

So, Everything Is Biased... Now What?!

Introducing the Bias-Aware Framework

Amber Zijlma^{1,*†}, Mrinalini Luthra^{1,*†}

¹*Huygens Institute, Koninklijke Nederlandse Akademie van Wetenschappen, Oudezijds Achterburgwal 185, 1012 DK Amsterdam*

Abstract

In the digital humanities, datasets inherit and perpetuate biases through multiple channels: individual and institutional biases, discriminatory language in archives, unequal representation in collection practices, and algorithmic biases in AI-assisted processing. These biases are compounded throughout the research process, yet the term “bias” itself lacks a clear definition, often causing “bias paralysis.” This paper proposes treating “bias” as a productive category of analysis for digital humanities research through the development of a “Bias-Aware Framework” for dataset creation and contextualisation. It has three components: a Bias Thesaurus creating shared vocabulary across disciplines to address the conceptual instability of “bias” by breaking down this nebulous concept into interrelated issues like representation, gaps, positionality, CARE, etc; a Bias-Aware Dataset Lifecycle Model showing where biases enter the research process; and Guidelines for documenting, describing, and mitigating bias. We approach bias not simply as an error, but as a revealing analytical lens that shapes knowledge production. By explicitly describing these conditions of production, researchers can improve transparency, improve dataset documentation, and enable more informed reuse of their data.

Keywords

Bias, Dataset, Data ethics, Archives, AI, Knowledge Production, Digital Humanities

Disclaimer. This paper contains records and metadata from the VOC archives which contain language, descriptions, and omissions that are offensive, biased, and distorted. Please be aware that engaging with this material may cause distress. We advise approaching the content with care and consideration.

1. Introduction: The State of the Arts and its Discontents

The digital humanities face a critical challenge. As computational methods increasingly intersect with humanities research, questions of bias take on renewed urgency. Bias has long been a concept of inquiry in the humanities [1, 2], with scholars examining how perspective and power shape historical narratives and knowledge production. However, the introduction of artificial intelligence (AI) in the field, ranging from automatic transcription to semantic contextualisation, transforms these traditional concerns into urgent technical and ethical challenges [3]. These technologies enable access to massive repositories of historical data but risk perpetuating historical and contemporary biases embedded in archival sources, metadata, tools and technology [4, 5]. This challenge is especially relevant for Semantic Web Technologies, as ontologies, knowledge graphs, and linked data structures can inherit and amplify biases present in original sources, potentially reinforcing problematic knowledge representations.

What makes this challenge particularly concerning is that these diverse forms of bias don’t simply exist in parallel—they compound and amplify each other throughout the research lifecycle. Historical biases in archival selection get encoded into metadata structures, which then inform algorithmic design, creating feedback loops that can dramatically magnify originally subtle imbalances. This cascading effect of bias—from historical sources through metadata design to algorithmic processing—creates a distinctive challenge that requires a systematic approach rather than isolated interventions. This

SemDH 2025: Second International Workshop of Semantic Digital Humanities. Co-located with ESWC 2025, June 02, 2025, Portoroz, Slovenia.

*Corresponding author.

† These authors contributed equally.

✉ amber.zijlma@huygens.knaw.nl (A. Zijlma); mrinalini.luthra@huygens.knaw.nl (M. Luthra)

id 0009-0005-8529-6190 (A. Zijlma); 0000-0001-5348-3673 (M. Luthra)



© 2025 Copyright for this paper by its authors. Use is permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

compounding effect means that seemingly minor biases at early stages can result in significantly skewed outcomes during analysis and interpretation.

