

The ROOTS Search Tool: Data Transparency for LLMs

Aleksandra Piktus^{1,2} Christopher Akiki^{3,4} Paulo Villegas⁵ Hugo Laurençon¹
Gérard Dupont⁶ Alexandra Sasha Luccioni¹ Yacine Jernite¹ Anna Rogers⁷

¹Hugging Face ²Sapienza University ³Leipzig University ⁴ScaDS.AI
⁵Telefonica I+D ⁶Mavenoid ⁷University of Copenhagen
piktus@huggingface.co

Abstract

ROOTS is a 1.6TB multilingual text corpus developed for the training of BLOOM, currently the largest language model explicitly accompanied by commensurate data governance efforts. In continuation of these efforts, we present the ROOTS Search Tool: a search engine over the entire ROOTS corpus offering both fuzzy and exact search capabilities. ROOTS is the largest corpus to date that can be investigated this way. The ROOTS Search Tool is open-sourced and available on [Hugging Face Spaces](#). We describe our implementation and the possible use cases of our tool.

1 Introduction

Large language models (LLMs) are ubiquitous in modern NLP, used directly to generate text and as building blocks in downstream applications. The ever-increasing size of the latest models inflates the demand for massive volumes of training data (Hoffmann et al., 2022), in practice sourced mainly from the Web. This raises questions concerning the quality of the data, the feasibility of curating and inspecting it, as well as documenting it in terms of what kinds of speech and speakers it represents (Jo and Gebru, 2020; Bender et al., 2021; Akiki et al., 2022). Without that level of characterization, we cannot tell for what varieties of language the resulting models can be expected to work well, whether the data was ethically sourced, how to interpret evaluation metrics, and to what degree a particular output was memorized directly from the training data. In an encouraging new trend, we see researchers exploring ways to quantitatively describe large datasets (Mitchell et al., 2022). However, user-friendly tools for an extensive qualitative analysis are still predominantly missing. In our current work, we aim to fill that gap for a specific, web-scale, textual corpus.

Building on the efforts of the BigScience work-

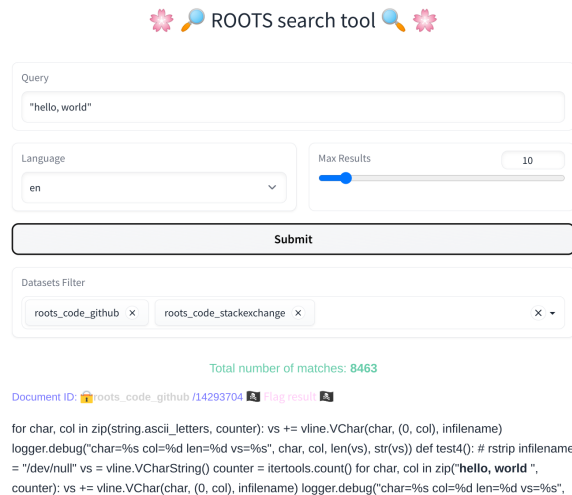


Figure 1: ROOTS search tool: user interface

shop,¹ we present the ROOTS Search Tool²—a search engine for the the 1.6TB multilingual ROOTS corpus (Laurençon et al., 2022). The ROOTS corpus was created to pre-train BLOOM (Scao et al., 2022)—the first LLM of its scale designed with commensurate efforts in responsible licensing³ and data governance (Jernite et al., 2022). We hope that our tool will facilitate qualitative analysis of the web-scale ROOTS corpus, and establish the qualitative analysis of training data—critical for the model understanding and governance work—as an essential step in the development of LLMs.

2 Related Work

Corpus linguistics. The core methodology for studying large volumes of text was developed in corpus linguistics (McEnery and Hardie, 2013), an area of research responsible for curating large text collections carefully designed to represent specific varieties of language. For example, the 100M

¹bigscience.huggingface.co

²hf.co/spaces/bigscience-data/roots-search

³bigscience.huggingface.co/blog/the-bigscience-rail-license

word British National Corpus (Leech, 1992) was created to represent the spoken and written British English of the late 20th century, with each text handpicked by experts, who also procured appropriate copyright exemptions. Similar national corpora were later created for many other languages, e.g. Japanese (Maekawa, 2008). The texts were often accompanied by multiple layers of annotations—syntactic, morphological, semantic, genre, source etc. This enabled valuable empirical research on the variants of represented languages, finding use in early distributional semantic models. Corpus linguistics developed sophisticated methodologies including concordances, word sketches and various word association measures (Stefanowitsch and Gries, 2003; Baker, 2004; Kilgarriff, 2014, among others). However, this methodology did not adapt well to Web-scale corpora due to the lack of tools and resources that could support such scale.

Web-scale corpora for LLM pre-training. As LLMs grew, so did the need for massive pre-training datasets. To date, there were several efforts to collect and clean large English and multilingual corpora (Raffel et al., 2020; Xue et al., 2021; Gao et al., 2020; Ortiz Suárez et al., 2020; Bañón et al., 2020; El-Kishky et al., 2020). Non-English, monolingual corpora of this scale have also started to emerge (Gutiérrez-Fandiño et al., 2022; Kummer-vold et al., 2022) However, the sheer scale of such datasets renders them hard to properly curate: we now know that the data used for training LLMs may contain synthetic data (Dodge et al., 2021), privacy-infringing data (Carlini et al., 2020; Huang et al., 2022), incorrect language codes or and translations (Kreutzer et al., 2022), as well as the ubiquitous issues with social biases (Blodgett et al., 2020; Field et al., 2021; Stanczak and Augenstein, 2021, among others). Another issue pertains to the permissions to use the data, which, perhaps the most famously, surfaced in relation to the BookCorpus (Zhu et al., 2015), used, among others, to train BERT (Devlin et al., 2019), but collected without author permissions and eventually taken down by the authors (Bandy and Vincent, 2021).

