

Understanding “Democratization” in NLP and ML Research

Arjun Subramonian*¹ Vagrant Gautam*²
Dietrich Klakow² Zeerak Talat³

¹University of California, Los Angeles, USA ²Saarland University, Germany

³Mohamed Bin Zayed University of Artificial Intelligence, UAE

arjunsu@cs.ucla.edu

{vgautam, dietrich.klakow}@lsv.uni-saarland.de

z@zeerak.org

Abstract

Recent improvements in natural language processing (NLP) and machine learning (ML) and increased mainstream adoption have led to researchers frequently discussing the “democratization” of artificial intelligence. In this paper, we seek to clarify how democratization is understood in NLP and ML publications, through large-scale mixed-methods analyses of papers using the keyword “democra*” published in NLP and adjacent venues. We find that democratization is most frequently used to convey (ease of) access to or use of technologies, without meaningfully engaging with theories of democratization, while research using other invocations of “democra*” tends to be grounded in theories of deliberation and debate. Based on our findings, we call for researchers to enrich their use of the term democratization with appropriate theory, towards democratic technologies beyond superficial access.¹

1 Introduction

As the influence of language technologies has grown, it has become increasingly popular to discuss “democratization” in natural language processing (NLP) and machine learning (ML) research (Seger et al., 2023); for instance, OpenAI has invested in a “democratic process for deciding what rules AI systems should follow” (Zaremba et al., 2023), Anthropic has explored how “democratic processes can influence artificial intelligence (AI) development” (Ganguli et al., 2023), and HuggingFace has stated their mission to be to “democratize good machine learning” (Simon, 2022). Indeed, a large number of NLP and ML papers mention terms related to democracy (see Figure 1), thereby raising the question: What do we understand by “democracy” and “democratization” when we invoke them in research?

*Equal contribution.

¹Our code is available at: <https://github.com/ArjunSubramonian/democratization-nlp>.




Figure 1: Frequency of mentions of democracy (>0) per paper in all work published in the ACL Anthology, ICLR, ICML, or NeurIPS before November 24, 2023. 76.1% of papers only mention democracy once.

So far, the treatment of democracy in NLP and ML literature, and particularly the term “democratization,” has not been subject to careful investigation. Our paper fills this gap by analyzing uses of “democratization” in NLP and ML papers, and their connections to democracy. We examine conceptualizations of these terms through a large-scale mixed-methods analysis of every use of “democra*” in papers published in the Anthology of the Association of Computational Linguistics (ACL Anthology), the International Conference on Learning Representations (ICLR), the International Conference on Machine Learning (ICML), and Neural Information Processing Systems (NeurIPS).

We find that on one hand, the use of “democratization” tends to indicate a broadening of access to research artifacts, particularly without domain expertise, while NLP and ML literature discussing democracy in other contexts is often rooted in theories of deliberation and debate. We also find that while authors associate democratization with positive values related to access and reducing costs, the term itself is rarely defined or operationalized.

Prior work has argued that the “democratization of AI revolves primarily around the notion of access” (Burkhardt, 2019; Sudmann, 2019; Sudmann and Waibel, 2019; Luchs, 2023); our work provides systematic evidence for this claim, and is grounded in a comparison to other democracy-related terms.

Next, we examine papers that mention democracy for their depth of engagement with the topic, by exploring their text and citations. We find that a majority of papers only invoke democracy once, do so outside of methods and results sections, and engage minimally with extra-disciplinary work.

We conclude that “democratization” constitutes a misnomer for “access,” and therefore encourage future work to either enrich their research by drawing on over 3000 years of scholarship on democracy and democratization, or use “access” instead. Lacking clear, consistent and responsible use of the term “democratization,” NLP and ML risk misrepresenting progress in capturing democratic values, the distribution of power, and public control of AI. Clearer conceptualizations of “democratization” can thus strengthen progress towards truly democratic technologies beyond just superficial access.

2 Related Work

2.1 Democratization beyond AI

The use of “democratization” extends far beyond AI. In conservation biology, for instance, it is discussed in the context of citizen and community science; public participation in processes such as data collection is seen as democratizing knowledge production (Bela et al., 2016), and reducing gaps between academia and wider society (Sauermaun et al., 2020). However, constraining community science to participation has also been criticized as “participation washing” (Sloane et al., 2022), as it often disregards local knowledge, prevents the public from formulating scientific questions, and fails to change the norms of institutions (Kimura and Kinchy, 2016). In contrast, political scientists examine the democratization of policy research through “collaborative citizen-expert inquiry” (Fischer, 1993) which has been considered essential to democratically tackling social issues (Weinberg, 2022). Internet scholars investigate the democratizing effects of online information and social media, i.e., how they have helped to spread pro-democratic ideas, discussions, and protests globally (Hill and Hughes, 1999; Weinstein, 2012). Beyond research, there have been calls towards protecting the in-

tegrity of democracy through the democratization of media and “free access to pluralistic information and opinion” (de Zayas, 2017). In relation to emerging democracies, the democratization of media is often linked to the diversification of news sources (Barnett, 1999; Tettey, 2001; Porto, 2012).

2.2 Conceptions of democratization in AI

Research in AI has presented access-centric conceptions of democratization, e.g., to identify criteria for democratizing the use of AI, such as affordability, accessibility, and fairness (Ahmed et al., 2020). Similarly, Ahmed and Wahed (2020) conceptualize democratization as equity in access to compute between tech companies and non-elite universities. However, this line of research has not examined the possible connections between democratization and democracy. Prior work has also challenged the conceptualization of democratization in AI. Seger et al. (2023) argue that disparate uses of the term “democratization” have caused a lack of recognition of shared “goals, methodologies, risks, and benefits.” Drawing from news articles and talks, they identify four notions of democratization: use, development, benefits, and governance. Similarly, in a study of 35 articles on the use of “democratization” and its connection to democracy within the scope of medical AI, Rubeis et al. (2022) uncover diverse conceptualizations, from increasing data access to AI governance.

Another line of work, focusing on AI governance and increased public control of AI development and deployment, argues that public participation is critical for democratizing AI, e.g., Gilman (2023) calls for institutions to budget for participation at all stages of AI development. Participation has also been operationalized by aligning models to a “constitution” based on the values of human representatives (Siddarth, 2023); by connecting open-source and democratic communities, and widening geographic diversity in public input processes (Collective Intelligence Project, 2024); and by leveraging “democratic” frameworks to gather AI uses, harms, and benefits from the public to guide the evaluation and regulation of AI (Mun et al., 2024). However, these approaches offer minimal opportunities for publics to contest the logics and power structures of the AI industry (Luchs, 2023).

