

# Participation Is not a Design Fix for Machine Learning

Mona Sloane\*

Institute for Public Knowledge, New York University, New York, USA; Tübingen AI Center, BMBF Competence Centre for Machine Learning (TUE.AI). Eberhard Karls University of Tübingen, Tübingen, Germany  
ms11521@nyu.edu

Olaitan Awomolo

Workday, Inc., Pleasanton, USA

Emanuel Moss

Cornell Tech Digital Life Initiative, New York, USA; Data & Society Research Institute, New York, USA  
em683@cornell.edu

Laura Forlano

Illinois Institute of Technology, Chicago, USA

## ABSTRACT

This paper critiques popular modes of participation in design practice and machine learning. It examines three existing kinds of participation in design practice and machine learning participation as work, participation as consultation, and as participation as justice – to argue that the machine learning community must become attuned to possibly exploitative and extractive forms of community involvement and shift away from the prerogatives of context independent scalability. Cautioning against “participation washing”, it argues that the notion of “participation” should be expanded to acknowledge more subtle, and possibly exploitative, forms of community involvement in participatory machine learning design. Specifically, it suggests that it is imperative to recognize design participation as work; to ensure that participation as consultation is context-specific; and that participation as justice must be genuine and long term. The paper argues that such a development can only be scaffolded by a new epistemology around design harms, including, but not limited to, in machine learning. To facilitate such a development, the paper suggests developing we argue that developing a cross-sectoral database of design participation failures that is cross-referenced with socio-structural dimensions and highlights “edge cases” that can and must be learned from.

## CCS CONCEPTS

• :: **Human-centered computing**; • **Ethnographic studies**; • **Social and professional topics**; • **Socio-technical systems**;

## KEYWORDS

machine learning, design, participatory methods

### ACM Reference Format:

Mona Sloane, Emanuel Moss, Olaitan Awomolo, and Laura Forlano. 2022. Participation Is not a Design Fix for Machine Learning. In *Equity and Access*

\*Mona Sloane’s work was supported by the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039A



This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 License.

EAAMO ’22, October 06–09, 2022, Arlington, VA, USA

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9477-2/22/10.

<https://doi.org/10.1145/3551624.3555285>

in *Algorithms, Mechanisms, and Optimization (EAAMO ’22)*, October 06–09, 2022, Arlington, VA, USA. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3551624.3555285>

## 1 INTRODUCTION

Over the past years, we have seen mounting evidence of the disparate impact of ML systems on already oppressed and disadvantaged groups [7, 11, 23, 50, 52]. The experiences of oppression and privilege are structural challenges that are incredibly complex, and they are not new particularly not to the communities that suffer from them. But they have heightened alongside the exponential growth of wealth inequality alongside planetary destruction [38, 55]. It is therefore both unsurprising and promising that the ML community wishes to build ML systems that are, for example, more democratic and more cooperative – often subsumed under the umbrella term of “more participatory” [8].

Whilst this is an honorable goal, we want to caution against a familiar-sounding impulse towards “participation washing” that we have seen in other areas of design and technology. For example, in the international development sector where “participation” of local communities at the receiving end of powerful agencies is based on manufactured consent and is based on (post-)colonial structures of global power [53]; in the corporate sector where “users” are invited into “co-creation” sessions in order to create new product ideas; in the philanthropic sector where “the public” is challenged to join in defining new problems and/or solutions to “wicked problems”; or in the urban design or architecture sector where stakeholder engagement protocols often legitimize injustices in the (material) planning of space and systematically devalue user needs as part of profit and scale-oriented design practices, or design inequality [65, 66].