These issues are a consequence of the fact that the current web-scale corpora are opportunistic samples of publicly available text, rather than artifacts curated to provide a representative snapshot of a specific language variety, as in the corpus linguistics work (Rogers, 2021). This highlights the general problem with the lack of documentation in

NLP datasets of all sizes (Bender and Friedman, 2018; Gebru et al., 2020), and the fact that data work has generally not been a priority in NLP recently (Sambasivan et al., 2021).

Information Retrieval for massive text corpora.

Inspecting large data collection is a central topic of study in another Machine Learning domain, namely Information Retrieval. Even though multiple techniques for analysing large document collections have been developed over the years, there has been little interest so far in applying them specifically to study LLM training data. The closest to our work is that of Dodge et al. (2021) who analyze the C4 dataset (Raffel et al., 2022) and also provide a searchable index.⁴ Similar tools emerge for smaller, more specialised corpora, e.g. COVID-related datasets (Zhang et al., 2020), news quotes (Vuković et al., 2022) and medical literature (Niezni et al., 2022). Razeghi et al. (2022) provide an interface to pre-computed term frequencies from the Pile, but it does not provide full-text corpus search. In the Computer Vision community, related efforts⁵ target large text and image datasets such as LAION (Schuhmann et al., 2022, 2021).

We believe our work to be the first to provide both fuzzy and exact search access to the training corpus of an existing large language model.

3 The ROOTS corpus

The ROOTS corpus (Laurençon et al., 2022) is a high-quality, heterogeneous, multilingual text corpus collected as part of the BigScience project to train the BLOOM LLM (Scao et al., 2022). ROOTS consists of 1.6TB of data in 46 natural and 13 programming languages. The full ROOTS dataset is open to the members of the BigScience Data organization on the Hugging Face hub, which the interested researchers can still apply to join⁶.

3.1 Data Governance

The development of the BLOOM model within the BigScience project was backed by significant work on data governance, as it was identified early on as one of the highest-impact levers of action to enable better accountability and data subject agency in modern ML technology⁷. Participants started by designing a new governance framework to meet

⁴<https://c4-search.apps.allenai.org/>

⁵<https://haveibeentrained.com/>

⁶Sign-up link is available [here](#)

⁷[Data governance and representation in BigScience.](#)

the unique needs of distributed data governance for web-scale data in terms of respecting data subject rights (Jernite et al., 2022). A partial implementation of this framework was used for the ROOTS data as described by Laurençon et al. (2022), focusing on explicit agreements with data custodians, extensive documentation of the data sources, technical tools for privacy-enhancing data handling, and purpose-specific access to subsets of the data.

The present tool goes one step further in implementing the proposed data governance feedback by enabling examination and feedback for the data sources from any interested parties; while still maintaining the controlled access necessary to the proposed governance. The tool only provides 128-word snippets of indexed documents, akin to regular web search engines, and hence provides no practical way to reconstruct the full corpus. The snippets are traceable to their origin in the full ROOTS corpus, and we additionally link to original source documents whenever possible.⁸ Finally, users of the tool are able to flag specific search results and provide an explanation outlining possible infringements of data subjects’ privacy or intellectual property rights. At this stage, the information collected from the flagging process is primarily intended to serve as a basis for future research on collaborative data governance processes. We provide more examples of use cases to support data examination and governance in Section 5.

3.2 Data Pre-processing

Documents vs snippets. ROOTS consists of documents of varying lengths, with outliers as long as 282,571 words. For fuzzy search, we split documents into short snippets of at most 128 words and index snippets rather than the original documents. This helps us follow the controlled access principle discussed in the previous section and makes indexed snippets more comparable in the context of fuzzy search. In exact search, we look for the exact occurrences of the input query within documents and construct snippets ad hoc, including words on both sides of the detected occurrence.

Unique Result IDs. In order to be able to trace search results back to their source, we construct result IDs, adopting the following convention: (a) we include the dataset name as defined on the Hugging Face Hub, followed by (b) the ID of the docu-

⁸The metadata in ROOTS is inconsistent and we only have access to URLs in the pseudocrawl datasets.

Document ID: roots_en_no_code_stackexchange/1495630?seg=para_128_8&seg_id=6
 Flag result
 Murch, REDACTED EMAIL Mustafa, REDACTED EMAIL Nathaniel
 Mahieu, REDACTED EMAIL Nick Bosma, REDACTED EMAIL Nick, REDACTED
 EMAIL Nicolas Benoit, REDACTED EMAIL NicolasDorier, REDACTED EMAIL
 NicolasDorier, REDACTED EMAIL Nicolas Dorier, REDACTED EMAIL
 NicolasDorier, REDACTED EMAIL Nicolas DORIER, REDACTED EMAIL Nils
 Schneider, REDACTED EMAIL Nils Schneider, REDACTED EMAIL Nils
 Schneider, REDACTED EMAIL Noel Tiernan, REDACTED EMAIL

Figure 2: PII leakage: example result for the query gmail.com. We indicate the redacted PII with green and pink treatment.

ment from which the given snippet came, (c) and a question mark. We then include parameters which differ depending on the search strategy used. In fuzzy search we introduce two parameters: the seg parameter describing the segmentation strategy applied during the pre-processing stage, and the seg_id parameter indicating the rank of the given snippet under the specified segmentation strategy. For exact search, we include a single id parameter indicating the the rank of the occurrence of the query in the current document.

PII redaction. During preliminary experiments on the ROOTS corpus, OSCAR (Ortiz Suárez et al., 2019) has been identified as a source of a large amount of documents containing personally identifiable information (PII). A regular-expression-based PII redaction script⁹ has been applied to OSCAR prior to BLOOM training. However, the dataset itself still contains unredacted text. In order to avoid leaking PII through our search tool, we apply an improved variant of the BigScience PII redaction script on the backend side and display results with PII redacted in a visible way - this way one can inspect the data and observe the problem, but personal information are predominantly removed. An example is shown in Figure 2.