In contrast to these bodies of work, we perform a large-scale mixed-methods analysis of papers published at NLP and ML venues. Similarly to Seger

et al. (2023) and Rubeis et al. (2022), we find distinct conceptualizations of democratization that obviate its benefits and risks, often due to a lack of theoretical engagement. Ultimately, our analysis shows that the dominant conception of democratization is access, and that a shared understanding of democratization and democracy, which is essential for democratic frameworks, remains absent within the NLP and ML community at large.

3 Data

Using the Semantic Scholar API (Kinney et al., 2023), we collect all papers published before November 24, 2023 in the ACL Anthology, ICML, ICLR, and NeurIPS, that mention terms related to “democracy.” We choose these venues, as they are top-tier NLP and ML conferences that influence practices in the field. We obtain 1,537 papers, which we filter for relevance, obtaining a final dataset of 506 papers and 916 excerpts for analysis.

Obtaining Excerpts We first collect all metadata and text from open-access PDFs using the Semantic Scholar API. We split the text of each paper using the punkt NLTK sentence tokenizer (Bird and Loper, 2004), and extract all sentences that contain the substring “democra” (excluding “democrats”), resulting in 4,203 excerpts across 1,709 papers. We do not include related terms (e.g., participatory governance, constitution, etc.) so that we do not inadvertently select irrelevant papers, and to keep our discussion firmly grounded in a comparison between democratization and democracy.

Filtering Irrelevant Excerpts In order to identify excerpts that reveal how authors conceptualize “democratization” and “democracy,” we remove unrelated uses of “democra,” such as those in named entities (e.g., “Center for Media and Democracy”), motivating examples (e.g., for textual entailment), modeling examples (e.g., LDA topics), examples from datasets (e.g., tweets), mentions in non-English languages, and references. We perform this filtering using a two-stage approach: automatic filtering and manual annotation for relevance.

We curate a list of terms (see Appendix A) for automatically filtering excerpts: We exclude named entities (e.g., “the Syrian Democratic Forces”) and terms that exclusively appear as examples of data (e.g., tweets containing “#democracy”). One author verified all automatically filtered excerpts.

After filtering, we manually annotate the re-

maining 2,273 excerpts, searching for instances where the authors deliberately use words containing “democra” as part of their argument or evidence, examining the full PDF in ambiguous cases. After concluding the two-stage filtering process, we obtain 916 excerpts from 506 papers for analysis.

4 Conceptualizations of Democracy

To understand how democracy and democratization are conceptualized by authors in NLP and ML papers, we inductively analyze our data for overarching themes, values, and concepts. We find that conceptualizations of democratization are distinct from democracy, and instead are closely related to access and financial costs.

4.1 Methodology

Two authors annotate the first 300 excerpts independently for themes, concepts, and values in an open-ended manner (see Table 1 for example excerpts and annotations). We then resolve inconsistencies and consolidate themes, concepts, and values, before annotating the remaining excerpts independently. Finally, we group the themes, concepts, and values, respectively, into sets per paper.

Themes We qualitatively code the excerpts to identify salient, overarching themes that characterize how they discuss democracy; this is a common inductive methodology from the social sciences described by Saldana (2021). Four major categories emerge after a first pass over all the excerpts:

- *Necessary/Beneficial*: things that are necessary for or beneficial to democracy (e.g., discourse, majority, voting)
- *Danger*: dangers to democracy (e.g., misinformation)
- *Democratization*: use of the words “democratize” or “democratization” (e.g., of ML)
- *Math*: mathematical or ML ways to operationalize democracy (e.g., democratic matrices, mathematical models of democracy)

Two authors then systematically annotate every excerpt with an explicit and, if applicable, an implicit theme. An explicit theme is assigned to excerpts that explicitly state, e.g., that something is necessary for or a danger to democracy; otherwise, it is classified as *other*. In contrast, the implicit theme requires annotators to make inferences about how researchers think about democracy.

Excerpt	Themes	Concepts
“The right to access judicial information is a fundamental component of Canadian democracy and its judicial process.”	necessary / beneficial	access, information
“An abundance of incorrect information can plant wrong beliefs in individual citizens and lead to a misinformed public, undermining the democratic process.”	danger	citizenship, misinformation
“This is a totally democratic method where each vote counts the same.”	math	equal contribution
“This helps to improve data literacy, democratizing accessibility to otherwise opaque public database systems.”	democratization	access, data

Table 1: Example excerpts each of the four themes, along with the associated concepts we annotate.

For example, the excerpt: “The most democratic option is to give each tagger one vote (Majority),” is assigned an explicit theme of *math*, as it discusses operationalizing NLP taggers in a “democratic” way. We also infer that the authors believe majority voting is necessary for democracy, hence *necessary/beneficial* is assigned as an implicit theme.

We do not differentiate between papers about the effect of democracy on technology (e.g., *danger*) and democratic principles in technology (e.g., *democratization*), as all papers that invoke democracy-related terms can engage with democratic theories, and both themes relate to participation. Not distinguishing between them and instead inductively looking for what patterns emerge allows us to identify how differently democratization and democracy may be conceptualized.

Values and Concepts The same two authors also label each excerpt for values (e.g., “consensus” and “equality”) and more broadly concepts (e.g., “misinformation” and “elections”) associated with democracy to explore conceptualizations of democracy more granularly. We focus on values (a subset of concepts) in our main analysis; see [Appendix B](#) for further discussion of values and concepts.

4.2 Results

Of the four themes, we find that *democratization* is by far the most frequent with 213 papers, followed by 67 for *necessary/beneficial*, 58 for *danger*, and 35 for *math*. In total, we identify 110 concepts (including 77 values) associated with democracy, with each paper containing an average of 1.16 themes and 1.036 concepts. For themes, annotation is highly consistent, with a Cohen’s kappa of 0.973 for explicit themes and 0.887 for implicit themes




Figure 2: Frequency of values, split by *democratization* papers and all other papers. Associations with *democratization* (top) are different from associations with all other mentions of democracy (bottom).

using the Jaccard distance metric (Cohen, 1960). Given the minimal disagreement between authors, we henceforth do not distinguish between explicit and implicit themes. For concepts, annotators have a Cohen’s kappa of 0.349. Although this score only indicates fair agreement in the binary classifi-

cation setting (McHugh, 2012), with 110 possible concepts, there is a much lower random chance of agreement, and thus 0.349 reflects moderate to high agreement in this context.