1.1. Case Study: Compounding Bias in VOC Testaments

The Dutch East India Company (Verenigde Oost-Indische Compagnie, henceforth VOC) testament archives provide a revealing case study of how digitisation can inadvertently preserve and amplify historical bias.¹ When the Dutch National Archives digitised this collection in 2017, they made the archive accessible online by scanning its pages and digitising a 19th-century index. However, this digital transformation preserved a significant historical bias: while the index includes approximately 10,000 European male testators, it omits female co-testators, individuals from diverse ethnic backgrounds, and enslaved persons who appear as beneficiaries, partners, debtors, properties, or witnesses [6]. Thus, while digitisation increased general accessibility by allowing access across the world, the preservation of the biased indexing structure perpetuates colonial and patriarchal hierarchies, making research on marginalised individuals more challenging [7]. Figure 1 illustrates this bias.²

To address these silences, Luthra et. al. [8] worked with transcribed versions of the testaments and used information retrieval methods to develop more inclusive finding aids. They used named entity recognition and classification (NERC)—a common natural language processing method for identifying and categorising entities such as people, organisations, and locations [9]. However, standard NERC schemas like CoNLL [10] and ACE [11] were insufficient for the complexity of colonial records, which often reference unnamed or marginalised individuals and include vital context like roles, gender, and legal status. By designing a custom typology tailored to colonial archives, the project was able to surface individuals who were previously obscured both by the original records and the digital tools built on them.

This example illustrates a critical point: while techniques like NERC are powerful for information extraction, they can also embed existing power dynamics, reinforcing or challenging historical biases depending on how they are designed and applied. Similarly, in semantic web technologies, choices in ontology design directly influence whose histories are made visible and whose remain obscured. When datasets and knowledge structures fail to account for marginalised perspectives, the colonial and patriarchal biases embedded in original archives are not only encoded into digital systems—they can become further entrenched. Awareness of these multiple types of bias and their compounding effects is essential for developing better digital humanities methodologies.

1.2. Need for a Transdisciplinary Understanding of Bias

Despite growing attention to bias and its mitigation [12, 13, 14, 15], there is still no coherent framework for understanding what bias actually *is*. The term ‘bias’ carries different meanings across contexts:




Figure 1: Silences in the historic index; a: 19th century index; b: Testator: only his name is indexed; c: “Free Christian woman Magdalena van Boegis” is present in the document but not findable in the digitised index (NA, VOC, 6847 folio number 40, page 119)

¹“1.04.02 Inventaris van Het Archief van de Verenigde Oost-Indische Compagnie (VOC), 1602–1795 (1811) | Nationaal Archief,” <https://www.nationaalarchief.nl/onderzoeken/archief/1.04.02>

²https://www.nationaalarchief.nl/onderzoeken/archief/1.04.02/invnr/6847/file/NL-HaNA_1.04.02_6847_0119

archivists emphasise issues of inventorisation and deceptive categorisation; historians examine historical power structures [2, 16]; digital humanists have focused on unfair representation [17]; semantic web researchers grapple with how ontologies reflect—and reproduce—dominant epistemologies [18], and machine learning focuses on ground truth bias or representation in training data [19, 20, 21]. Even within specific academic fields, the concept of bias proves elusive. Blodgett et al.'s [22] analysis of 146 papers in the field of natural language processing revealed significant confusion in defining 'bias', while in digital cultural heritage, the characterisation of offensive terminology as bias remains unclear.³ Yet despite growing awareness of bias in the digital humanities, many researchers and institutions find themselves paralysed by the concept's complexity and apparent ubiquity. Without a coherent framework for understanding and addressing bias, there is a risk of either oversimplifying it or becoming overwhelmed by it - what we term "bias paralysis".

1.3. Bias Mitigation Approaches and their Limitations

Several valuable interventions have emerged, from documentation templates to tools for identifying harmful language and replacing offensive terminology [25, 14, 26, 27, 28, 29, 23]. However, these approaches typically address specific manifestations of bias. Digital humanities researchers regularly encounter an entangled spectrum: historical, technical, descriptive, and representational biases—many of which are reinforced through infrastructures like metadata schemas, machine learning pipelines, and semantic web technologies. While these semantic systems can encode dominant world-views, recent work also shows their potential for mitigating bias—by enabling more nuanced representations, identifying disparities across groups, and supporting fairer information retrieval and classification [8, 18].