4 Implementation

Fuzzy Search Backend. The ROOTS corpus is organized in 498 datasets, each annotated with a language identifier. There are two types of identifiers: those indicating an individual language (e.g. pt for Portuguese), and those indicating a language within a language group (e.g. indic-mr for Marathi, as part of the Indic language group). All programming languages are collected under a common code tag. We build 13 sparse, BM25 (Robertson, 2009) indices: one per language group for the

⁹The BigScience PII redaction script is available [here](#)

ROOTS language tag	# documents	Data size (GB)	# snippets	Index size (GB)	Analyzer
zh, zhs, zht	88,814,841	259.01	111,284,681	682	zh
indic	84,982,982	70.45	100,810,124	714.08	whitespace
en	77,010,827	470.47	695,521,432	766.14	en
es	67,005,817	172.40	267,542,136	264.35	es
fr	58,847,091	204.03	299,938,546	305.29	fr
vi	34,110,375	42.83	76,164,552	72.89	whitespace
pt	31,969,891	77.59	122,221,863	119.98	pt
code	26,176,998	173.16	365,424,222	206.96	whitespace
ar	15,234,080	73.75	68,509,441	93.71	ar
id	12,514,253	19.63	29,531,873	27.16	id
ca	6,142,390	17.42	26,844,600	29.65	es
eu	5,149,797	2.36	6,219,039	4.56	whitespace
nigercongo	1,162,568	0.48	1,462,238	0.89	whitespace
total	597,936,751	1583.59	2,171,474,747	2518.99	

Table 1: Each row represents a single BM25 index we build.

indic and nigercongo groups, one for code, and one for each of the remaining languages (except Chinese, where we combine the tags zh, zht, and zhs into a single index). Table 1 presents the basic information per index. We index respective subsets of the corpus using Pyserini (Lin et al., 2021), a leading toolkit for reproducible IR research. Tokenization is performed with native Lucene¹⁰ analyzers available via Pyserini API (see Table 1 to check which analyzers were used for specific indices).

Exact Search Backend. We leverage a suffix array implementation¹¹ proposed by Lee et al. (2022). We build the suffix array for the whole ROOTS corpus, this time without the split into languages or language groups. We host both the BM25 indices and the suffix array on Hugging Face-provisioned machines. The server code is open-sourced¹².

Frontend and User Experience. The ROOTS Search Tool user interface is built with Gradio (Abid et al., 2019) and served via Hugging Face Spaces.¹³ By default, searches are performed in fuzzy mode, in order to move to the exact search one can enclose the query in double quotes. Fuzzy searches can be performed in a user-specified language, or in *all* languages (in that case results are surfaced separately for each language). We also provide an option to auto-detect the language of the query with a FastText classifier (Joulin et al., 2017). Results are displayed in the order of decreasing relevance; users can control the maximum

number of results they want to see using a slider. In exact search mode, the backend returns all documents matching a given query exactly irrespective of the language, and they are displayed over multiple pages in a random order, with the max results parameter controlling the size of a single page. The total number of matched results is displayed at the top of the results page. PII redaction is applied to all results on the backend side. The tool also allows users to filter out all results from a specific dataset appearing on a given page.

5 Use cases

Detecting PII issues to improve obfuscation. BLOOM was trained with efforts to detect and obfuscate personally identifiable information, or PII (e.g. email and personal addresses, age, phone numbers or government-issued identifiers such as license plates) in the original ROOTS documents. As described in section 3.2, we build on that effort when obfuscating PII in search results. However, it is still possible that some such data was not detected. The tool allows searching for the specific PII by concerned individuals, which is the first step for requesting removal of their data. One could also simply search for their name to see if they are represented in the corpus, and how.

Detecting undesired content. Text from Web crawls contains all kinds of undesired content (Lucioni and Viviano, 2021). Examples of possible classes of problems include hate speech, pornography, synthetic text (e.g. machine-translated text, AI-generated text), word lists that are not meaningful and are meant to trick search engines (Hamilton, 2013), factually incorrect text such as fake news

¹⁰<https://lucene.apache.org/>

¹¹<https://github.com/google-research/deduplicate-text-datasets>

¹²<https://github.com/huggingface/roots-search-tool>

¹³<https://huggingface.co/docs/hub/spaces>

or conspiracy theories, among others. For example, we found at least 5 snippets from the OSCAR source incorrectly arguing that Barack Obama was born in Kenya. While the creators of ROOTS employed filtering strategies targeted specifically at spam and machine-generated content (Laurençon et al., 2022), developing filters for such content is a never-ending arms race with its producers, and the only way to keep improving them is to look at the data—which our tool enables.

Studying representation of dialects and social groups.

When LLM-based systems are deployed, the implicit assumption is often that they are general-purpose and can serve all of its potential users equally well. But there is no such thing as a “neutral”, one-size-fits-all corpus (Rogers, 2021). An obvious issue is dialects, and in case of multilingual models like BLOOM another obvious problem is language imbalance. Besides that, the training data may not equally represent the topics and sources associated with different demographic groups, and hence the LLM would likely not cater to them equally well. Bender et al. (2021) cite the example of GPT-2: the filter for its sources was that they were shared on Reddit, which overrepresents the interests of the typical Reddit user (of whom in the US 67% are men, and 64% are 18-29 years old).

Training data that is then likely to reinforce social stereotypes harmful to marginalized populations. For example, GPT-3 has been shown to over-associate Muslims with violence (Abid et al., 2021). In particular, prompting the model to continue “*Two Muslims walked into...*” tends to lead to mentions of terrorism or assault. BLOOM is not free from these biases: we sampled 10 completions and found 4 that mentioned guns or death (compared to 66% reported for GPT-3). Exact search for “*Two Muslims Walked into...*” returned examples of papers studying this very phenomenon, but a search for just “*Two Muslims*” shows that many passages in OSCAR mention violence or terrorism, whereas mentions in Semantic Scholar, pseudo-crawled websites, and Wikipedia are more varied.

Detecting the presence of specific information.

Where the suitability of a model to a given application depends on it being up-to-date with the latest events, or knowledge about a given fact, a tool like ours can help to quickly find out if the model even theoretically could “learn” a given fact.