Values Associated with Democracy in NLP and ML Figure 2 shows the values associated with “democratization” compared to all other mentions of democracy. We find that some values contradict each other. For instance, work has conceptualized “random selection,” “consensus,” and “majority (voting)” as democratic, however these are all mutually exclusive of one another. Yet, researchers conceptualize NLP and ML systems operating in these three manners as “democratic,” showing the need to explicitly consider how different systems require different conceptualizations of democracy.

We find that *non-democratization* papers identify values and concepts that readily connect to widespread theoretical notions of democracy, e.g., decision-making, deliberation, debate, and diversity. In contrast, *democratization* papers are overwhelmingly associated with increasing access, ease of use, and reducing costs and barriers. That is, democratization papers share values with radical egalitarian theories of democracy (see Section 7), but do not distinguish or make apparent the relationship between access and equal access to democratic processes. Thus, in contrast to other fields (see Section 2.1), NLP and ML researchers who use these words seem to conceive of *democratization* quite differently from democracy, associating them with different and sometimes conflicting values, and agreeing primarily that both are aspirational.

5 Democratization in NLP and ML

Given that *democratization* in NLP and ML is markedly different from the other themes, we examine the politics of democratization in papers with this theme. Specifically, we consider *what* is being democratized, *how*, and *to what end*?

5.1 Methods

To examine the politics of *democratization*, one author annotates all excerpts with an explicit theme of *democratization* for **targets of democratization**, i.e., what the object of democratization is; **causes of democratization**, i.e., how is an object being democratized, or what engenders its democratization; and the **goals of democratization**, i.e., why or to what ends an object is being democratized. See example excerpts and annotations in Table 2.

5.2 Results

We find that 59% of the papers do not state causes of democratization and 75% do not state the goals. A subset of authors describe democratization as a separate autonomous process that is, at best, minimally affected by their contributions.

Other authors posit that their research democratizes a technology, but do not elaborate on how that occurs, e.g., in terms of digital infrastructure, governance structures, participatory methods, etc. When stated, popular causes (see Table 3) for democratization are reductions in required compute, time, and cost. Targets for democratization are more nebulous; for instance, authors indicate NLP, AI, or research and access are the target for democratization, however what it means for any of these to be democratized is unclear at such a level of abstraction. In contrast, the primary goals of democratization are increasing access and use, particularly without requiring expertise. However, without consideration of the causes and the targets of democratization, such goals appear inherently elusive.

We validate our excerpt-based results by sampling papers for close readings of the entire articles. We identify set of papers by using the Huggingface (Wolf et al., 2020) all-mpnet-base-v2 sentence transformer (Reimers and Gurevych, 2019) to embed all excerpts related to *democratization*. Then, we apply spectral clustering to the embeddings (Figure 5 in Appendix A) and select 3 clusters using the spectral gap heuristic. We select 5 papers from the cluster centers and 5 from the boundaries from each cluster, for a total of 30 papers. Our close reading of our sampled papers confirm our excerpt analysis: none of the selected papers consider *what* is being democratized, or plan for *how* to democratize. Indeed, very few even comment on democratization outside of the excerpts.

6 Engagement with Democratic Theories

Given such a lack of consideration within the *democratization* theme, we examine how NLP and ML papers engage with literature on democracy to understand its influence on the conceptualizations in Section 4. We argue that discussing democracy or democratization without connecting to established theories reflects subpar interdisciplinarity and citational praxis, and risks misrepresenting how grounded AI is in democratic values.

Excerpt	Cause	Target	Goal
“We aim at an ambitious goal of democratizing the cost of pretraining.”		cost	
“We narrow our purview to open source and accessible data collections, motivated by the goal of democratizing accessibility to research.”	data	access, research	access
“With everyone being able to create data for their model training, we can pave the way for the democratization of AI.”		AI	access, use without expertise

Table 2: Top causes, targets and goals of democratization in the 213 papers that mention it.

Causes	<i>None specified (59%)</i> , compute reduction, data, cost reduction, social media, time reduction, open source, internet, access, tools, research, model hubs, libraries
Targets	Research, access, NLP, AI, ML, content creation, DL, language models, MT, internet, information, RL, data
Goals	<i>None specified (75%)</i> , use without expertise, access, increased language use, social good, reduce barriers, multilingual, sociological phenomena, quality issues, broader audience, fake news, commodification

Table 3: Top causes, targets and goals of democratization in the 213 papers that mention it.

6.1 Methods

We measure the depth of engagement with democracy by counting where and how often “democra*” terms are mentioned in papers. We extract section names using the Semantic Scholar API and normalize them across papers, e.g., mapping “Related Works” to “Related Work.” For a complementary view of engagement that is not limited to words containing the substring “democra,” we also study the references these papers cite: the fields they belong to, the proportion of extra-disciplinary citations, and citational *intent*, i.e., whether the citation is used to provide background, inform the methodology of the paper, or is related to the results. This analysis allows us to evaluate engagement with theories of democracy. We obtain field, venue and intent metadata using the Semantic Scholar API; we classify references as *intra-disciplinary* if they are from Computer Science, Mathematics, or Lin-




Figure 3: Frequency of paper sections in which mentions of democracy occur.

guistics, and as *extra-disciplinary* otherwise. Finally, we confirm the results of our computational analyses with a close reading of 24 papers.

6.2 Results

Where and How Often is Democracy Invoked?

We find that the vast majority of papers that mention the “democra*” tokens only mention it once (see Figure 1), and most mentions occur in the abstract, introduction, and conclusion sections (see Figure 3). These results support our findings in Section 5.2 that democracy is under-discussed in NLP and ML literature. Additionally, we find via a close reading of the nine papers with seven or more mentions that a larger number of mentions does *not* necessarily signal higher engagement; for example, mathematical papers frequently refer to “democratic” mathematical objects without connecting them to democratic theories.

What Type of Papers are Cited and Why?

Our citation analysis reveals that the vast majority of citations are from computer science (48.4%), which




Figure 4: Proportion of fields of study of references cited by papers that mention democracy.

is cited three times more than the second most frequent field, linguistics (12.9%) (see Figure 4). Following these, mathematics (9.7%), political science (5.8%), and sociology (4.1%) are most prominently cited. Given the direct relationship between ML, computer science and mathematics, and NLP and linguistics, less than 29% of references are directed towards other disciplines. This number may not indicate low engagement in and of itself, as it represents an aggregated view of our entire corpus of papers. However, when we look at individual papers, we find that the majority of them cite zero or one extra-disciplinary works: 181 papers cite **zero** extra-disciplinary papers, and another 88 cite exactly one. Given the invocation of far-reaching social concepts such as democracy and democratization, such poor levels of engagement are particularly surprising. The remaining 220 papers constitute a long-tail that engages more extensively with literature outside of NLP and ML.