## 2 PARTICIPATORY DESIGN

Participatory design methods can be traced to the 1970s when workers in Scandinavia worked together collaboratively to design the technologies that they would use in corporate settings [59, 60, 67]. Over the past several decades, participatory design and related concepts such as codesign and co-creation have been introduced as a way of engaging with ethics, values in design [49], value-sensitive design (Friedman 1996), and values levers in design [63]. Participatory design with its rich history in socially democratic countries in Europe, has sought to engage multiple stakeholders in deliberative processes in order to achieve consensus. At the same time, other

approaches have emphasized agonism and the importance of dissent, friction and disagreement [20, 36, 40, 47]. In this tradition of participatory design, the focus has been on designing publics [19] to engage in matters of concern around complex socio-technical systems. In order to facilitate the engagement of multiple stakeholders in participatory design processes, designers often use prototypes, games [24] and other structured activities.

More recently, scholars have argued that nonhuman actors such as algorithms and machines [14] as well as the multispecies (microbes, plants, animals and the natural environment) be considered as stakeholders in participatory design processes [25, 28, 29, 37]. Finally, with the introduction of critical and speculative design and experiential futures in the early 2000s, design researchers have become interested in the ways in which participatory design and design futures might come together to create new modes of experiential futures [12], design fiction [6], speculative design [21], speculative civics [20] and critical futures [25, 28, 29] in order to think through the social consequences of emerging technologies.

Participatory design methods have often been seen as a way of overcoming supposed difficulties that users have in understanding ostensibly complex technologies, particularly in healthcare settings [48]. Participatory methods have also been employed where designers anticipate public resistance or skepticism to a product or service [3]. The use of participatory methods in technology settings follows the development of participatory methods in other domains, particularly international development [53] where participation was seen as a means for overcoming local resistance to international development schemes [33].

ML already incorporates certain forms of participation throughout the design of models and their integration into society, however participatory design practices from other domains hold important lessons for ML. We will expand the notion of “participation” beyond the forms of involvement that are commonly understood as participatory design. Following the review of key literature on participatory design and ML, we will introduce three different forms of participation: participation as work, participation as consultation, participation as justice, each illustrated with a list of examples. Through this framing, it becomes possible to understand how participatory design, a necessarily situated and context-dependent endeavor, articulates with industrial prerogatives of context-independent scalability. It also becomes possible to recognize where the discourse of participation fails to account for existing power dynamics and obscures the extractive nature of collaboration, openness, and sharing, particularly in corporate contexts. We conclude the paper with a set of recommendations drawn from considering a more expansive definition of participation in the context of ML.

### 3 DIFFERENT FORMS OF PARTICIPATION

#### 3.1 Participation as Work

Much of ML plays out upon what is an intensely participatory field. Whether acknowledged or not, a broad range of participants play an important role in producing the data that is used to train and evaluate ML models. For example, ImageNet, which laid the foundations for deep learning and most image recognition applications and is still used for ML benchmarking, is a dataset of millions of images,

taken by hundreds of thousands of people, scraped from the open web and labeled by mTurk workers [41]. Image classification tools are often built on top of models trained on the ImageNet dataset. Photographers, web designers, and mTurk workers all participate in every such application. A similar case presents itself for Natural Language Processing applications which, for over a decade, have sourced from Wikipedia for training language corpora [31].

Recently, industrial players in ML have taken steps that appear to acknowledge the ethical concerns that have been raised around indiscriminate aggregation of accessible data into large datasets. These concerns are that such datasets contain harmful, stereotypical labels [5], that they were collected without adequately informed consent [17], and that they are imbalanced across demographic groups in ways that lead to bias in ML systems trained on such datasets [69]. Subsequently, Google has sent legions of contractors out into the streets to gather ‘selfies’ of minority groups’ faces, offering nominal rewards in exchange for signatures on consent forms (Dillon & Collett). Other efforts have sought to curate datasets that are balanced with respect to demographic categories and are well-documented as to their appropriate uses and limitations. As has been pointed out [30], subjects of these data efforts are often offered compensation they are relatively powerless to refuse, and are often given partial or deceptive information to obtain their consent. While these curatorial efforts are motivated by concerns about potential biases and informed by best practices in research ethics, they point to the several ways ordinary people are incorporated into the work of producing ML, despite having little power to affect the terms of their participation.