What remains missing is a cohesive approach that makes visible how these forms of bias interact and transform across the data lifecycle. Without such a framework, mitigation efforts risk remaining fragmented—treating symptoms rather than confronting the underlying systems through which bias is produced, sustained, and reproduced in digital knowledge production.

1.4. From Paralysis to Practice: The Bias-Aware Framework

To effectively identify, articulate, and mitigate bias in digital humanities research, three fundamental questions need to be answered:

1. What exactly do we mean by "bias" in digital humanities research?
2. Where does bias occur in the dataset creation process?
3. How can researchers effectively address bias within resource constraints?

In response to these questions, we are developing a "Bias-Aware Framework" for dataset creation:

1. **A Bias Thesaurus:** A comprehensive list of the concepts connected to bias (such as representation, offensive language, FAIR, CARE, silences, etc.) that creates a shared vocabulary for discussing bias across disciplines.
2. **A Bias-Aware Data Lifecycle Model:** A visual and conceptual model mapping where and how different types of bias arise across the research process, enabling targeted reflection and intervention.
3. **Practical Guidelines:** A set of reflective questions, examples, and "good-better-best" recommendations tailored to each stage of the data lifecycle, supporting practices of bias identification, description, and redress.

This framework addresses a recognised gap in digital humanities: "a set of guidelines is missing, a serious lack when one might want to think through ethical concerns" [30]. It is a framework that demystifies 'bias' and transforms it into a productive tool for improving knowledge production.

³For instance the Words Matter [23], a publication on sensitive words in the museum sector, doesn't use the term 'bias', but projects such as DE-BIAS [24], based at the Dutch Institute of Sound and Vision, use the term in context of developing an automated tool to identify harmful language in archives.

2. Methodology

Focus on Datasets as Critical Intersection Points. Our framework centres datasets as the primary unit of analysis within the digital humanities landscape. This focus is strategic for several reasons. Datasets function as critical nexus points where four key elements converge: the **data** itself (from archives, born-digital sources, or interviews); the **researchers** who structure this data; the **users** who access and build upon these resources; and the **computational methods** that process this information. For Semantic Web Technologies in particular, datasets form the foundation upon which ontologies and knowledge graphs are constructed, making them crucial sites for bias intervention before problematic representations become encoded in semantic structures.

Most importantly, datasets should not be viewed merely as areas of ‘risk’ requiring intervention, but as sites of tremendous opportunity when created with critical awareness of biases. Our project highlights how dataset creation can function as a form of “deconstructing archival sources”, enabling researchers to view historical materials through new analytical lenses [31]. For example, it was through dataset creation and subsequent analysis that researchers uncovered the pivotal role of a 17-year-old woman named Flora in orchestrating the escape of nineteen enslaved people—effectively re-inscribing her into historical narratives despite her name appearing only fleetingly in primary sources [32].

The Bias-Aware framework development follows a three-phase approach combining theoretical analysis, practitioner insights, and practical validation:

2.1. Literature Review

To gain a better overview of current theories about and strategies against bias, we systematically reviewed literature in the fields of archival studies [23, 26, 33, 34, 32], epistemology [35, 36] and computer sciences [37, 38, 39, 40, 18]. These are fields in which bias has received much attention. From here, we expanded out to include sources that were not academic and/or text-based, such as videos, art installations, and fiction [41, 42, 43]. The importance of including these resources was to critically confront our own biases for the written and academic. For each resource, we focused on what forms of bias or strategies to mitigate bias were present, and made a note of that under the column ‘concepts used’. We compiled these in an open-to-access list of resources.⁴

2.2. Insights from Partner Projects

Our framework development draws on semi-structured interviews with partners from four major digital infrastructure projects focusing on colonial and slavery archives: Slave Voyages⁵, GLOBALISE⁶, Exploring Slave Trade in Asia⁷, and the Historical Database of Suriname and Caribbean⁸. These partnerships provide crucial insights into practical implementation challenges. We also engaged advisors with diverse expertise across cultural heritage, critical archival studies, community memory work, ethnomusicology, natural language processing, and FAIR data principles. This plurivocal approach [44] ensures our framework offers adaptable strategies and examples suitable for diverse project contexts and resource levels.