Analyzing the extra-disciplinary citations, we see that most citations are from the social sciences, especially political science. However, when considering citational intent, we find that most (82.3%) are cited as background, and 15.5% and 2.14% are in the context of methods and results, respectively. Thus, even when work on democracy and democratization consults extra-disciplinary research, it may be primarily used to frame work rather than engage with methods or analyses of results. Simply citing scholarship on democracy (e.g., “hit and run” citations; Gorelik (2019)) is not equivalent to meaningfully engaging with it. To evaluate our use of extra-disciplinary citational intent as a proxy for meaningful engagement with democratic theories, we closely read the 15 papers with background, methods, *and* results citations of extra-disciplinary work. We observe that only the papers citing political science and economics literature, in particular, for methods and results, exhibit deeper considera-

tion of theories of democracy and democratization. Among citations of political science and economics references, most are still background (84.4%), with 13.5% and 2.15% in the context of methods and results. We confirm that the nine papers with background, methods, *and* results citations of political science and economics literature indeed meaningfully engage with democracy and democratization.

6.3 Qualitative Examples

Below, we present some examples of papers with higher engagement with democratic theories:

Public Dialogue: Analysis of Tolerance in Online Discussions Mukherjee et al. (2013) discusses tolerance in online discussion, citing work on public spheres, deliberative democracy, tolerance, public dialogue, deliberation, disagreement, and consensus. The authors ground their methodology in this cited work and perform a “computational study of tolerance” in online discussions. They also interpret their results in the context of this literature, discussing consequences for deliberative discussion and society at large when sustained disagreement turns into intolerance.

A Mathematical Model For Optimal Decisions In A Representative Democracy Magdon-Ismail and Xia (2018) propose a mathematical model for decision-making under representative democracy. Their extra-disciplinary citations include social sciences and mathematical social sciences. Through discussion of the differences between direct democracy and representative democracy, the authors motivate a new mathematical model to study the quality-quantity tradeoff with different numbers of representatives, for different types of voting issues, and with different levels of public competence.

Asking Too Much? The Rhetorical Role of Questions in Political Discourse Zhang et al. (2017) identify a rich source of rhetorical information in questions in UK parliamentary debates, and thus analyze these as an example of the rhetorical aspects of question in political discourse. They cite extra-disciplinary literature that establishes the role of questions in democratic processes, and motivate their work as a quantitative examination of aspects that have mostly been qualitatively examined before. Their unsupervised approach discovers clusters of question types asked in parliamentary discussions. The authors present an analysis

of these clusters and the members of parliament posing these questions (in terms of their tenure and their affiliation to the governing party or the opposition), grounded within extra-disciplinary literature about UK politics and history.

Legal and Political Stance Detection of SCOTUS Language Bergam et al. (2022) studies the Supreme Court of the United States (SCOTUS) using text analysis and stance detection on publicly available documents. Grounding their motivation and analysis in literature about public opinion and democratic principles, the authors also compare their approach to existing metrics from the social sciences, and show how a result about case salience parallels existing findings in political science research. Finally, they note a trade-off common to the quantitative social sciences in their ethics statement, i.e., that quantitatively analyzing text at scale erases many aspects of its complexity, even as it helps to uncover patterns that cannot feasibly be uncovered by a single qualitative researcher.

7 Democratic Theories and NLP and ML

As democratization has had a long history of study starting from 1100 BCE in ancient Phoenicia (Jacobsen, 1943; Glassman, 2017; Graeber and Wenigrow, 2021), in this section, we consider select theories of democracy as a basis for how NLP and ML research has understood and operationalized democracy. We argue that these theories can provide foundations for more democratic NLP and ML technologies by making democratic discussions representative and efficient, diversifying forums for democratic dialogues, and dismantling barriers to participation in democratic processes.

Deliberative Democracies Deliberation and inclusion in the democratic process are often highlighted as goals for democratic societies (Roberts, 2004) and technologies (Gilman, 2023). Indeed, in our surveyed papers, democratic deliberation often appears as a goal (see Section 4). Deliberative democracy is a form of democracy that emphasizes processes where participants can debate a particular object (e.g., a policy or technology) on its merits and make collective decisions about its implementation (Goodin, 2000). Deliberative democratic theory thus provides an avenue for obtaining more legitimacy of decisions by engaging wider publics in conversation about the use and application of research artifacts (Rosenberg, 2007).

Democratic Spheres As diversity and equal representation are values often associated with democracy in NLP and ML, this raises the question of how we might achieve such goals. While deliberative democracy provides an avenue for engaging publics, creating a single democratic arena—or sphere, as argued for by Habermas (1991)—for a large and diverse group gives weight to the loudest voices and majoritarian perspectives. This risks relegating many communities to the margins, particularly when the publics are large. In contrast, a plurality of public spheres, which each represent smaller communities, can afford better representation of all communities (Fraser, 1990). In practice, if NLP and ML research is consulting a larger group, it can be useful to divide the group into smaller segments, for all voices to be heard.

Democracy and Power Mumford (1964) has argued that technology can either afford access, agency, and distribute power, i.e., be democratic, or consolidate power within a small set of actors, i.e., be authoritarian. Therefore, efforts towards operationalizing the democratization of NLP and ML need to understand and address barriers to public participation and uneven distributions of power. In relation to discriminatory ML, Kalluri (2020) and D’Ignazio and Klein (2020) have argued that searching for fair ML can serve as a distraction to considering how ML distributes power.

Radical Egalitarian Democracies One approach towards dismantling power differentials and barriers to participation is egalitarian democratic theory. Under this framework, all humans must have equal access to participate in democratic processes, and these processes should in turn institute programs that dismantle systems of oppression (Wright, 2010). However, the development, operation, and control of NLP and ML technologies are currently determined by the interests of privately held companies (Zaremba et al., 2023; Ganguli et al., 2023; Talat et al., 2022; Gray Widder et al., 2023), under processes that consolidate impact within a small segment of society. Addressing barriers to public participation and power differentials, as seen through egalitarian democratic theory, would require rethinking processes of public engagement in all stages of the development lifecycle.