Billions of ordinary web users also continually participate in the production and refinement of ML, as their online (and offline) activities produce neatly labeled rows of data on how they click their way around the web, navigate their streets, and engage in any number of other commercial, leisure, or romantic activities [45]. Users also improve the performance of ML models as they interact with them, a single unanticipated click can update a model’s parameters and future accuracy. This work sometimes is so deeply integrated into the ways in which users navigate the Internet that it is performed unconsciously, e.g. when using Google Maps and producing data movement patterns that enable traffic predictions. But other times it becomes more conscious, e.g. when classifying photos when completing a reCAPTCHA [51], or ranking Uber drivers [57]. Where ML technology does not live up to its mythos, people work behind the veil to complete tasks as if by the magic of AI. Behind some mobile apps claiming to use AI are real people transcribing images of paper receipts and populating a purchase history database [34], moderating content [56]. The labor of integrating new technologies, such as AI applications, into everyday life and existing work processes and even out their rough edges, e.g. in healthcare [62], is the “human infrastructure” without which the socio-technical system cannot function [22]. Labor, here, is multi-layered and includes affective and emotional labor, e.g. coping with stress and sleep-deprivation when integrating medical devices into everyday life [27], or social labor, e.g. when explaining ML outcomes to users or even out their glitches such as when chatbots fail. All this work often happens without consent or acknowledgement, and remains uncompensated. Such ML design processes are cases of “designing

for”, i.e. processes that are void of a genuine integration of design users, relying on them to make the design product work ex post. The role data subjects, crowdworkers, and end users play in producing ML is often unacknowledged. Even when acknowledged, these roles are not seen as “participatory” in the senses commonly given in participatory research and development literatures. However, many of the strongest claims about the power and promise of ML depend on how closely linked an ML system is to the fine-grained, intimate data it draws upon. We argue that this closeness, between developer, technology, and subject, is much the same in ML as it is in participatory research and development work, and ought to be acknowledged as such. Acknowledging collaborative work that contributes to ML as participation is crucial for drawing attention to imbalanced power relations between developers and users, identifying responsibilities owed to users by developers by virtue of these relationships [46], and working toward greater equity across the ML ecosystem.

### 3.2 Participation as Consultation

In the case of participation as consultation, cf. [44], designers and technologists engage in episodic, short-term projects in which diverse stakeholders might be consulted at various stages of the process. This model is most common in architecture and urban planning as well as among major philanthropic foundations and private corporations. Architecture and urban planning practices use citizen participation approaches to engage different stakeholder groups in project development. As these projects are complex and have significant socio-economic impacts on communities, participatory workshops can provide an integrated framework where experts work with stakeholder groups to identify context-specific needs [9, 56]. Here, participation might be facilitated through small, face-to-face workshops or larger design sprints or hackathons as well as through the use of online platforms for crowdsourcing ideas.

There are several challenges that can limit the effectiveness of participation as consultation. For a variety of reasons including intellectual property concerns, in this model, long-term partnerships are either impossible, undesirable, unnecessary or cost prohibitive. As this type of top-down design process also takes the form of “designing for” a particular group without an ongoing commitment to their inclusion in the process, systemic inequalities that can be hard-coded into consultation and representation protocols [65]. Experts do not often have a good understanding of how to design effective participatory processes or engage the right stakeholders to achieve the desired outcomes. A third challenge occurs as cities begin to require participation workshops as part of the permitting and approvals process. Participation workshops can become performative, where experts do not actually take the needs or recommendations of the different stakeholder groups into consideration [15].