For instance, ROOTS contains 231 references to the *death of Queen Elizabeth*, but they refer to the death Elizabeth I in 1603 and not to the recent passing of Elizabeth II in 2022.

Detecting plagiarism/memorization. Generative LLMs can memorize part of their training sets and repeat it verbatim in their outputs. We can probe an LLM to elicit candidates for data memorization (Carlini et al., 2020), and the ROOTS Search Tool can help in different ways:

- By conditioning model probing on actual training data, so that we can more easily check whether such data has been memorized;
- By providing the ground truth to verify that model output was part of the training data;
- By providing the ground truth to verify that model did have a chance to memorize something that it should have memorized;
- By providing match counts to identify which data was more likely to be memorized (since the number of copies in the training data influences memorization (Kandpal et al., 2022)).

For example, BLOOM correctly completes Prince Hamlet’s *To be or not to be* soliloquy—both using greedy decoding and nucleus sampling—but not the less popular Shakespeare quote *I am in this earthly world, where to do harm... is often laudable, to do good sometime accounted dangerous folly*. With our tool we verified that BLOOM had access to at least 7 sources for the *Macbeth* quote (vs at least 47 for *Hamlet*), but did not “learn” it.

Verifying originality. An important question about generative AI models is to what extent their output – that is not a verbatim copy of training data – can be considered original. Consider the above quote from *Macbeth*, which BLOOM completed for us as follows: “*I am in this earthly world, where to do harm... is to do good, and to do good is to do harm.*” With our tool, we could easily verify that the suggested completion does not exist in the corpus verbatim. However, there are dozens of contexts where the concepts of “good” and “harm” are mentioned close to each other (esp. in the phrase “do more harm than good”), so they were the likely indirect sources for this completion. To what degree that completion can be considered new, original text is a key question for the current discussions on plagiarism in AI writing assistants and the legal status of their output.

Non-existing facts. When the same associative mechanism generates factoid text, the model may “hallucinate” events that never occurred—or at least, there was no evidence on which the model could draw. This, too, becomes easy to verify with our tool. BLOOM completed the prompt “*When was the Golden Gate Bridge transported for the second time across Egypt?*” (Hofstadter, 2022) with “*The first time was in the late 19th century, when the bridge was transported from San Francisco to Cairo*”. Of course, this “fact” is untrue, and was not mentioned in the corpus. But we could not even find mentions of anything else transported from San Francisco to Cairo. How exactly LLMs come up with such generations is an interesting research problem, for which tools like ours could be useful.

Enabling data removal requests. The authors of texts that were included in web crawls could use such a tool to identify that fact and request the removal of their texts. For ROOTS, the data governance structure set up for Big Science workshop operated only for its duration, but should there be any future work relying on the same data hosts and agreements, the flagged data collected through our tool can be used to honor the removal requests.

Benchmark data contamination. To interpret benchmark results, we need to know whether they reflect training data memorization or generalization. One approach is for the model authors to specifically plan for the evaluation benchmarks prior to training, and try to exclude the benchmark data (Brown et al., 2020), but this limits the options for external evaluation. Our tool enables sampled checks of benchmark data, and was already successfully used to find¹⁴ that BLOOM should not be evaluated on XNLI (Conneau et al., 2018).

Language contamination. According to Laurençon et al. (2022), ROOTS contains data in 46 languages. But this is clearly not the full story. For example, neither Danish nor Ukrainian are listed, but we found examples in these languages (stackexchange, OSCAR, parsed academic pdf data). The tool can thus be useful for investigating the transfer to “unseen” languages in multilingual evaluation.

Word sense disambiguation. Since the ROOTS Search Tool provides context paragraphs, it can be used to check in what sense a word was used in

¹⁴<https://twitter.com/WilliamBarrHeld/status/1586090252946448384>

the training data. For example, the acronym LLM in ROOTS is used as “large language model” in the parsed academic article data, but in OSCAR it means predominantly “limited liability company” or “Legum Magister”. If future work extends our approach to providing search results through API, then quantitative research would also be possible with techniques like context clustering and classification.

Pre-processing issues. By searching for phrases occurring in different parts of the same document, it is possible to verify that the entire document made it through the pre-processing pipeline – which is useful for improving it. For example, we found a news article in OSCAR, the initial paragraphs of which are missing from ROOTS.

6 Limitations and Future Work

A major limitation of this work is that to mitigate possible issues on the data governance side, we can only provide short snippets of the indexed texts, as is typical of web search engines. We strive to provide links to the original text sources, but this metadata is not consistently available in ROOTS.

Implementation-wise, the current version of exact search is exact down to capitalization and punctuation, and fuzzy search can be noticeably slower. These issues will be addressed in future versions.

The current tool is heavily influenced by the UX of search engines, and its core functionality is similar. In future we intend to review classic corpus analysis tools for ideas of different presentation modes, such as concordance and word sketches. We would like to add more quantitative information, e.g. term frequency information, number of hits, and co-occurrence statistics. Community feedback and suggestions are welcome in the [Community](#) tab of the demo. We are also pursuing a spin-off collaboration with Pyserini to make large scale indexing and hosting of textual data even more seamless.

7 Acknowledgements

We thank the Pyserini team—Odunayo Ogundepo, Xinyu Zhang, Akintunde Oladipo and Jimmy Lin, for their indexing insights. Big thanks to the Gradio team, especially Pete Allen, Abubakar Abid and Freddy Boulton for their support on the front-end side, and to the Hugging Face infra team for answering questions regarding hosting the tool. We thank Carlos Muñoz Ferrandis and Meg Mitchell for valuable discussions.

8 Impact Statement

Our tool aims to improve the current state of documentation search for large corpora of web-scraped text, starting with the ROOTS corpus. However, it also comes with ethical considerations: for instance, it can also inadvertently display sensitive information such as PII and harmful content, and help malicious actors find information about a given topic from multiple sources (which is more difficult given only the raw text of the corpus). We are aware of these limitations, and have taken precautions to compensate for them, such as the PII redaction measures we present in Figure 2. We also present only a snippet of the raw text, which means that for accessing the full documents, users must sign up to be a part of the Big Science organization on the Hugging Face Hub, which also reduces the amount of information that potentially malicious anonymous users can access.