8 Discussion, Conclusion, and Recommendations

Our thematic and large-scale mixed-methods analyses show that democracy is used in NLP and ML with infrequent operationalization of democratization, vastly different views of what democracy means, and low levels of interdisciplinary engagement. Overall, our results show that when invoking democracy, NLP and ML researchers only shallowly engage with the centuries of literature from philosophy and social science devoted to it. It is thus necessary that NLP and ML researchers describe what they mean by and how they intend to operationalize democratization, to avoid misrepresenting public control of AI and bolstering “utopian-idealistic” AI hype (Sudmann, 2019).

In particular, researchers should reflect on what values and concepts they associate with democratization, how their understanding of democratization may be contested, and how their usage of “democratization” may be overloaded or overhyped. We also echo Seger et al.’s (2023) call to simply use the word “access” rather than “normatively loaded language” like “democratization” when discussing access-related questions. Researchers should go beyond “‘access’ as the sole condition for participation” (Luchs, 2023) and discuss the processes for “democratic oversight and control” of their artifacts (Verdegem, 2022). To this end, they should explicate the causes, targets, methods, and goals of democratization, what it means for their research to be fully democratic, and which opportunities and limits to public participation and control emerge.

Moreover, when invoking democracy and related concepts, researchers should detail how their understanding is informed by underlying theory and ensure to draw from and cite relevant literature. Conversely, if it is not, they should explicitly indicate this in their work. In both cases, researchers should reflect on where their conceptualizations fail with respect to their research and goals, and which challenges remain unresolved by their work. For example, when invoking democratization, researchers should explicitly note what remains unresolved in their goal of democratized technologies. Some efforts, e.g., OpenAI’s call for democratic inputs to AI (Zaremba et al., 2023) and Anthropic AI’s *Collective Intelligence Project* (2024), appear to engage more deeply with definitions and implications of democratic AI, yet do not critically examine questions of power and control. Similarly,

Djeffal (2019) operationalizes AI democratization in line with democratic traditions, including “parliamentary processes to debate and regulate artificial intelligence.” However, on the whole, we must urgently “reflect on [our] engagement with other fields” (Wahle et al., 2023). While engagement with democratic theory is a necessary precondition for research towards democratizing NLP and ML technologies, it is also necessary to address the hegemonic praxis of NLP and ML, and how it begets or hinders democratic technologies.

Limitations

In our analysis, we may miss relevant NLP and ML literature that treats democratization or democracy due to our focus on the ACL Anthology, ICLR, ICML and NeurIPS. In our choice of these venues, we are not explicitly controlling for differences in prestige (e.g., workshop papers in the ACL anthology, c.f. main conference papers) or focus (most notably, NLP versus ML), an analysis of which we leave to future work. We further cannot account for the perspectives of NLP and ML researchers who have richer conceptualizations of democratization but are not writing about it. In addition, our filtering of excerpts based on keywords like “democra” may cause us to exclude important discussions of democracy-adjacent concepts that do not use the word. This may be worsened by parsing errors stemming from our methods and the Semantic Scholar API. The Semantic Scholar API can also fail to correctly predict scholarly metadata, including fields of study and intent, which may affect our results. Furthermore, our discussion of theories of democracy (see Section 7) is far from exhaustive, given the rich history of the subject.

Ethical Considerations

Our paper emphasizes careful consideration and usage of the term “democratization,” especially given its relation to democracy, and urges drawing from extra-disciplinary literature on democratic theories. This is important for accurately representing the distribution of power, public control, and progress in NLP and ML. In light of our findings, we stress that our analysis only captures a snapshot in time and that researchers’ perspectives on democratization and democracy can evolve; moreover, the text of papers may not wholly reflect the perspectives of their authors, given the diversity of opinions among authors and reviewing incentives.

Acknowledgements

We thank the anonymous reviewers for their insightful feedback. We also greatly appreciate Luca Soldaini, Lucy Li, Maria Antoniak, and Shaily Bhatt for their constructive comments on the presentation and organization of the paper. We further thank Shreya Chowdhary and Skyler Wang for preliminary discussions about this project, and Luca Soldaini for help with Semantic Scholar.

References

- Nuri Mahmoud Ahmed and Muntasir Wahed. 2020. The de-democratization of ai: Deep learning and the compute divide in artificial intelligence research. *ArXiv*, abs/2010.15581.
- Shakkeel Ahmed, Ravi Mula, and Soma S. Dhavala. 2020. A framework for democratizing ai. *ArXiv*, abs/2001.00818.
- Clive Barnett. 1999. **The limits of media democratization in south africa: politics, privatization and regulation.** *Media, Culture & Society*, 21:649–671.
- Györgyi Bela, Taru Peltola, Juliette Claire Young, Bálint Balázs, Isabelle Arpin, György Pataki, Jennifer Hauck, Eszter Kelemen, Leena Kopperoinen, Ann Van Herzele, Hans Keune, Susanne Hecker, Monika Suškevičs, Helen E. Roy, Pekka Itkonen, Mart Külvik, Miklós László, Corina Basnou, Joan Pino, and Aletta Bonn. 2016. **Learning and the transformative potential of citizen science.** *Conservation Biology*, 30.
- Noah Bergam, Emily Allaway, and Kathleen Mckeown. 2022. **Legal and political stance detection of SCOTUS language.** In *Proceedings of the Natural Legal Language Processing Workshop 2022*, pages 265–275, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Steven Bird and Edward Loper. 2004. **NLTK: The natural language toolkit.** In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pages 214–217, Barcelona, Spain. Association for Computational Linguistics.
- Marcus Burkhardt. 2019. *Mapping the Democratization of AI on GitHub: A First Approach*, page 209–222. transcript Verlag.
- Jacob Cohen. 1960. **A coefficient of agreement for nominal scales.** *Educational and Psychological Measurement*, 20:37–46.
- Collective Intelligence Project. 2024. **A roadmap to democratic ai.**
- Alfred de Zayas. 2017. UN expert calls for democratization of the media. <https://www.ohchr.org/en/press-releases/2017/09/un-expert-calls-democratization-media>.
- C. D’Ignazio and L.F. Klein. 2020. *Data Feminism*. Strong Ideas. MIT Press.
- Christian Djeflal. 2019. *AI, Democracy and the Law*, page 255–284. transcript Verlag.
- Frank Fischer. 1993. **Citizen participation and the democratization of policy expertise: From theoretical inquiry to practical cases.** *Policy Sciences*, 26:165–187.
- Nancy Fraser. 1990. **Rethinking the public sphere: A contribution to the critique of actually existing democracy.** *Social Text*, (25/26):56–80.
- Deep Ganguli, Saffron Huang, Liane Lovitt, Divya Siddarth, Thomas Liao, Amanda Askell, Yuntao Bai, Saurav Kadavath, Jackson Kernion, Cam McKinnon, Karina Nguyen, and Esin Durmus. 2023. **Collective constitutional ai: Aligning a language model with public input.**
- Michele Gilman. 2023. **Democratizing ai: Principles for meaningful public participation.**
- Ronald M Glassman. 2017. *The origins of democracy in tribes, city-states and nation-states*, 1 edition. Springer International Publishing, Basel, Switzerland.
- Robert E. Goodin. 2000. **Democratic deliberation within.** *Philosophy & Public Affairs*, 29(1):81–109.
- Boris Gorelik. 2019. **The problem with citation count as an impact metric.**
- David Graeber and David Wengrow. 2021. *The Dawn of Everything: A New History of Humanity*. Penguin UK.
- David Gray Widder, Sarah West, and Meredith Whitaker. 2023. **Open (for business): Big tech, concentrated power, and the political economy of open ai.** *SSRN Electronic Journal*.
- J. Habermas. 1991. *The Structural Transformation of the Public Sphere: An Inquiry into a Category of Bourgeois Society*. Studies in Contemporary German Social Thought. MIT Press.
- Kevin Anthony Hill and John E. Hughes. 1999. **Is the internet an instrument of global democratization.** *Democratization*, 6:99–127.
- Thorkild Jacobsen. 1943. Primitive democracy in ancient mesopotamia. *Journal of Near Eastern Studies*, 2(3):159–172.
- Pratyusha Kalluri. 2020. **Don’t ask if artificial intelligence is good or fair, ask how it shifts power.** *Nature*, 583(7815):169–169.
- Aya Hirata Kimura and Abby Kinchy. 2016. **Citizen science: Probing the virtues and contexts of participatory research.** *Engaging Science, Technology, and Society*, 2:331–361.