2.3. Framework Validation and Refinement

We are validating and refining the framework through two parallel tracks: expert consultations and interactive workshops with digital humanities and social science projects. The workshops serve as practical testing grounds where participants apply the framework and its methodology to analyse bias in

⁴Combating Bias Resources List. We are now working on the bias thesaurus to establish better categories for organising the readings before making the resource list fully collaborative.

⁵<https://www.slavevoyages.org/>

⁶<https://globalise.huygens.knaw.nl/>

⁷<https://esta.iisg.nl/>

⁸<https://www.ru.nl/onderzoek/onderzoeksprojecten/historische-database-van-suriname-en-de-cariben>

their own datasets. This implementation phase aims to reveal the framework’s strengths and limitations across different domains and identify potential blind spots. Participant feedback and documented use cases will drive iterative improvements, ensuring its broader applicability and effectiveness.

3. The Bias-Aware Framework

3.1. From Bias Paralysis to Bias as a Category of Analysis

To conceptualise bias meaningfully, we first examined its etymology. “Bias” entered English in the 1570s from the game of *boules*, describing balls weighted to curve obliquely. This technical meaning evolved into the figurative sense of “a one-sided tendency of the mind” and later “undue propensity or prejudice.” The French origin *biais* means “sideways, askance, against the grain”—suggesting movement contrary to an expected direction. This etymology raises a crucial question: when labelling something “biased,” what is our assumed “true” path? What constitutes an unbiased space, description, or archive—and is such a thing even possible? Rather than pursuing an impossible “bias-free” ideal, we draw inspiration from the textile meaning of bias: the diagonal stretch between warp and weft where fabric shows greatest flexibility. Garments cut “on the bias” follow this diagonal orientation, creating fluidity and adaptability.

We employ this sense of bias in our framework: just as fabric’s bias exposes structural tensions and possibilities, biases in datasets highlight gaps, conditions of production, overlooked questions, and unconsidered perspectives. This shifts our focus from attempting to “solve” bias to using it as a critical tool for systematic analysis.

3.2. Bias-Aware Framework

The Bias-Aware Framework consists of three integrated components designed to transform how researchers understand, identify, and address bias throughout the dataset creation process. Each component builds upon the others to create a comprehensive approach to bias as a category of analysis.

3.2.1. Bias Thesaurus

Our interviews with dataset creators reveal that bias functions as a heuristic addressing interconnected concerns about power, inequality, positionality, silences, and representation. Drawing from Scott’s [36, 45] concept of gender as an analytical category and Foucault’s [46] understanding of power as relational, we view bias as dynamic—actively shaping and being shaped by social and historical contexts. The bias Thesaurus maps the various *expressions of bias*—concrete forms bias takes in research practices, such as harmful language, uneven descriptive depth, or limiting categorisation schemes. The thesaurus creates a shared vocabulary across disciplines, visualises interconnections between different expressions of bias, and provides researchers with a conceptual map for navigating bias-related concerns.




Figure 2: Bias-Aware Dataset Lifecycle

3.2.2. Bias-Aware Dataset Lifecycle Model

The dataset creation lifecycle forms the structural backbone of this framework, grounding abstract bias considerations in familiar research workflows while addressing a gap in digital humanities methodology. Our model builds upon the Research Data Alliance’s harmonised Research Data Lifecycle (RDL) model⁹

⁹https://www.rd-alliance.org/wp-content/uploads/2024/09/D1_The-creation-of-a-harmonised-research-data-lifecycle-RDL-model-and-crosswalk.pdf

which identifies five key stages (Set Up, Collect, Process, Analyse, Preserve & Share). We extend this framework by mapping how different expressions of bias defined in the Thesaurus manifest at each stage.

