 pages over one year, finding that 40% of them were still accessible after one year, noticing only insignificant changes on most pages.

Near-Duplicate Detection

There are syntactic, URL-based, and semantic algorithms for detecting near-duplicates [3], from which the detection of syntactic near-duplicates received the most attention, resulting in many effective algorithms based on fingerprinting techniques [7, 8, 15, 17]. The CopyCat framework implements syntactic near-duplicate detection in large web crawls with the SimHash algorithm using a fingerprint size of 64 bit and a Hamming-threshold of 3 bits as suggested by Manku et al. [17], while reducing the number of calculated pairwise similarities with the partitioning scheme proposed by Henzinger et al. [15]. Complementary to estimating the similarity of documents with SimHash, CopyCat can calculate the lossless S_3 fingerprint similarity [5] for near-duplicate detection in

small sets of documents, such as run and qrel files frequently used in information retrieval experiments.

NEAR-DUPPLICATES IN WEB CRAWLS: RISKS AND POTENTIALS

We recapitulate two risks and one potential of near-duplicate pages in web crawls that the CopyCat framework addresses. Please note that we here focus on information retrieval and that other risks, e.g., in the training of large language models [21], exist.

Risk: Evaluation of Search Engines

Bernstein and Zobel [6] found that near-duplicates cause problems in information retrieval evaluations because search engine users do not benefit from seeing near-duplicates. Therefore, they introduce the so-called novelty principle, which states that a document, though relevant in isolation, is irrelevant if it is a near-duplicate to a document the user has already seen on the search engine result page. Especially on web crawls with many near-duplicates, the novelty principle has a non-negligible impact on evaluating search engines [14]. E.g., applying the novelty principle on the runs submitted to the Terabyte track 2004 decreases mean average precision scores by 20% on average [6].

The classical evaluation setup of search engines employs the Cranfield paradigm, making it pretty easy to oversee negative impacts caused by near-duplicates. Relevance assessors judge the relevance of documents to a query in isolation, seeing only one document at a time. Hence, situations that would severely reduce the experience for searchers, e.g., when many near-duplicates occur at subsequent positions in the ranking, can be overlooked because assessors do not look at the ranking. Topic 194 of the ClueWeb09 Web Tracks includes a particularly striking example, where among 47 relevant documents, there are 40 near-duplicates of the same Wikipedia article.

Risk: Training of Learning to Rank Models

Near-duplicates form a kind of oversampling because multiple identical or very similar copies of a page are in the dataset. As recently exemplified [22], oversampling data before partitioning it into training and test sets can invalidate evaluations in machine learning because models may see the same object during training and test. This leakage of information is not possible during the training of learning to rank models because the train/test partitioning is done per query. Still, not removing near-duplicates during the training of learning to rank models decreases the effectiveness of models and biases the trained models [13].

A study [13] on the ClueWeb09 with 42 ranking features using popular algorithms finds that near-duplicates in the training data harm the retrieval performance, since the presence of near-duplicates is unaccounted for in the loss-minimization of learning to rank and in subsequent evaluations. Furthermore, by varying the number of Wikipedia near-duplicates in the training set, the study showed that models might be biased towards retrieving near-duplicates at higher positions. Hence, these observations make a strong case that learning to rank pipelines benefit from removing duplicate documents from the data before training the model.

Mitigating the negative effects of near-duplicates during the training of retrieval models is easily possible with the CopyCat framework, which can deduplicate the training and test set.

Potential: Transfer of Relevance Labels

In contrast to the previous two risks to the validity and robustness of search engine evaluation and tuning, near-duplicate detection enables the transfer of relevance judgments between different editions (or updates) of web crawls [12]. Relevance judgments—obtained from click logs or expert assessments—are an important and costly resource for the development of search engines. E.g., the effort for the 73,883 relevance judgments for the TREC Web tracks on the ClueWeb09 crawl can be estimated at a manual labor of about 4–8 full-time person-months (assuming 40-hour weeks with 30–60 seconds per judgment [23]).

To “reduce” the costs of keeping the relevance judgments up-to-date for ever-evolving web-indices, search engines might transfer relevance judgments from the previous version of a crawl to the next version when they find the judged documents (or near-duplicates of them) in the newer version of the crawl. In a showcase [12] using precision-oriented near-duplicate detection with the CopyCat framework, 10% of the ClueWeb09 relevance judgments could be transferred to the ClueWeb12. The number transferred relevance judgments would even increase to 15% when the ClueWeb12 crawling process would have ensured that the URLs judged in the ClueWeb09 are part of the URL seeds for the next crawling round. More frequent updates (compared to the gap of three years in the relevance transfer showcase) would likely further increase the amount of transferrable relevance judgments. Additionally, the reported experiments on the transfer of relevance labels have used only the full content of the pages. Hence, further improvements, e.g., by leveraging main content extraction to increase the recall while maintaining good precision, are possible.