References

- Abubakar Abid, Ali Abdalla, Ali Abid, Dawood Khan, Abdulrahman Alfozan, and James Zou. 2019. Gradio: Hassle-free sharing and testing of ml models in the wild. *arXiv preprint arXiv:1906.02569*.
- Abubakar Abid, Maheen Farooqi, and James Zou. 2021. [Persistent anti-muslim bias in large language models](#). In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '21, page 298–306, New York, NY, USA. Association for Computing Machinery.
- Christopher Akiki, Giada Pistilli, Margot Mieskes, Matthias Gallé, Thomas Wolf, Suzana Ilic, and Yacine Jernite. 2022. [Bigscience: A case study in the social construction of a multilingual large language model](#). In *Workshop on Broadening Research Collaborations 2022*.
- Paul Baker. 2004. [Querying Keywords: Questions of Difference, Frequency, and Sense in Keywords Analysis](#). *Journal of English Linguistics*, 32(4):346–359.
- Jack Bandy and Nicholas Vincent. 2021. [Addressing "documentation debt" in machine learning research: A retrospective datasheet for bookcorpus](#).
- Marta Bañón, Pinzhen Chen, Barry Haddow, Kenneth Heafield, Hieu Hoang, Miquel Esplà-Gomis, Mikel L. Forcada, Amir Kamran, Faheem Kirefu, Philipp Koehn, Sergio Ortiz Rojas, Leopoldo Pla Sempere, Gema Ramírez-Sánchez, Elsa Sarrías, Marek Strelec, Brian Thompson, William Waites, Dion Wiggins, and Jaume Zaragoza. 2020. [ParaCrawl: Web-Scale Acquisition of Parallel Corpora](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4555–4567, Online. Association for Computational Linguistics.
- Emily M. Bender and Batya Friedman. 2018. [Data Statements for Natural Language Processing: Toward Mitigating System Bias and Enabling Better Science](#). *Transactions of the Association for Computational Linguistics*, 6:587–604.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?](#) In *FACt '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 610–623.
- Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. [Language \(Technology\) is Power: A Critical Survey of "Bias" in NLP](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5476, Online. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, Alina Oprea, and Colin Raffel. 2020. [Extracting Training Data from Large Language Models](#). *arXiv:2012.07805 [cs]*.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. [XNLI: Evaluating cross-lingual sentence representations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485, Brussels, Belgium. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [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)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jesse Dodge, Maarten Sap, Ana Marasović, William Agnew, Gabriel Ilharco, Dirk Groeneveld, Margaret Mitchell, and Matt Gardner. 2021. [Documenting](#)

- Large Webtext Corpora: A Case Study on the Colossal Clean Crawled Corpus. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1286–1305, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Ahmed El-Kishky, Vishrav Chaudhary, Francisco Guzmán, and Philipp Koehn. 2020. **CCAligned: A Massive Collection of Cross-Lingual Web-Document Pairs**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5960–5969, Online. Association for Computational Linguistics.
- Anjalie Field, Su Lin Blodgett, Zeerak Waseem, and Yulia Tsvetkov. 2021. **A Survey of Race, Racism, and Anti-Racism in NLP**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1905–1925, Online. Association for Computational Linguistics.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, Shawn Presser, and Connor Leahy. 2020. **The Pile: An 800GB Dataset of Diverse Text for Language Modeling**. *arXiv:2101.00027 [cs]*.
- Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2020. **Datasheets for Datasets**. *arXiv:1803.09010 [cs]*.
- Asier Gutiérrez-Fandiño, David Pérez-Fernández, Jordi Armengol-Estapé, David Griol, and Zoraida Callejas. 2022. **esCorpius: A Massive Spanish Crawling Corpus**.
- Peter A. Hamilton. 2013. Google-bombing - manipulating the pagerank algorithm. *Information Retrieval*.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Jack W. Rae, Oriol Vinyals, and Laurent Sifre. 2022. **Training compute-optimal large language models**.
- Douglas Hofstadter. 2022. **Artificial neural networks today are not conscious, according to Douglas Hofstadter**. *The Economist*.
- Jie Huang, Hanyin Shao, and Kevin Chen-Chuan Chang. 2022. **Are Large Pre-Trained Language Models Leaking Your Personal Information?**
- Yacine Jernite, Huu Nguyen, Stella Biderman, Anna Rogers, Maraim Masoud, Valentin Danchev, Samson Tan, Alexandra Sasha Luccioni, Nishant Subramani, Isaac Johnson, Gerard Dupont, Jesse Dodge, Kyle Lo, Zeerak Talat, Dragomir Radev, Aaron Gokaslan, So-maieh Nikpoor, Peter Henderson, Rishi Bommasani, and Margaret Mitchell. 2022. **Data governance in the age of large-scale data-driven language technology**. In *2022 ACM Conference on Fairness, Accountability, and Transparency, FAccT '22*, page 2206–2222, New York, NY, USA. Association for Computing Machinery.
- Eun Seo Jo and Timnit Gebru. 2020. **Lessons from archives**. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. ACM.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. **Bag of tricks for efficient text classification**. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 427–431, Valencia, Spain. Association for Computational Linguistics.
- Nikhil Kandpal, Eric Wallace, and Colin Raffel. 2022. **Deduplicating Training Data Mitigates Privacy Risks in Language Models**. In *Proceedings of the 39th International Conference on Machine Learning*, pages 10697–10707. PMLR.
- Adam Kilgarriff. 2014. **The Sketch Engine: Ten years on**. *Lexicography*, pages 1–30.
- Julia Kreutzer, Isaac Caswell, Lisa Wang, Ahsan Wahab, Daan van Esch, Nasanbayar Ulzii-Orshikh, Allahsera Tapo, Nishant Subramani, Artem Sokolov, Claytone Sikasote, Monang Setyawan, Supheakmungkol Sarin, Sokhar Samb, Benoît Sagot, Clara Rivera, Annette Rios, Isabel Papadimitriou, Salomey Osei, Pedro Ortiz Suarez, Iro-ro Orife, Kelechi Ogueji, Andre Niyongabo Rubungo, Toan Q. Nguyen, Mathias Müller, André Müller, Shamsuddeen Hassan Muhammad, Nanda Muhammad, Ayanda Mnyakeni, Jamshidbek Mirzakhlov, Tapiwanashe Matangira, Colin Leong, Nze Lawson, Sneha Kudugunta, Yacine Jernite, Mathias Jenny, Orhan Firat, Bonaventure F. P. Dossou, Sakhile Dlamini, Nisansa de Silva, Sakine Çabuk Ballı, Stella Biderman, Alessia Battisti, Ahmed Baruwa, Ankur Bapna, Pallavi Baljekar, Israel Abebe Azime, Ayodele Awokoya, Duygu Ataman, Orevaoghene Ahia, Oghenefego Ahia, Sweta Agrawal, and Mofetoluwa Adeyemi. 2022. **Quality at a Glance: An Audit of Web-Crawled Multilingual Datasets**. *Transactions of the Association for Computational Linguistics*, 10:50–72.
- Per Kummervold, Freddy Wetjen, and Javier de la Rosa. 2022. **The Norwegian Colossal Corpus: A Text Corpus for Training Large Norwegian Language Models**. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 3852–3860, Marseille, France. European Language Resources Association.
- Hugo Laurençon, Lucile Saulnier, Thomas Wang, Christopher Akiki, Albert Villanova del Moral, Teven Le Scao, Leandro Von Werra, Chenghao Mou, Eduardo González Ponferrada, Huu Nguyen, Jörg