A key insight from our research is the “stickiness” of certain bias expressions across multiple stages, though they manifest differently depending on the stage’s focus. For example, representation concerns appear throughout the lifecycle: in Set Up, they relate to whose scholarship informs the project; in Collect, they concern whose perspectives are captured in the data; in Process, they involve how categories represent complex realities; and in Analysis, they address whose stories are highlighted in figure 2.

3.2.3. Practical Guidelines

The final component transforms theoretical understanding into practical action through structured guidelines for each stage of the dataset lifecycle. These guidelines provide reflective questions, curated resources, documentation templates, “good-better-best” recommendations that accommodate varying resource constraints, and example strategies drawn from successful digital humanities projects.

Figure 3 illustrates our guideline approach for considering CARE as a principle at the funding step, while figure 4 illustrates our guideline approach for addressing archival silences, offering tiered intervention strategies from basic documentation to participatory community engagement.

The guidelines emphasise that addressing bias is not an all-or-nothing proposition—even resource-constrained projects can implement basic bias-aware practices. This scaffolded approach helps prevent “bias paralysis” by making intervention accessible regardless of project scale or resources.

4. Conclusion and Future Work

4.1. Conclusions

Our approach reconceptualises bias not as noise to be eliminated but as parameters that reveal how knowledge is constructed. By making bias visible and analysable, we transform it from a technical problem into a productive analytical lens that enhances both the integrity and reuse potential of digital humanities research. The Bias-Aware Framework provides a structured vocabulary, makes visible critical intervention points in the research process, and offers actionable strategies adaptable to various resource constraints. Rather than eliminating bias—an impossible task—we aim to establish bias awareness as a fundamental aspect of scholarship comparable to citation practices or methodological transparency.

4.2. Future Work

Our work opens several promising directions for future research:




Figure 3: Bias-Aware Dataset Lifecycle with reflective questions

2. Archival Silences		draft example guidelines
<p>Scenario: A dataset omits or underrepresents certain voices, groups, or perspectives due to gaps in (historical) records or structural biases in data collection processes.</p> <p>Reflective questions</p> <ul style="list-style-type: none"> What groups, voices, or perspectives are missing from your sources? How do these silences impact potential research questions and findings? What alternative sources might help address these gaps? How can silences be made visible to users? <p>Example dataset: GLOBALESE Namebooks</p>		
<p>Good practice</p> <ul style="list-style-type: none"> Recognize archival silences in source utilized and notify users about them in the data-envelope and documentation. Reflect on the impact that the gender politics of the source has on research and down-stream machine learning tasks. Include this information in the data-envelope and documentation accompanying the dataset. 	<p>Better practice</p> <p>In addition to 'good practices':</p> <ul style="list-style-type: none"> Scout for and incorporate alternative sources to ameliorate gender imbalances of the source. 	<p>Best practice</p> <p>In addition to 'better practices':</p> <ul style="list-style-type: none"> Develop a plan for how the dataset creators conceptualise gender, incorporate this in dataset design. Consider how the gender plan can be applied to the source material used. Document all decisions in documentation and data-envelopes. Maintain versioned documentation with clear update plans. Create data visualisations, both as a way to identify gaps in the sources, and to make the gaps easily readable to users.