CONTENT EXTRACTION EXPERIMENTS

To experimentally compare the impact of main content extraction on near-duplicate detection, we construct a dataset of 186 819 ClueWeb document pairs with redundant content as indicated by canonical URLs. For each document pair, we calculate its syntactic similarity with the lossless S_3 fingerprinting [5] for four main content extraction algorithms and the full content of pages. We label 900 document pairs as near-duplicates or not sampled with two approaches: (1) with 100 document pairs stratified sampled from the S_3 distribution of each of the five content extraction methods, and (2) with 50 document pairs with maximal positive/negative S_3 differences between each of the four main content extraction methods and the full content of a page.

Dataset Construction

We aim at constructing a manageable dataset for our experiments that allows us to draw a sharp line between “only similar” documents and near-duplicates. Therefore, we identify document pairs in the ClueWeb09 and ClueWeb12 that should contain redundant content because they share the same canonical URL. We inspect all documents in the ClueWeb, group them by their canonical URL, and select 5000 groups having the same canonical link at random. From each group, we select all possible pairs (with a maximum 50 document pairs per group) giving us 186,819 document pairs.

Document Preprocessing

We preprocess all documents with the CopyCat framework. CopyCat provides five content extraction approaches that transform




Figure 2: Cumulative distribution plot showing the proportion of document pairs in our dataset below a given similarity measured by their S_3 score for all five considered content extraction methods.

the raw HTML of a page into text: four main content extraction approaches and one full content extraction. The four provided main content extraction approaches are Boilerpipe [16], Jericho,[§] Justext [20], and Trafilatutura [4]. The full content extraction uses JSoup[¶] to extract the plain text—without any main content extraction—from the HTML. After extracting the documents text, we remove stop words using Lucene’s default stop word list for English, apply stemming with the Porter Stemmer, and lower case the remaining words.

Similarity in our Dataset

We use the lossless S_3 fingerprint similarity [5] using word-8-grams to calculate the similarities between all document pairs for all five content extraction methods in our dataset. An S_3 score of 0 indicates no overlap between documents, and an S_3 score of 1 means equality. Figure 2 shows the cumulative distribution plot for all five content extraction methods regarding the portion of document pairs below a given S_3 score.

We can identify two groups of content extraction methods that share similar overall behavior. The first group consists of the main content extraction methods Trafilatutura and Boilerpipe that show many document pairs with an S_3 score of 0 (26% for Boilerpipe and 27% for Trafilatutura), indicating that the main content extraction produces disjoint main contents from the considered documents in a pair. This group additionally contains many document pairs with an S_3 score of 1 (55% for Boilerpipe and 39% for Trafilatutura), indicating that for documents in a pair often the same main content is extracted.

The second group consists of the full content, Justext, and Jericho methods of content extraction. Approaches in this group have very few document pairs with an S_3 score of 0 (the Justext approach from this group has the maximum of 0.35% of document pairs with an S_3 score of 0), and much fewer document pairs with an S_3 score of 1 (all three have around 15.6% of document pairs with an S_3 score of 1). We calculated the Pearson correlation between the full content extraction and all other methods. We found a very

[§]<http://jericho.htmlparser.net>

[¶]<https://jsoup.org/>

Table 1: Overview of (a) precision and recall for near-duplicates in the uniform sampled document pairs at high syntactic similarity ($S_3 = 1$ and $S_3 \geq 0.9$), and (b) near-duplicates per S_3 similarity in the uniform sampled document pairs for all five content extraction methods. Lastly, (c) shows near-duplicates per S_3 similarity in document pairs with large S_3 differences to the full content extraction.