- Frohberg, Mario Šaško, Quentin Lhoest, Angelina McMillan-Major, Gérard Dupont, Stella Biderman, Anna Rogers, Loubna Ben allal, Francesco De Toni, Giada Pistilli, Olivier Nguyen, Somaieh Nikpoor, Maraim Masoud, Pierre Colombo, Javier de la Rosa, Paulo Villegas, Tristan Thrush, Shayne Longpre, Sebastian Nagel, Leon Weber, Manuel Romero Muñoz, Jian Zhu, Daniel Van Strien, Zaid Alyafeai, Khalid Almubarak, Vu Minh Chien, Itziar Gonzalez-Dios, Aitor Soroa, Kyle Lo, Manan Dey, Pedro Ortiz Suarez, Aaron Gokaslan, Shamik Bose, David Ifeoluwa Adelani, Long Phan, Hieu Tran, Ian Yu, Suhas Pai, Jenny Chim, Violette Lepercq, Suzana Ilic, Margaret Mitchell, Sasha Luccioni, and Yacine Jernite. 2022. [The bigscience ROOTS corpus: A 1.6TB composite multilingual dataset](#). In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Katherine Lee, Daphne Ippolito, Andrew Nystrom, Chiyuan Zhang, Douglas Eck, Chris Callison-Burch, and Nicholas Carlini. 2022. [Deduplicating training data makes language models better](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics.
- Geoffrey Neil Leech. 1992. 100 million words of English: The British National Corpus (BNC). *Language Research*, 1/4.
- Jimmy Lin, Xueguang Ma, Sheng-Chieh Lin, Jheng-Hong Yang, Ronak Pradeep, and Rodrigo Nogueira. 2021. [Pyserini: A Python toolkit for reproducible information retrieval research with sparse and dense representations](#). In *Proceedings of the 44th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2021)*, pages 2356–2362.
- Alexandra Luccioni and Joseph Viviano. 2021. [What’s in the box? an analysis of undesirable content in the Common Crawl corpus](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 182–189, Online. Association for Computational Linguistics.
- Kikuo Maekawa. 2008. [Balanced Corpus of Contemporary Written Japanese](#). In *Proceedings of the Third International Joint Conference on Natural Language Processing (IJCNLP)*, pages 101–102.
- Tony McEnery and Andrew Hardie. 2013. [The History of Corpus Linguistics](#). In *The Oxford Handbook of the History of Linguistics*. Oxford University Press.
- Margaret Mitchell, Alexandra Sasha Luccioni, Nathan Lambert, Marissa Gerchick, Angelina McMillan-Major, Ezinwanne Ozoani, Nazneen Rajani, Tristan Thrush, Yacine Jernite, and Douwe Kiela. 2022. [Measuring data](#). *CoRR*, abs/2212.05129.
- Danna Niezni, Hillel Taub-Tabib, Yuval Harris, Hagit Sason-Bauer, Yakir Amrusi, Dana Azagury, Maytal Avrashami, Shaked Launer-Wachs, Jon Borchardt, M Kusold, Aryeh Tiktinsky, Tom Hope, Yoav Goldberg, and Yosi Shamay. 2022. [Extending the boundaries of cancer therapeutic complexity with literature data mining](#). *bioRxiv*.
- Pedro Javier Ortiz Suárez, Laurent Romary, and Benoît Sagot. 2020. [A Monolingual Approach to Contextualized Word Embeddings for Mid-Resource Languages](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1703–1714, Online. Association for Computational Linguistics.
- Pedro Javier Ortiz Suárez, Benoît Sagot, and Laurent Romary. 2019. [Asynchronous pipelines for processing huge corpora on medium to low resource infrastructures](#). In *7th Workshop on the Challenges in the Management of Large Corpora (CMLC-7)*, Proceedings of the Workshop on Challenges in the Management of Large Corpora (CMLC-7) 2019. Cardiff, 22nd July 2019, pages 9 – 16, Mannheim. Leibniz-Institut für Deutsche Sprache.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2022. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *The Journal of Machine Learning Research*, 21(1):140:5485–140:5551.
- Yasaman Razeghi, Raja Sekhar Reddy Mekala, Robert L Logan Iv, Matt Gardner, and Sameer Singh. 2022. [Snoopy: An Online Interface for Exploring the Effect of Pretraining Term Frequencies on Few-Shot LM Performance](#). In *Proceedings of the The 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 389–395, Abu Dhabi, UAE. Association for Computational Linguistics.
- S. Robertson. 2009. [The Probabilistic Relevance Framework: BM25 and Beyond](#). *Foundations and Trends® in Information Retrieval*, 3(4):333–389.
- Anna Rogers. 2021. [Changing the World by Changing the Data](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2182–2194, Online. Association for Computational Linguistics.
- Nithya Sambasivan, Shivani Kapania, Hannah Highfill, Diana Akrong, Praveen Paritosh, and Lora M Aroyo. 2021. ["Everyone wants to do the model work, not the data work": Data Cascades in High-Stakes AI](#). In *Proceedings of the 2021 CHI Conference on Human*