- Rodney Kinney, Chloe Anastasiades, Russell Authur, Iz Beltagy, Jonathan Bragg, Alexandra Buraczynski, Isabel Cachola, Stefan Candra, Yoganand Chandrasekhar, Arman Cohan, et al. 2023. The semantic scholar open data platform. *arXiv preprint arXiv:2301.10140*.
- Inga Luchs. 2023. [Ai for all?: Challenging the democratization of machine learning](#). *A Peer-Reviewed Journal About*, 12(1):135–147.
- Malik Magdon-Ismail and Lirong Xia. 2018. [A mathematical model for optimal decisions in a representative democracy](#). In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- M. L. McHugh. 2012. Interrater reliability: the kappa statistic. *Biochemia Medica*, 22:276–282.
- Arjun Mukherjee, Vivek Venkataraman, Bing Liu, and Sharon Meraz. 2013. [Public dialogue: Analysis of tolerance in online discussions](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1680–1690, Sofia, Bulgaria. Association for Computational Linguistics.
- Lewis Mumford. 1964. [Authoritarian and Democratic Technics](#). *Technology and Culture*, 5(1):1.
- Jimin Mun, Liwei Jiang, Jenny Liang, Inyoung Cheong, Nicole DeCario, Yejin Choi, Tadayoshi Kohno, and Maarten Sap. 2024. [Particip-ai: A democratic surveying framework for anticipating future ai use cases, harms and benefits](#). *arXiv preprint arXiv:2403.14791*.
- Mauro P. Porto. 2012. [Media power and democratization in brazil: Tv globo and the dilemmas of political accountability](#).
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Nancy Roberts. 2004. [Public deliberation in an age of direct citizen participation](#). *The American Review of Public Administration*, 34(4):315–353.
- Shawn W. Rosenberg. 2007. [Rethinking Democratic Deliberation: The Limits and Potential of Citizen Participation](#). *Polity*, 39(3):335–360.
- Giovanni Rubeis, Keerthi Dubbala, and Ingrid Metzler. 2022. [“democratizing” artificial intelligence in medicine and healthcare: Mapping the uses of an elusive term](#). *Frontiers in Genetics*, 13.
- Johnny Saldana. 2021. *The coding manual for qualitative researchers*, 4 edition. SAGE Publications, London, England.
- Henry Sauermann, Katrin Vohland, Vyrion Antoniou, Bálint Balázs, Claudia Göbel, Kostas D. Karatzas, Peter Mooney, Josep Perelló, Marisa Ponti, Roeland Samson, and Silvia Winter. 2020. [Citizen science and sustainability transitions](#). *EcoRN: Citizen Science (Topic)*.
- Elizabeth Seger, Aviv Ovadya, Ben Garfinkel, Divya Siddarth, and Allan Dafoe. 2023. [Democratizing ai: Multiple meanings, goals, and methods](#). *ArXiv*, abs/2303.12642.
- Divya Siddarth. 2023. [How ai and democracy can fix each other](#).
- Julien Simon. 2022. [Intel and hugging face partner to democratize machine learning hardware acceleration](#).
- Mona Sloane, Emanuel Moss, Olaitan Awomolo, and Laura Forlano. 2022. [Participation is not a design fix for machine learning](#). In *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, pages 1–6.
- Andreas Sudmann. 2019. [The Democratization of Artificial Intelligence: Net Politics in the Era of Learning Algorithms](#), page 9–32. transcript Verlag.
- Andreas Sudmann and Alexander Waibel. 2019. [“That is a 1984 Orwellian future at our doorstep, right?”: Natural Language Processing, Artificial Neural Networks and the Politics of \(Democratizing\) AI](#), page 313–324. transcript Verlag.
- Zeerak Talat, Aurélie Névéol, Stella Biderman, Miruna Clinciu, Manan Dey, Shayne Longpre, Sasha Luccioni, Maraim Masoud, Margaret Mitchell, Dragomir Radev, Shanya Sharma, Arjun Subramonian, Jaesung Tae, Samson Tan, Deepak Tunuguntla, and Oskar Van Der Wal. 2022. [You reap what you sow: On the challenges of bias evaluation under multilingual settings](#). In *Proceedings of BigScience Episode #5 – Workshop on Challenges & Perspectives in Creating Large Language Models*, pages 26–41, virtual+Dublin. Association for Computational Linguistics.
- Wisdom J. Tettey. 2001. [The media and democratization in africa: contributions, constraints and concerns of the private press](#). *Media, Culture & Society*, 23:31–5.
- Pieter Verdegem. 2022. [Dismantling ai capitalism: the commons as an alternative to the power concentration of big tech](#). *AI & SOCIETY*, 39(2):727–737.
- Jan Philip Wahle, Terry Ruas, Mohamed Abdalla, Bela Gipp, and Saif Mohammad. 2023. [We are who we cite: Bridges of influence between natural language processing and other academic fields](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12896–12913, Singapore. Association for Computational Linguistics.
- Lindsay Weinberg. 2022. [Rethinking fairness: An interdisciplinary survey of critiques of hegemonic ML fairness approaches](#). *Journal of Artificial Intelligence Research*, 74:75–109.