	$S_3 = 1$		$S_3 \geq 0.9$	
	Pr.	Re.	Pr.	Re.
Full Content	1.00	0.06	0.98	0.26
JusText	1.00	0.07	1.00	0.24
Trafilatura	0.76	0.33	0.77	0.41
Jericho	0.93	0.07	0.97	0.26
Boilerpipe	0.60	0.50	0.61	0.59

Figure 2 consists of two violin plots, (b) and (c), showing the distribution of S_3 scores for five content extraction methods: Full Content, Justext, Trafilatura, Jericho, and Boilerpipe. The y-axis represents the S_3 Score from 0.0 to 1.0. The x-axis is labeled 'Content Extraction'. In plot (b), for uniform sampling, Full Content, Justext, and Jericho show a high concentration of 'No Near-Duplicate' (blue) scores near 1.0, while Trafilatura and Boilerpipe show a high concentration of 'Near-Duplicate' (orange) scores near 1.0. In plot (c), for large S_3 differences, Full Content shows a high concentration of 'No Near-Duplicate' (blue) scores near 0.0, while the other methods show a high concentration of 'Near-Duplicate' (orange) scores near 1.0.

high correlation to Justext and Jericho (0.97 respectively 0.99) and only a moderate correlation to Boilerpipe and Trafilatura (0.65 respectively 0.73). Overall, Figure 2 shows that the S_3 scores in our dataset differ substantially between the two groups, which motivates us to manually verify which of the document pairs are indeed near-duplicates.

Labeling Near-Duplicates

After calculating the S_3 scores for all document pairs with all content extraction methods, we sample two sets of document pairs for manual review. First, we sample 100 document pairs uniformly covering S_3 scores between 0 and 1 for all five content extraction methods. Second, we sample document pairs with large S_3 differences between a main content extraction method and the full content extraction aiming at identifying document pairs where main content extraction yields opposite S_3 scores to full content extraction. Therefore, we select the 50 document pairs with the largest positive and largest negative S_3 difference for all four main content extraction methods for manual review.

We use the near-duplicate definition and review guidelines of Bernstein and Zobel [6] to label near-duplicates: A document pair is considered as near-duplicate when both documents are content-equivalent, and users would be able to extract the same information from either one for all reasonable queries. Two versions of the same Wikipedia article with only minor non-content changes are an example of near-duplicates under this definition.

We labeled the two document pair samples with two assessors. We applied a κ -test on the 100 document pairs sampled for S_3 similarities between 0 and 1 for the full content method, finding a high Fleiss' κ of 0.78, indicating good agreement between both assessors. In a follow-up discussion among the annotators, we discussed all 11 document pairs with different near-duplicate judgments, finally agreeing in all cases. After our κ test, each annotator judged the document pairs for the same two main content extraction methods for both user studies.

Evaluation

Table 1a and Table 1b shows the ability of all five content extraction methods to identify near-duplicate documents in our set of 500 manually reviewed document pairs that uniformly cover S_3 scores between 0 and 1. In Table 1a, we report precision and recall

for S_3 thresholds of 1 (for exact duplicates after content extraction) and 0.9 (highly similar extracted content). As in our initial discussion on similarity scores produced by the five content extraction methods, Trafilatura and Boilerpipe (the group with many document pairs at an S_3 score of 1 in Figure 2) as well as the full content, Justext, and Jericho (the group with fewer document pairs with an S_3 score of 1 in Figure 2) show similar behavior in terms of precision and recall. The full content and Justext approaches show a perfect precision of 1.0 at an S_3 threshold of 1, and Justext even has a perfect precision at an S_3 threshold of 0.9. On the other side, Trafilatura and Boilerpipe show a very high recall. Even for an S_3 score of 1, Boilerpipe achieves a remarkable Recall of 0.5.

Table-1b shows the correctly and wrongly identified near-duplicates per S_3 score for all content extraction methods in our set of 500 manually reviewed document pairs that uniformly cover S_3 scores between 0 and 1. Again, we can see similar behavior for the full content, Justext, and Jericho methods which make almost no mistakes at high respectively low S_3 scores. In the opposite group, with Trafilatura and Boilerpipe, we see quite some mistakes (even at $S_3 = 1$ and $S_3 = 0$).

Table 1c shows the correctly and wrongly identified near-duplicates per S_3 score for all content extraction methods in our set of 400 manually reviewed document pairs for which the main content extraction changes the similarity drastically. In almost all cases, barring few exceptions, we find that for such large differences, the S_3 score calculated on the full content correctly identifies near-duplicates and non-near-duplicates. This is visible since the full content method assigns, almost perfectly, non-near-duplicates an S_3 score near 0, and near-duplicates an S_3 score near 1. All other approaches make substantial mistakes in this selection of document pairs, indicated by assigning many non-near-duplicates an S_3 score near 1 (for which the full content method assigned scores near 0, since we selected large differences), and many near-duplicates an S_3 score near 0 (for which the full content method assigned scores near 1). Especially for cases in which highly similar documents get dissimilar main content extracted, we often found that the main content extraction had problems in identifying the correct main content. Overall, Trafilatura is the most vulnerable in this setting (the most near-duplicates near S_3 of 0, and most non-near-duplicates near S_3 of 1). Still, even main content extraction approaches with a very high correlation to the full content

extraction method, like Jericho and Justext in our experiments, make substantial mistakes.