Factors in Computing Systems, CHI '21, pages 1–15, New York, NY, USA. Association for Computing Machinery.

Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, Jonathan Tow, Alexander M. Rush, Stella Biderman, Albert Webson, Pawan Sasanka Amanamanchi, Thomas Wang, Benoît Sagot, Niklas Muennighoff, Albert Villanova del Moral, Olatunji Ruwase, Rachel Bawden, Stas Bekman, Angelina McMillan-Major, Iz Beltagy, Huu Nguyen, Lucile Saulnier, Samson Tan, Pedro Ortiz Suarez, Victor Sanh, Hugo Laurençon, Yacine Jernite, Julien Launay, Margaret Mitchell, Colin Raffel, Aaron Gokaslan, Adi Simhi, Aitor Soroa, Alham Fikri Aji, Amit Alfassy, Anna Rogers, Ariel Kreisberg Nitzav, Canwen Xu, Chenghao Mou, Chris Emezue, Christopher Klamm, Colin Leong, Daniel van Strien, David Ifeoluwa Adelani, Dragomir Radev, Eduardo González Ponferrada, Efrat Levkovich, Ethan Kim, Eyal Bar Natan, Francesco De Toni, Gérard Dupont, Germán Kruszewski, Giada Pistilli, Hady Elsahar, Hamza Benyamina, Hieu Tran, Ian Yu, Idris Abdulmumin, Isaac Johnson, Itziar Gonzalez-Dios, Javier de la Rosa, Jenny Chim, Jesse Dodge, Jian Zhu, Jonathan Chang, Jörg Frohberg, Joseph Tobing, Joydeep Bhattacharjee, Khalid Almubarak, Kimbo Chen, Kyle Lo, Leandro Von Werra, Leon Weber, Long Phan, Loubna Ben allal, Ludovic Tanguy, Manan Dey, Manuel Romero Muñoz, Maraim Masoud, María Grandury, Mario Šaško, Max Huang, Maximin Coavoux, Mayank Singh, Mike Tian-Jian Jiang, Minh Chien Vu, Mohammad A. Jauhar, Mustafa Ghaleb, Nishant Subramani, Nora Kassner, Nuru-laqilla Khamis, Olivier Nguyen, Omar Espejel, Ona de Gibert, Paulo Villegas, Peter Henderson, Pierre Colombo, Priscilla Amuok, Quentin Lhoest, Rheza Harliman, Rishi Bommasani, Roberto Luis López, Rui Ribeiro, Salomey Osei, Sampo Pyysalo, Sebastian Nagel, Shamik Bose, Shamsuddeen Hassan Muhammad, Shanya Sharma, Shayne Longpre, So-maieh Nikpoor, Stanislav Silberberg, Suhas Pai, Sydney Zink, Tiago Timponi Torrent, Timo Schick, Tristan Thrush, Valentin Danchev, Vassilina Nikoulina, Veronika Laippala, Violette Lepercq, Vrinda Prabhu, Zaid Alyafeai, Zeerak Talat, Arun Raja, Benjamin Heinzerling, Chenglei Si, Davut Emre Taşar, Elizabeth Salesky, Sabrina J. Mielke, Wilson Y. Lee, Abheesht Sharma, Andrea Santilli, Antoine Chaffin, Arnaud Stiegler, Debajyoti Datta, Eliza Szczechla, Gunjan Chhablani, Han Wang, Harshit Pandey, Hendrik Strobelt, Jason Alan Fries, Jos Rozen, Leo Gao, Lintang Sutawika, M. Saiful Bari, Maged S. Al-shaibani, Matteo Manica, Nihal Nayak, Ryan Teehan, Samuel Albanie, Sheng Shen, Srulik Ben-David, Stephen H. Bach, Taewoon Kim, Tali Bers, Thibault Fevry, Trishala Neeraj, Urmish Thakker, Vikas Raunak, Xiangru Tang, Zheng-Xin Yong, Zhiqing Sun, Shaked Brody, Yallow Uri, Hadar Tojarieh, Adam Roberts, Hyung Won Chung, Jaesung Tae, Jason Phang, Ofir Press, Conglong Li, Deepak Narayanan, Hatim Bourfoune, Jared Casper,