Figure 4: Example guidelines to deal with archival silences

1. **Disciplinary Expansion and Empirical Validation:** We aim to extend the framework beyond colonial archives to other humanities domains while validating its effectiveness through diverse case studies. This parallel expansion and validation will test the framework's flexibility, identify domain-specific adaptations, and document how bias awareness transforms research outcomes across project types.
2. **Formalising Vocabulary and Knowledge Structures:** A critical next step involves formalising the Bias Thesaurus, which will capture the relationships between different bias expressions, their manifestations across the data lifecycle, and appropriate mitigation strategies.
3. **Theoretical Foundations:** Further research will justify and reflect on digital humanities' unique position at the intersection of computational methods and humanistic inquiry. This work will contribute to ongoing debates about how computational approaches can be informed by critical humanities perspectives, particularly regarding knowledge representation and classification systems.
4. **Sustainable Infrastructure** Long-term maintenance of the framework requires developing sustainable infrastructure through community governance and versioning systems. We envision creating a collaborative platform where researchers can contribute examples, adaptations, and extensions to the framework, ensuring it evolves alongside changing research practices and emerging technologies.
5. **Implementation Formats:** To maximise accessibility and adoption, we will explore various presentation formats from open-access platforms to downloadable templates, interactive tools, and integration with existing data management frameworks.

Acknowledgments

The authors are the main researchers of the Combatting Bias project¹⁰, which is a collaborative initiative based at the Huygens Institute and International Institute of Social History in Amsterdam, Netherlands, funded by the NWO via the Thematic Digital Competence Centre Social Sciences and Humanities¹¹. We thank the reviewers for their critical feedback.

Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

References

- [1] C. B. McCullagh, Bias in historical description, interpretation, and explanation 39 (2000) 39–66. URL: <https://www.jstor.org/stable/2677997>, publisher: [Wesleyan University, Wiley].
- [2] M.-R. Trouillot, *Silencing the Past: Power and the Production of History*, Beacon Press, 1995.
- [3] A. Prescott, Bias in big data, machine learning and AI: What lessons for the digital humanities? 17 (2023). URL: <https://www.digitalhumanities.org/dhq/vol/17/2/000689/000689.html>.
- [4] N. B. Thylstrup, *The politics of mass digitization*, MIT Press, 2019.
- [5] R. Navigli, S. Conia, B. Ross, Biases in large language models: Origins, inventory, and discussion 15 (2023) 10:1–10:21. doi:10.1145/3597307.
- [6] C. Jeurgens, M. Karabinos, Paradoxes of curating colonial memory 20 (2020) 199–220. doi:10.1007/s10502-020-09334-z.
- [7] M. Luthra, C. Jeurgens, Humanising digital archival practice. access to archives guided by social justice, in: M. Ginés-Blasi (Ed.), *Intentional Invisibilization in Modern Asian History: Concealing and Self-Concealed Agents*, De Gruyter, 2025, pp. 161–196. doi:10.1515/9783111381831-008.