CONCLUSION AND FUTURE WORK

We have recapitulated two risks and one potential application of near-duplicates in web search to motivate the maintenance of metadata on duplicate and near-duplicate documents. Given metadata on near-duplicates, it is easy to remove risks such as overestimated evaluation scores of retrieval systems or overfitting learning to rank models. Additionally, updating relevance judgments to the next version of the underlying web crawl can be done at lower costs because relevance labels might automatically be transferred to near-duplicates in the newer version.

In a first attempt to simplify the difficult decision of choosing an appropriate similarity threshold, we investigated how removing parts of documents that are rather irrelevant for the retrieval impacts the similarity of documents. Therefore, we have compared document similarities after preprocessing documents with five (main) content extraction methods. We found that main content extraction can yield very high recall for near-duplicate detection, even when only documents with identical main content are considered as near-duplicates.

An interesting prospect for future work is to include more main content extraction methods and expand the experiments to more document pairs. Another interesting direction for future work might be a further inspection of our observation that highly similar documents having very dissimilar extracted main contents were in most cases caused by mistakes in the main content extraction. This technique might help bootstrap a distant supervision dataset of documents with main content that is difficult to extract.

REFERENCES

- [1] E. Adar, J. Teevan, S. T. Dumais, and J. L. Elsas. The Web Changes Everything: Understanding the Dynamics of Web Content. In *Proceedings of the Second International Conference on Web Search and Web Data Mining, WSDM 2009, Barcelona, Spain, February 9-11, 2009*, pages 282–291, 2009.
- [2] M. Alshomary, M. Völske, T. Licht, H. Wachsmuth, B. Stein, M. Hagen, and M. Potthast. Wikipedia Text Reuse: Within and Without. In L. Azzopardi, B. Stein, N. Fuhr, P. Mayr, C. Hauff, and D. Hiemstra, editors, *Advances in Information Retrieval. 41st European Conference on IR Research (ECIR 2019)*, volume 11437 of *Lecture Notes in Computer Science*, pages 747–754, Berlin Heidelberg New York, Apr. 2019. Springer.
- [3] B. Alsulami, M. Abulkhair, and F. Eassa. Near Duplicate Document Detection Survey. *International Journal of Computer Science and Communications Networks*, 2(2):147–151, 2012.
- [4] A. Barbaresi. Generic web content extraction with open-source software. In *Proceedings of the 15th Conference on Natural Language Processing, KONVENS 2019, Erlangen, Germany, October 9-11, 2019*, 2019.
- [5] Y. Bernstein and J. Zobel. A Scalable System for Identifying Co-derivative Documents. In *String Processing and Information Retrieval, 11th International Conference, SPIRE 2004, Padova, Italy, October 5-8, 2004, Proceedings*, pages 55–67, 2004.
- [6] Y. Bernstein and J. Zobel. Redundant Documents and Search Effectiveness. In *Proceedings of the 2005 ACM CIKM International Conference on Information and Knowledge Management, Bremen, Germany, October 31 - November 5, 2005*, pages 736–743, 2005.
- [7] A. Z. Broder. On the Resemblance and Containment of Documents. In *Compression and Complexity of SEQUENCES 1997, Positano, Amalfitan Coast, Salerno, Italy, June 11-13, 1997, Proceedings*, pages 21–29, 1997.
- [8] M. Charikar. Similarity Estimation Techniques from Rounding Algorithms. In *Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada*, pages 380–388, 2002.
- [9] J. Cho and H. Garcia-Molina. The Evolution of the Web and Implications for an Incremental Crawler. In *VLDB 2000, Proceedings of 26th International Conference on Very Large Data Bases, September 10-14, 2000, Cairo, Egypt*, pages 200–209, 2000.
- [10] D. Fetterly, M. S. Manasse, and M. Najork. On the Evolution of Clusters of Near-Duplicate Web Pages. In *1st Latin American Web Congress (LA-WEB 2003), Empowering Our Web, 10-12 November 2003, Sanitago, Chile*, pages 37–45, 2003.
- [11] D. Fetterly, M. S. Manasse, M. Najork, and J. L. Wiener. A Large-Scale Study of the Evolution of Web Pages. In *Proceedings of the Twelfth International World Wide Web Conference, WWW 2003, Budapest, Hungary, May 20-24, 2003*, pages 669–678, 2003.
- [12] M. Fröbe, J. Bevendorff, L. Gienapp, M. Völske, B. Stein, M. Potthast, and M. Hagen. CopyCat: Near-Duplicates within and between the ClueWeb and the Common Crawl. In *44th International ACM Conference on Research and Development in Information Retrieval (SIGIR 2021)*. ACM, July 2021.
- [13] M. Fröbe, J. Bevendorff, J. H. Reimer, M. Potthast, and M. Hagen. Sampling Bias Due to Near-Duplicates in Learning to Rank. In J. Huang, Y. Chang, X. Cheng, J. Kamps, V. Murdock, J. Wen, and Y. Liu, editors, *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, July 25-30, 2020*, pages 1997–2000. ACM, 2020.
- [14] M. Fröbe, J. P. Bittner, M. Potthast, and M. Hagen. The Effect of Content-Equivalent Near-Duplicates on the Evaluation of Search Engines. In J. M. Jose, E. Yilmaz, J. Magalhães, P. Castells, N. Ferro, M. J. Silva, and F. Martins, editors, *Advances in Information Retrieval - 42nd European Conference on IR Research, ECIR 2020, Lisbon, Portugal, April 14-17, 2020, Proceedings, Part II*, volume 12036 of *Lecture Notes in Computer Science*, pages 12–19. Springer, 2020.
- [15] M. R. Henzinger. Finding Near-Duplicate Web Pages: A Large-Scale Evaluation of Algorithms. In *SIGIR 2006: Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Seattle, Washington, USA, August 6-11, 2006*, pages 284–291, 2006.