### 3.3 Participation as Justice

In the case of participation as justice, designers and technologists typically see to engage in more-long term partnerships with diverse stakeholders, because they are interested in addressing structural issues that lead to harm, for example by creating mechanisms for recourse. Here, all members of the design process engage in more tightly coupled relationships with more frequent communication

(which often happens through a blended communication and interaction approach, e.g. online/offline). The canon of participation as collaboration notably comprises participatory action research, which is focused on researchers and participants undertake action-oriented and self-reflexive practices that leads to them having more control over their lives [4]; infrastructuring, which centers designers’ locations, the materials and systems intrinsic to designing, as well as (community) capacity building [1, 39, 42]; design justice, which goes beyond value-focused design and centers typically marginalized groups in collaborative and creative design processes that challenge and dismantle the matrix of domination, i.e. white supremacy, heteropatriarchy, capitalism, and settler colonialism [16]; crip technoscience, which refuses demands to eliminate disability, underscores that disabled people are expert designers of everyday life, and centers technoscientific activism, critical design practices, and disability justice [35]; data feminism, which focuses on ideas of intersectional feminism [18]; and tech activism and resistance, both from people affected by potentially harmful technology, such as the Atlantic Towers Resident Association in Brooklyn, NY [32] and those designing it, see for example the Tech Worker Movement [68], or a mix of both, such as Data for Black Lives, Black in AI, or LatinX in AI. What ties these approaches together is favoring using language around “designing *with*” in order to ensure that outcomes are valuable to people from diverse backgrounds and communities, including the disability community. Participation as justice has social and political importance, but it may be difficult to do it well, especially in a corporate context where design justice can almost be seen as an oxymoron: given the extractive and often oppressive capitalist logics and contexts of ML systems in corporate environments, it appears impossible to design ML products that are genuinely “just” and “equitable”. A similar trajectory can unfold when ML products are deployed in the public sector.

## 4 CRITIQUES OF PARTICIPATION

The dominant mode of extraction within the ML industry is deeply entangled with the capitalist paradigm of scale, referring to the ability to gain revenue at a greater proportion per unit cost of inputs [13]. Participation, which can be extractive at the level of individual micro-interactions, is *necessarily* extractive at the level of ML design, development, and deployment. As a tech industry buzzword, the verb “to scale” refers to the ability of products to spread far beyond the context of development to new applications in new markets. Part of the promise of ML is that statistical generalizations learned from finite datasets will allow for inferences to be made across broader contexts, and that capabilities engendered by ML can be applied to additional settings without adding proportional costs. However, datasets are deeply context-bound, and that context, as well as the appropriateness of the use of those datasets, is lost in the scaling of ML applications [9].

Acknowledging the modes of participation that are already components of ML challenges understandings of how these tools are able to scale. As such technologies scale across contexts, the generalizations that are learned inevitably require updating, by providing additional training data or correcting errors [61]. This often requires the participation of users interacting with the system who experience the friction of providing additional information to the system

(as with CAPTCHAs) or bearing the burden of system errors. As discussed above, representation/consultation is often prohibitively costly. Where a cost-benefit analysis may encourage such forms of participation in the earlier stages of product development, in later stages that product is expected to scale without incurring additional costs. The initial utility of representative and consultative forms of participation are thus diluted as products scale beyond the context in which that mode of participation contributed to the overall design of the product in earlier stages. For ML products to simultaneously scale and engage in meaningful partnerships oriented toward justice, they also require additional inputs of participation, and budgets must be set aside for that.

This can be thought of as levelling the playing field of futuring: product futures are often made very concrete for venture capitalists. But what kind of imaginative work do entrepreneurs do when it comes to the communities that they seek out as users (or targets) of their products? There is an existing imbalance between market-fit and community-fit. To address that and pave the way for design justice processes to become integral to ML, it is key to expand the notion of value beyond monetary value and the extractive logics underpinning the invasive data collection that is necessitated by most ML system designs. Promising developments have recently been made in the context of Indigenous data sovereignty which includes access, control and governance of Indigenous data [2]. Against that backdrop we suggest three cues for considering participation in ML in a more equitable way:

- Recognize participation as work. Users already labor in, for, and through ML systems across a number of dimensions (affective, social, emotional). This labor upholds and improves ML systems and therefore is valuable for the owners of the ML systems. To acknowledge that, users should be asked for consent, be provided with opt-out options or alternatives, and, if they chose to participate through labor, be offered compensation. This could mean to clarify when and how data generated by user behavior is used for the training and improvement of ML systems (e.g. via a banner on the Wikipedia page, or in Google Maps); to give an alternative security option for reCAPTCHA; to not punish users for refusing to leave reviews; to provide appropriate support for content moderators; to compensate “ghost workers” fairly [33]; to develop reward systems for users that labor to integrate technologies into their lives and thereby provide rich data for profit-oriented ML companies.
- Participation as consultation must be designed for specific contexts. If short-term participation is the most feasible and desired version for ML participation, then there needs to be a commitment to context-specificity, especially in terms of how the participation is facilitated. Every context is different, so participation has to be designed to address these different contexts. Rather than a one-size-fits-all approach, consultation and representation processes must be revisited and reexamined to ensure they are gathering the right information from the right people. As ML systems affect a wide range of groups, marginalized stakeholders should be given the space and voice to co-design and co-produce these systems

[16]. Documenting these processes and their contexts can form a knowledge base for long term, effective participation.

- Participation as justice must be genuine and long term. This means to engage in creating processes that provide transparency and genuine knowledge sharing. This can be difficult particularly for proprietary design cases. Further, using the language of design justice without actually engaging in actual design justice processes and practices can only lead to corporate co-optation. For example, the ML field has seen a hype of “ethical AI” serving as a smokescreen for continuing with non-participatory and non-justice oriented ML design approaches [63], despite good intentions. To avoid that, it may be helpful to make the tensions that characterize the goal of long term participation in ML visible, acknowledging that partnerships and justice do not scale in frictionless ways, but require constant maintenance and articulation with existing social formation in new contexts (Tsing).

As we discuss in this paper using the “participatory design” framework, harms can be produced by the same ways of thinking that produce the technology that causes the harms. To address this issue, we argue that it is crucial to enhance the ability for lateral thinking across applications, domains, and (academic) disciplines in order to enable “holistic futuring”. Such an approach could facilitate the development of an epistemology of (design) harms [54] or “design inequalities” [64]. Here, the notion of design harms or design inequalities should be linked to concrete cases, including cases in which participatory design approaches were deployed, but ended up perpetuating design harms. A new epistemology around “holistic futuring” should be focused on demonstrating the connection between social structures and design practice. To facilitate these efforts, we propose to develop a searchable database of design precedents across applications and disciplines that highlights design failures, especially failures of design participation, cross-referenced with socio-structural dimension (e.g. issues pertaining to racial inequality, or class-based inequity). This database should cover design projects across all sectors and domains, not just ML, and explicitly acknowledge deliberate absences and outliers which often are the most interesting and relevant social phenomena we can learn from (e.g. transgender identities). It may also acknowledge and educate on the deliberate refusal to “get counted” [18].

## 5 CONCLUSION

In this paper, we have cautioned against “participation-washing” of ML by critically examining the existing kinds of participation in design practice and ML. Existing forms of participation can be classified as work, as consultation, and as justice, but we have argued that the notion of “participation” should be expanded to acknowledge more subtle, and possibly exploitative, forms of community involvement in participatory ML design. This framing allows for understanding participatory design as a necessarily situated and context-dependent endeavor which is at odds with industrial prerogatives of extraction and context-independent scalability. Against that backdrop, it is imperative to recognize design participation as work; to ensure that participation as consultation is context-specific; and that participation as justice must be genuine and long

term. Therefore, we argue that developing a cross-sectoral database of design participation failures that is cross-referenced with socio-structural dimensions and highlights “edge cases” that can and must be learned from.