¹⁰<https://combattingbias.huygens.knaw.nl/>

¹¹<https://tdcc.nl/about-tdcc/ssh/>

- [8] M. Luthra, K. Todorov, C. Jeurgens, G. Colavizza, Unsilencing colonial archives via automated entity recognition (2023). doi:10.1108/JD-02-2022-0038.
- [9] M. Ehrmann, A. Hamdi, E. L. Pontes, M. Romanello, A. Doucet, Named entity recognition and classification in historical documents: A survey (2023). doi:10.1145/3604931.
- [10] E. F. T. K. Sang, F. De Meulder, Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition (2003) 142–147. URL: <https://aclanthology.org/W03-0419>.
- [11] G. Doddington, A. Mitchell, M. Przybocki, L. Ramshaw, S. Strassel, R. Weischedel, The automatic content extraction (ACE) program - tasks, data, and evaluation, in: M. Teresa Lino, M. Francisca Xavier, F. Ferreira, R. Costa, R. Silva (Eds.), Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC'04), European Language Resources Association (ELRA), 2004.
- [12] C. H. Chu, S. Donato-Woodger, S. S. Khan, R. Nyrup, K. Leslie, A. Lyn, T. Shi, A. Bianchi, S. A. Rahimi, A. Grenier, Age-related bias and artificial intelligence: a scoping review 10 (2023) 1–17. doi:10.1057/s41599-023-01999-y, publisher: Palgrave.
- [13] A. Ortolja-Baird, J. Nyhan, Encoding the haunting of an object catalogue: on the potential of digital technologies to perpetuate or subvert the silence and bias of the early-modern archive 37 (2022) 844–867. doi:10.1093/llc/fqab065.
- [14] E. M. Bender, B. Friedman, Data statements for natural language processing: Toward mitigating system bias and enabling better science, Transactions of the Association for Computational Linguistics 6 (2018) 587–604. doi:10.1162/tacl_a_00041.
- [15] M. K. Scheuerman, A. Hanna, E. Denton, Do datasets have politics? disciplinary values in computer vision dataset development 5 (2021) 317:1–317:37. doi:10.1145/3476058.
- [16] M. van Rossum, Labouring transformations of amphibious monsters: Exploring early modern globalization, diversity, and shifting clusters of labour relations in the context of the dutch east india company (1600–1800) 64 (2019) 19–42. doi:10.1017/S0020859019000014.
- [17] I. Kizhner, M. Terras, M. Rumyantsev, V. Khokhlova, E. Demeshkova, I. Rudov, J. Afanasieva, Digital cultural colonialism: measuring bias in aggregated digitized content held in google arts and culture 36 (2021) 607–640. doi:10.1093/llc/fqaa055.
- [18] P. R. Lobo, E. Daga, H. Alani, M. Fernandez, Semantic web technologies and bias in artificial intelligence: A systematic literature review, Semantic Web 14 (2023) 745–770. doi:10.3233/SW-223041.
- [19] J. Buolamwini, T. Gebru, Gender shades: Intersectional accuracy disparities in commercial gender classification, in: Proceedings of the 1st Conference on Fairness, Accountability and Transparency, PMLR, 2018, pp. 77–91. URL: <https://proceedings.mlr.press/v81/buolamwini18a.html>.
- [20] H. Suresh, J. lobo, A framework for understanding sources of harm throughout the machine learning life cycle, in: Equity and Access in Algorithms, Mechanisms, and Optimization, ACM, 2021, pp. 1–9. doi:10.1145/3465416.3483305.
- [21] A. Søgaard, B. Plank, D. Hovy, Selection bias, label bias, and bias in ground truth, in: Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Tutorial Abstracts, 2014, pp. 11–13. URL: <https://aclanthology.org/C14-3005.pdf>.
- [22] S. L. Blodgett, S. Barocas, H. Daumé III, H. Wallach, Language (technology) is power: A critical survey of “bias” in NLP, in: D. Jurafsky, J. Chai, N. Schluter, J. Tetreault (Eds.), Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, 2020-07, pp. 5454–5476. URL: <https://aclanthology.org/2020.acl-main.485>. doi:10.18653/v1/2020.acl-main.485.
- [23] W. Modest, R. Lelijveld, Words matter: an unfinished guide to word choices in the cultural sector, 2018. URL: <https://amsterdam.wereldmuseum.nl/en/about-wereldmuseum-amsterdam/research/words-matter-publication>.
- [24] DE-BIAS - detecting and cur(at)ing harmful language in cultural heritage collections, 2024. URL: <https://pro.europeana.eu/project/de-bias>.
- [25] H. Alkemade, S. Claeysens, G. Colavizza, N. Freire, J. Lehmann, C. Neudecker, G. Osti, D. v. Strien, Datasheets for digital cultural heritage datasets 9 (2023) 17.