- [16] C. Kohlschütter, P. Fankhauser, and W. Nejdl. Boilerplate Detection Using Shallow Text Features. In B. D. Davison, T. Suel, N. Craswell, and B. Liu, editors, *Proceedings of the Third International Conference on Web Search and Web Data Mining, WSDM 2010, New York, NY, USA, February 4-6, 2010*, pages 441–450. ACM, 2010.
- [17] G. S. Manku, A. Jain, and A. D. Sarma. Detecting Near-Duplicates for Web Crawling. In C. L. Williamson, M. E. Zurko, P. F. Patel-Schneider, and P. J. Shenoy, editors, *Proceedings of the 16th International Conference on World Wide Web, WWW 2007, Banff, Alberta, Canada, May 8-12, 2007*, pages 141–150. ACM, 2007.
- [18] A. Ntoulas, J. Cho, and C. Olston. What’s new on the Web?: The Evolution of the Web from a Search Engine Perspective. In S. I. Feldman, M. Uretsky, M. Najork, and C. E. Wills, editors, *Proceedings of the 13th international conference on World Wide Web, WWW 2004, New York, NY, USA, May 17-20, 2004*, pages 1–12. ACM, 2004.
- [19] C. Olston and S. Pandey. Recrawl Scheduling Based on Information Longevity. In J. Huai, R. Chen, H. Hon, Y. Liu, W. Ma, A. Tomkins, and X. Zhang, editors, *Proceedings of the 17th International Conference on World Wide Web, WWW 2008, Beijing, China, April 21-25, 2008*, pages 437–446. ACM, 2008.
- [20] J. Pomikálek. *Removing boilerplate and duplicate content from web corpora*. PhD thesis, Masarykova univerzita, Fakulta informatiky, 2011.
- [21] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21:140:1–140:67, 2020.
- [22] G. Vandewiele, I. Dehaene, G. Kovács, L. Sterckx, O. Janssens, F. Ongenaë, F. D. Backere, F. D. Turck, K. Roelens, J. Decruyenaere, S. V. Hoecke, and T. Demeester. Overly Optimistic Prediction Results on Imbalanced Data: A Case Study of Flaws and Benefits when Applying Over-Sampling. *Artif. Intell. Medicine*, 111:101987, 2021.
- [23] E. M. Voorhees. The Philosophy of Information Retrieval Evaluation. In C. Peters, M. Braschler, J. Gonzalo, and M. Kluck, editors, *Evaluation of Cross-Language Information Retrieval Systems, Second Workshop of the Cross-Language Evaluation Forum, CLEF 2001, Darmstadt, Germany, September 3-4, 2001, Revised Papers*, volume 2406 of *Lecture Notes in Computer Science*, pages 355–370. Springer, 2001.