## REFERENCES

- [1] Agid, S. “...it’s your project, but it’s not necessarily your work...”: infrastructuring, situatedness, and designing relational practice. In Proceedings of the 14th Participatory Design Conference on Full papers PDC ’16, pp. 81–90. ACM Press. ISBN 978-1-4503-4046-5. doi: 10.1145/2940299.2940317.
- [2] Anderson, J. and Hudson, M. IP and indigenous data sovereignty: The traditional knowledge and biocultural labels and notice system. In NYU Engelberg Innovation Colloquium Paper.
- [3] Asaro, P. M. Transforming society by transforming technology: the science and politics of participatory design. 10:34.
- [4] Baum, F., MacDougall, C., and Smith, D. Participatory action research. 60(10):854–857. ISSN 0143-005X. doi: 10.1136/jech.2004.028662.
- [5] Birhane, A., Prabhu, V. U., and Kahembwe, E. Multimodal datasets: misogyny, pornography, and malignant stereotypes. arXiv preprint arXiv:2110.01963, 2021.
- [6] Bleecker, J. Design fiction: A short essay on design, science, fact and fiction. URL <https://bit.ly/2B9rqjg>. Type: Journal Article.
- [7] Bolukbasi, T., Chang, K.-W., Zou, J. Y., Saligrama, V., and Kalai, A. T. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In Proceedings of the 30th Conference on Neural Information Processing Systems, pp. 9.
- [8] Bondi, E., Xu, L., Acosta-Navas, D., and Killian, J. A. Envisioning communities: A participatory approach towards AI for social good. In Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society, pp. 425–436. ACM. ISBN 978-1-4503-8473-5. doi: 10.1145/3461702.3462612. URL <https://dl.acm.org/doi/10.1145/3461702.3462612>.
- [9] boyd, d. and Crawford, K. Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon. 15(5):662–679. ISSN 1369-118X, 14684462. doi: 10.1080/1369118X.2012.678878.
- [10] Bratteteig, T. and Verne, G. Does AI make PD obsolete?: exploring challenges from artificial intelligence to participatory design. In Proceedings of the 15th Participatory Design Conference on Short Papers, Situated Actions, Workshops and Tutorial PDC ’18, pp. 1–5. ACM Press. ISBN 978-1-4503-5574-2. doi: 10.1145/3210604.3210646.
- [11] Buolamwini, J. and Gebru, T. Gender shades: Intersectional accuracy disparities in commercial gender classification. In Proceedings of Machine Learning Research, volume 18.
- [12] Candy, S. The futures of everyday life: Politics and the design of experiential scenarios. Candy, S. and Dunagan, J. Designing an experiential scenario: the people who vanished. 86:136–153. ISSN 0016-3287. Type: Journal Article.
- [13] Chandler, A. D. and Hikino, T. Scale and scope: the dynamics of industrial capitalism. Belknap Press. ISBN 978-0-674-78994-4.
- [14] Choi, J. H.-j., Forlano, L., and Kera, D. Situated automation: Algorithmic creatures in participatory design. In Proceedings of the Participatory Design Conference.
- [15] Costanza-Chock, S. Design Justice. MIT Press. Type: Book.
- [16] Crosby, Bryson, Quick, and Slotterback. Designing public participation processes. In Leadership for the common good: tackling public problems in a shared-power world, Public Administration Review.
- [17] Denton, E., Hanna, A., Amironesei, R., Smart, A., and Nicole, H. On the genealogy of machine learning datasets: A critical history of imagenet. *Big Data & Society*, 8(2):20539517211035955, 2021.
- [18] D’Ignazio, C. and Klein, L. F. Data Feminism. Strong ideas series. The MIT Press. ISBN 978-0-262-04400-4. Dillon, S. and Collett, C. AI and gender: Four proposals for future research. URL <https://www.repository.cam.ac.uk/handle/1810/294360>. Accepted: 2019-07-04T11:32:59Z.
- [19] DiSalvo, C. Design and the construction of publics. 25(1): 48–63. Type: Journal Article.