- [26] A. Chilcott, Towards protocols for describing racially offensive language in UK public archives, in: V. Frings-Hessami, F. Foscari (Eds.), *Archives in a Changing Climate - Part I & Part II*, Springer Nature Switzerland, 2022, pp. 151–168. doi:10.1007/978-3-031-19289-0_10.
- [27] M. Luthra, M. Eskevich, Data-envelopes for cultural heritage: Going beyond datasheets, in: I. Siegert, K. Choukri (Eds.), *Proceedings of the Workshop on Legal and Ethical Issues in Human Language Technologies @ LREC-COLING 2024, ELRA and ICCL*, 2024, pp. 52–65. URL: <https://aclanthology.org/2024.legal-1.9>.
- [28] A. Masschelein, F. Truyen, S. Taes, J. van Mulder, A. Stynen, R. Pireddu, Report on research into bias types and patterns, including a typology applied to europeana use cases and a vocabulary co-created with communities, 2023-12-31.
- [29] M. K. Scheuerman, K. Spiel, O. L. Haimson, F. Hamidi, S. M. Branham, HCI guidelines for gender equity and inclusivity, Maryland Shared Open Access Repository, 2020. doi:10.13016/M2NW1F-P0JX.
- [30] J. O'Sullivan, *The bloomsbury handbook to the digital humanities*, Bloomsbury Publishing, 2024.
- [31] N. L. Peluso, Whose woods are these? counter-mapping forest territories in kalimantan, indonesia, *Antipode* 27 (1995) 383–406.
- [32] C. W. van Galen, B. Quanjer, The wolf, the island and the sea: truancy and escaping slavery in curacao (1837–1863) 29 (2024) 262–279. doi:10.1080/1081602X.2024.2340542, publisher: Routledge.
- [33] A. L. Stoler, *Along the Archival Grain: Epistemic Anxieties and Colonial Common Sense*, Princeton University Press, 2010.
- [34] G. C. Spivak, The rani of sirmur: An essay in reading the archives 24 (1985) 247–272. doi:10.2307/2505169, publisher: [Wesleyan University, Wiley].
- [35] M. Foucault, *Archaeology of Knowledge*, 2 ed., Routledge, 2002. doi:10.4324/9780203604168.
- [36] J. W. Scott, Gender: A useful category of historical analysis 91 (1986) 1053–1075. doi:10.2307/1864376, publisher: [Oxford University Press, American Historical Association].
- [37] T. Gebru, J. Morgenstern, B. Vecchione, J. W. Vaughan, H. Wallach, H. D. III, K. Crawford, Datasheets for datasets 64 (2021) 86–92. doi:10.1145/3458723.
- [38] E. S. Jo, T. Gebru, Lessons from archives: strategies for collecting sociocultural data in machine learning, in: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, FAT* '20*, Association for Computing Machinery, 2020, pp. 306–316. doi:10.1145/3351095.3372829.
- [39] W. Orr, K. Crawford, The social construction of datasets: On the practices, processes and challenges of dataset creation for machine learning, 2023. doi:10.31235/osf.io/8c9uh.
- [40] R. Brate, A. Nesterov, V. Vogelmann, J. van Ossenbruggen, L. Hollink, M. van Erp, Capturing contentiousness: Constructing the contentious terms in context corpus, in: *Proceedings of the 11th Knowledge Capture Conference, K-CAP '21*, Association for Computing Machinery, 2021, pp. 17–24. doi:10.1145/3460210.3493553.
- [41] Chimamanda Ngozi Adichie, The danger of a single story, 2009. URL: https://www.ted.com/talks/chimamanda_ngozi_adichie_the_danger_of_a_single_story.
- [42] C. Kring, K. KU Leuven, DE-BIAS, Face/surface. metamorphosis of colonial perspectives, 2024. URL: https://kadoc.kuleuven.be/3_onderzoek/33_onzeonderzoeksoutput/tentoonstellingen/2024/tt_2024_bias.
- [43] L. V. Belle, In the place of shadows, 2022. URL: <https://www.lavaughnbelle.com/home-1#/in-the-place-of-shadows/>.
- [44] E. Sitzia, Multiple narratives and polyvocality as strategies of inclusive public participation: Challenges and disruption in the history museum 10 (2023) 51–63. doi:10.7202/1108037ar.
- [45] J. W. Scott, Gender: Still a useful category of analysis? 57 (2010) 7–14. doi:10.1177/0392192110369316.
- [46] M. Foucault, R. Hurley, *The history of sexuality. Volume 1, The will to knowledge*, Popular Penguins, Penguin, 2008.