Jack B. Weinstein. 2012. [The democratization of mass actions in the internet age](#). *Columbia Journal of Law and Social Problems*, 45:451.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Huggingface's transformers: State-of-the-art natural language processing](#).

Erik Olin Wright. 2010. *Envisioning Real Utopias*. Verso Books, London, England.

Wojciech Zaremba, Arka Dhar, Lama Ahmad, Tyna Eloundou, Shibani Santurkar, Sandhini Agarwal, and Jade Leung. 2023. [Democratic inputs to ai](#).

Justine Zhang, Arthur Spirling, and Cristian Danescu-Niculescu-Mizil. 2017. [Asking too much? the rhetorical role of questions in political discourse](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1558–1572, Copenhagen, Denmark. Association for Computational Linguistics.

A Methodological Details

Table 4 lists all false positive terms that we use in our first stage of manual filtering. Figure 5 shows the results of our PCA and clustering of embedded excerpts, with the darkest colour indicating the papers we select for reading and annotating fully.

democrat	Democratic National Committee	Project ANR Democrat
democrats	Liberal Democratic Party	Democrat system
Republican Democrat	Democratic Party	Description, Modélisation et Détection Automatique Des Chaînes de Référence
Democrat Republican	German Democratic Republic	DEMOCRAT
Republican and Democrat	Getman Democratic Republic	Democratic
Democrat and Republican	Democratic People's Republic of Korea	christian democratic parliamentary group
Republicans and Democrats	Christian Democratic Union	#democracy
Democrats and Republicans	Democratic Alliance	Democracy party
Republican or Democrat	United Democratic Front	Democrazia Cristiana / Christian Democracy
Democrat or Republican	Democratic Governors Association	#democratic_party
Republicans or Democrats	China Democracy Party	social-democratic political party
Democrats or Republicans	Christian Democrat	social-democratic leader
the Republican and the Democrat	Democratic primary	Center for Media and Democracy
the Democrat and the Republican	Democratic primaries	democratic president candidate
the Republicans and the Democrats	Somali Democratic Party	Stichting Democratie and Media (Democracy & Media Foundation)
the Democrats and the Republicans	New Democratic Party	Swedish social democratic politician
the Republican or the Democrat	Democratic Socialist Party	democratic congressman
the Democrat or the Republican	Liberal Democrat	social democratic movement
the Republicans or the Democrats	Democratic Left Alliance	Christian democratic
the Democrats or the Republicans	Alliance for Democracy in Mali	social democratic, centre-left political party
democratic and republican parties	Syrian Democratic Forces	Democratic Labour Party
Democratic Party of Japan	Democracy Now!	democratic republic of germany
Liberal Democratic Party of Japan	Movement for Democratic Change	Historical Press of the German Social Democracy Online
Social Democratic Party	Democracy Week	Forum voor Democratie, 'Forum for Democracy'
Democratic candidate	Democratic-controlled	centre-right party New Democracy
Democratic candidates	Croatian Democratic Union	Partito Democratico
Democratic republic of the Congo	Kurd Democratic Party	Social Democracy (S)
Democratic presidential candidate	New Democratic Union	Forum Migration and Democracy (MIDEM)
Democratic presidential candidates	ANR Democrat	

Table 4: False positives when matching “democra*” in corpus.

B Additional Results

B.1 All concepts and values

Values are understood to be a subset to concepts, which are regarded explicitly or implicitly as pertinent to a specific context (e.g., in relation to democracy). The authors undertake a subjective, direct democratic process to distinguish concepts from values. Tables 5 and 6 shows all concepts and values we find during excerpt annotation. Figure 6 shows the top concepts and values for each theme.

B.2 Where do (extra-disciplinary) references come from?

When considering where references are published, we find that the top five venues of all references are: NLP venues (*CL conferences), machine learning venues (ICLR/ICML/NeurIPS), arXiv, computer vision venues (CVPR/ICCV/ECCV), and the AAAI Conference on Artificial Intelligence. In contrast to this general pattern, extra-disciplinary references are mostly published in political science and social science journals, i.e., the American Political Science Review, Political Analysis, Nature, PloS ONE, Social Science Research Network, Science, and so on. The most frequently cited extra-disciplinary references are typically cited for methods, e.g., content analysis, agreement computations, discourse network analysis, or related to fake news and polarization. The most cited extra-disciplinary references in our corpus are:

1. Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts
2. A Coefficient of Agreement for Nominal Scales
3. Social Media and Fake News in the 2016 Election
4. Extracting Policy Positions from Political Texts Using Words as Data
5. Discourse Network Analysis: Policy Debates as Dynamic Networks




Figure 5: PCA and spectral clustering of excerpt embeddings, along with selected papers. Points that are the same color belong to the same cluster.

6. Measuring Political Deliberation: A Discourse Quality Index
7. Discourse Coalitions and the Institutionalization of Practice: The Case of Acid Rain in Great Britain
8. CUNY Academic Works
9. Exposure to Opposing Views on Social Media Can Increase Political Polarization