Jeff Rasley, Max Ryabinin, Mayank Mishra, Minjia Zhang, Mohammad Shoeybi, Myriam Peyrounette, Nicolas Patry, Nouamane Tazi, Omar Sanseviero, Patrick von Platen, Pierre Cornette, Pierre François Lavallée, Rémi Lacroix, Samyam Rajbhandari, San-chit Gandhi, Shaden Smith, Stéphane Reuena, Suraj Patil, Tim Dettmers, Ahmed Baruwā, Amanpreet Singh, Anastasia Cheveleva, Anne-Laure Ligozat, Arjun Subramonian, Aurélie Névél, Charles Lovering, Dan Garrette, Deepak Tunuguntla, Ehud Reiter, Ekaterina Taktasheva, Ekaterina Voloshina, Eli Bogdanov, Genta Indra Winata, Hailey Schoelkopf, Jan-Christoph Kalo, Jekaterina Novikova, Jessica Zosa Forde, Jordan Clive, Jungo Kasai, Ken Kawamura, Liam Hazan, Marine Carpuat, Miruna Clinciu, Najoung Kim, Newton Cheng, Oleg Serikov, Omer Antverg, Oskar van der Wal, Rui Zhang, Ruochen Zhang, Sebastian Gehrmann, Shachar Mirkin, Shani Pais, Tatiana Shavrina, Thomas Scialom, Tian Yun, Tomasz Limisiewicz, Verena Rieser, Vitaly Protasov, Vladislav Mikhailov, Yada Pruksachatkun, Yonatan Belinkov, Zachary Bamberger, Zdeněk Kasner, Alice Rueda, Amanda Pestana, Amir Feizpour, Ammar Khan, Amy Faranak, Ana Santos, Anthony Hevia, Antigona Undreaaj, Arash Aghagol, Are-zoo Abdollahi, Aycha Tammour, Azadeh HajiHosseini, Bahareh Behroozi, Benjamin Ajibade, Bharat Saxena, Carlos Muñoz Ferrandis, Danish Contractor, David Lansky, Davis David, Douwe Kiela, Duong A. Nguyen, Edward Tan, Emi Baylor, Ezinwanne Ozoani, Fatima Mirza, Frankline Onon-iwu, Habib Rezanejad, Hessie Jones, Indrani Bhat-tacharya, Irene Solaiman, Irina Sedenko, Isar Ne-jadgholi, Jesse Passmore, Josh Seltzer, Julio Bonis Sanz, Livia Dutra, Mairon Samagaio, Maraim El-badri, Margot Mieskes, Marissa Gerchick, Martha Akinlolu, Michael McKenna, Mike Qiu, Muhammed Ghauri, Mykola Burynok, Nafis Abrar, Nazneen Ra-jani, Nour Elkott, Nour Fahmy, Olanrewaju Samuel, Ran An, Rasmus Kromann, Ryan Hao, Samira Al-izadeh, Sarmad Shubber, Silas Wang, Sourav Roy, Sylvain Viguier, Thanh Le, Tobi Oyebade, Trieu Le, Yoyo Yang, Zach Nguyen, Abhinav Ramesh Kashyap, Alfredo Palasciano, Alison Callahan, Anima Shukla, Antonio Miranda-Escalada, Ayush Singh, Benjamin Beilharz, Bo Wang, Caio Brito, Chenxi Zhou, Chirag Jain, Chuxin Xu, Clémentine Fourrier, Daniel León Perriñán, Daniel Molano, Dian Yu, Enrique Manjavac-cas, Fabio Barth, Florian Fuhrmann, Gabriel Altay, Giyaseddin Bayrak, Gully Burns, Helena U. Vrabec, Imane Bello, Ishani Dash, Jihyun Kang, John Giorgi, Jonas Golde, Jose David Posada, Karthik Rangas-sai Sivaraman, Lokesh Bulchandani, Lu Liu, Luisa Shinzato, Madeleine Hahn de Bykhovetz, Maiko Takeuchi, Marc Pàmies, Maria A. Castillo, Mari-anna Nezhurina, Mario Sängler, Matthias Samwald, Michael Cullan, Michael Weinberg, Michiel De Wolf, Mina Mihaljcic, Minna Liu, Moritz Freidank, Myung-sun Kang, Natasha Seelam, Nathan Dahlberg, Nicholas Michio Broad, Nikolaus Muellner, Pascale Fung, Patrick Haller, Ramya Chandrasekhar, Renata Eisenberg, Robert Martin, Rodrigo Canalli, Rosaline Su, Ruisi Su, Samuel Cahyawijaya, Samuele Garda,

- Shlok S. Deshmukh, Shubhanshu Mishra, Sid Kiblawi, Simon Ott, Sinee Sang-aaronsiri, Srishti Kumar, Stefan Schweter, Sushil Bharati, Tanmay Laud, Théo Gigant, Tomoya Kainuma, Wojciech Kusa, Yanis Labrak, Yash Shailesh Bajaj, Yash Venkatraman, Yifan Xu, Yingxin Xu, Yu Xu, Zhe Tan, Zhongli Xie, Zifan Ye, Mathilde Bras, Younes Belkada, and Thomas Wolf. 2022. [BLOOM: A 176B-Parameter Open-Access Multilingual Language Model](#). In *Thirty-Sixth Conference on Neural Information Processing Systems*, New Orleans, Louisiana. arXiv.
- Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade W. Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa R. Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. 2022. [LAION-5B: An open large-scale dataset for training next generation image-text models](#). In *Thirty-Sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. 2021. [LAION-400M: Open Dataset of CLIP-Filtered 400 Million Image-Text Pairs](#). In *Data Centric AI NeurIPS Workshop 2021*.
- Karolina Stanczak and Isabelle Augenstein. 2021. [A Survey on Gender Bias in Natural Language Processing](#).
- Anatol Stefanowitsch and Stefan Th Gries. 2003. [Collostructions: Investigating the interaction of words and constructions](#). *International Journal of Corpus Linguistics*, 8(2):209–243.
- Vuk Vuković, Akhil Arora, Huan-Cheng Chang, Andreas Spitz, and Robert West. 2022. [Quote erat demonstrandum: A web interface for exploring the quotebank corpus](#). In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A Massively Multilingual Pre-trained Text-to-Text Transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.
- Edwin Zhang, Nikhil Gupta, Raphael Tang, Xiao Han, Ronak Pradeep, Kuang Lu, Yue Zhang, Rodrigo Nogueira, Kyunghyun Cho, Hui Fang, and Jimmy Lin. 2020. [Covidex: Neural ranking models and keyword search infrastructure for the COVID-19 open research dataset](#). In *Proceedings of the First Workshop on Scholarly Document Processing*, pages 31–41, Online. Association for Computational Linguistics.
- Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. [Aligning Books and Movies: Towards Story-Like Visual Explanations by Watching Movies and Reading Books](#). In *Proceedings of the IEEE International Conference on Computer Vision*, pages 19–27.