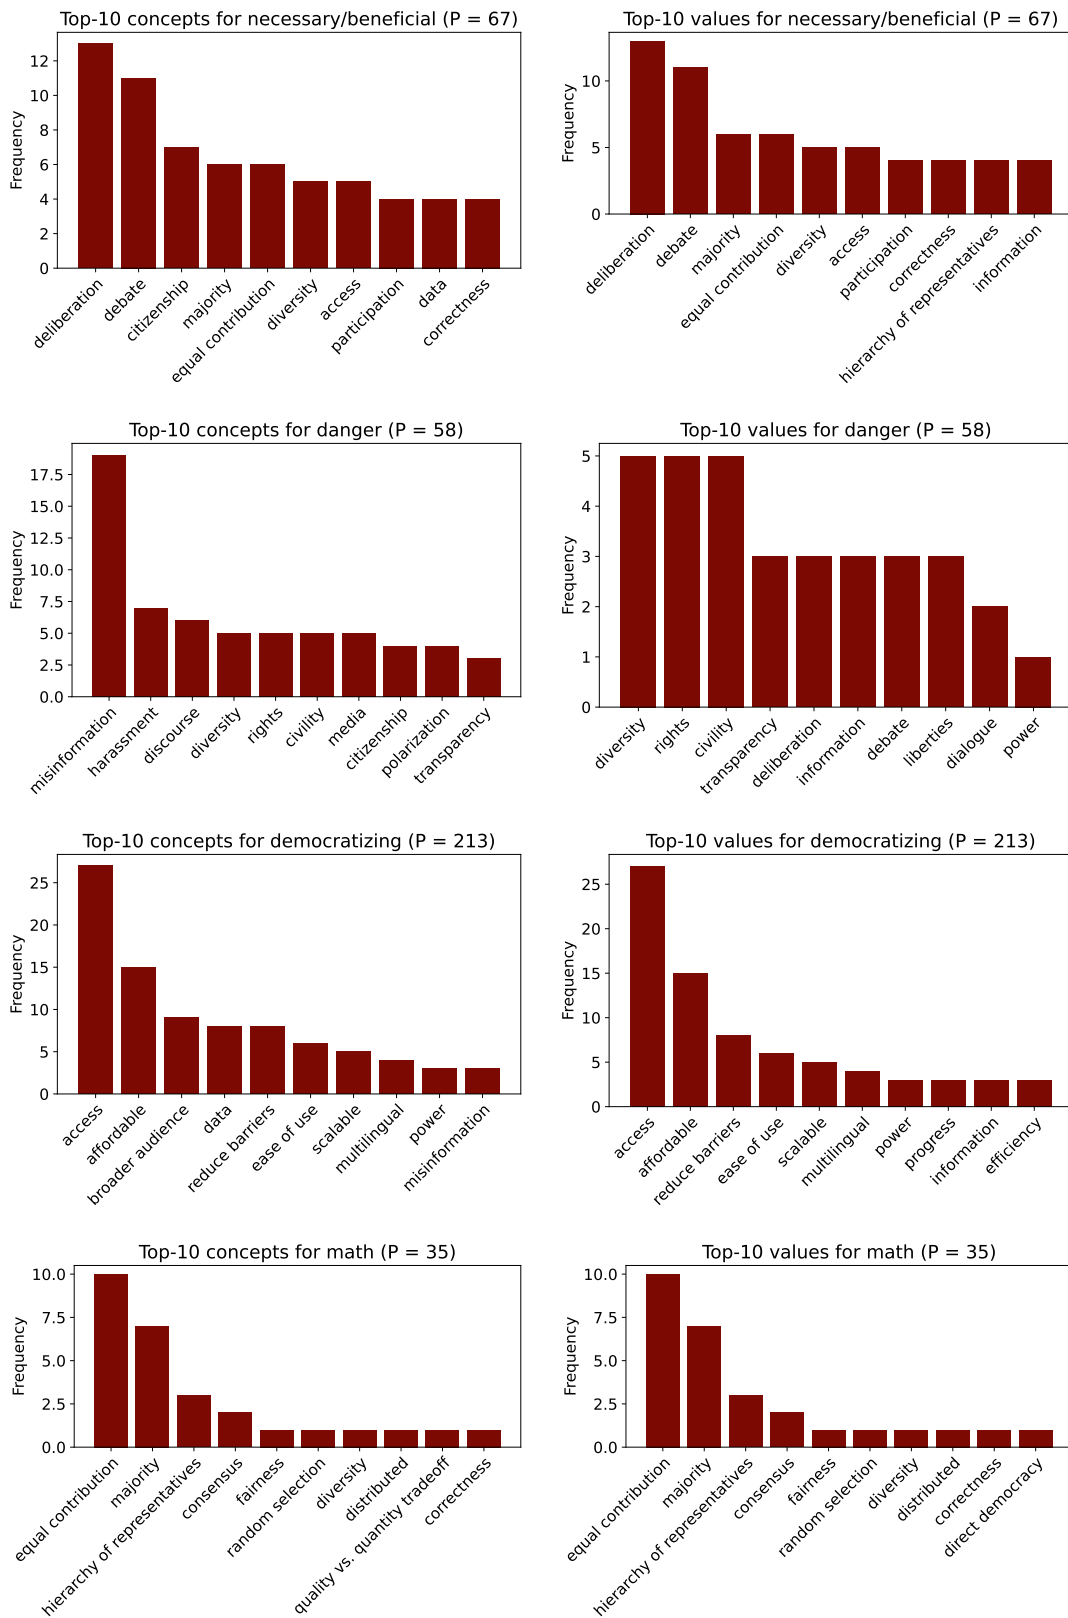


Figure 6: Frequency of concepts (left) and values (right) associated with democracy in papers, stratified by paper themes. For each theme, P refers to the number of papers annotated as having that type of theme.

generalizability	protection	dialogue
literacy	debate	decentralization
public opinion	freedom	sustainability
fairness	moderation	emotion
WEIRD	replicability	justice
liberties	environment	voting
anti-power	integrity	citizenship
equal contribution	resource-efficient	low-resource
interaction	engagement	broader audience
hierarchy of representatives	multilingual	scalable
rights	news	efficiency
governance	transparency	caution
acceleration	disagreement	civility
reduce barriers	protest	anxiety
discrimination	progress	data
translation	quality	access
happiness	reasoning	power
constitution	harassment	accountability
questioning	majority	consistency
competence	value	social good
reflection	open-source	cohesion
equal representation	evolving	polarization
informed	argument	campaign
fast	available	cooperation
representation	trust	information
responsibility	random selection	inclusion
diversity	quality vs. quantity tradeoff	direct democracy
political party	election	bill writing
correctness	affordable	choice
conflict	ease of use	discourse
equality	distributed	media
education	misinformation	discussion
privacy	participation	propaganda
complexity	critical	benefit
proficiency	ensorship	AI
rational	consensus	lack of prejudice
disinformation	deliberation	

Table 5: All associated concepts found when annotating excerpts.

sustainability	disagreement	moderation
fairness	caution	reduce barriers
argument	choice	justice
progress	optimality	direct democracy
trust	participation	rational
random selection	proficiency	resource-efficient
consensus	inclusion	diversity
available	critical	liberties
multilingual	engagement	cooperation
reasoning	interaction	efficiency
generalizability	benefit	open-source
integrity	accountability	reflection
literacy	transparency	access
social good	evolving	decentralization
civility	cohesion	informed
conflict	equal representation	equal contribution
majority	replicability	representation
correctness	equality	debate
privacy	power	distributed
quality	hierarchy of representatives	protection
deliberation	lack of prejudice	affordable
information	rights	discussion
ease of use	dialogue	happiness
responsibility	fast	anti-power
education	value	consistency
scalable	competence	

Table 6: All associated values found when annotating excerpts.