- [20] DiSalvo, C., Jenkins, T., and Lodato, T. Designing speculative civics. In Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, pp. 4979–4990. ACM. ISBN 1-4503-3362-1. Type: Conference Proceedings.
- [21] Dunne, A. and Raby, F. Speculative Everything: Design, Fiction, and Social Dreaming. MIT Press. ISBN 0-26201984-1. Type: Book.
- [22] Elish, M. C. and Mateescu, A. AI in context: The labor of integrating new technologies.
- [23] Eubanks, V. Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor. St. Martin’s Press.
- [24] Flanagan, M. and Nissenbaum, H. Values at Play in Digital Games. MIT Press. Type: Book.
- [25] Forlano, L. Posthumanism and design. 3(1):16–29. Type: Journal Article.
- [26] Forlano, L. Stabilizing/destabilizing the driverless city: Speculative futures and autonomous vehicles. 13. b. Type: Journal Article.
- [27] Forlano, L. The Danger of Intimate Algorithms. Public Books, April 2020. Type: Magazine Article.
- [28] Forlano, L. and Halpern, M. Reimagining work: Entanglements and frictions around future of work narratives. (26):32–59. Type: Journal Article.
- [29] Forlano, L. and Mathew, A. From design fiction to design friction: Speculative and participatory design of valueembedded urban technology. 21(4):7–24. Type: Journal Article.
- [30] Fussell, S. How an attempt at correcting bias in tech goes wrong. URL <https://www.theatlantic.com/technology/archive/2019/10/googleallegedly-used-homeless-train-pixelphone/599668/>. Section: Technology.
- [31] Gabrilovich, E. and Markovitch, S. Wikipedia-based semantic interpretation for natural language processing. 34:443–498. ISSN 1076-9757. doi: 10.1613/jair.2669.
- [32] Gagne, Y. How we fought our landlord’s secretive plan for facial recognition—and won. URL <https://bit.ly/2Ce9ODt>. Library Catalog: [www.fastcompany.com](http://www.fastcompany.com).
- [33] Goldman, M. Imperial Nature: The World Bank and Struggles for Social Justice in the Age of Globalization. Yale Agrarian Studies Series. Yale University Press.
- [34] Gray, M. L. and Suri, S. Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass. Houghton Mifflin Harcourt. ISBN 978-1-328-56624-9.
- [35] Hamraie, A. and Fritsch, K. Crip technoscience manifesto. 5(1):1–33. ISSN 2380-3312, 2380-3312. doi: 10.28968/cftt.v5i1.29607.
- [36] Hansson, K., Forlano, L., Choi, J. H.-j., DiSalvo, C., Pargman, T. C., Bardzell, S., Lindtner, S., and Joshi, S. Provocation, conflict, and appropriation: the role of the designer in making publics. 34(4):3–7. ISSN 0747-9360. Type: Journal Article.
- [37] Heitlinger, S., Foth, M., Clarke, R., DiSalvo, C., Light, A., and Forlano, L. Avoiding eocidal smart cities: participatory design for more-than-human futures. In Proceedings of the 15th Participatory Design Conference: Short Papers, Situated Actions, Workshops and Tutorial. Volume 2, pp. 51. ACM. ISBN 1-4503-5574-9. Type: Conference Proceedings.
- [38] Hickey, J. The Divide Global Inequality from Conquest to Free Markets. W. W. Norton & Company.
- [39] Hillgren, P.A., Seravalli, A., and Emilson, A. Prototyping and infrastructuring in design for social innovation. 7(3): 169–183. ISSN 1571-0882. Type: Journal Article.
- [40] Keshavarz, M. and Maze, R. Design and dissensus: framing and staging participation in design research. 11(1): 7–29. ISSN 1448-7136. Type: Journal Article.
- [41] Krizhevsky, A., Sutskever, I., and Hinton, G. E. ImageNet classification with deep convolutional neural networks. 60(6):84–90. ISSN 0001-0782, 1557-7317. doi: 10.1145/3065386. URL <https://dl.acm.org/doi/10.1145/3065386>.
- [42] Langley, P. Crafting papers on machine learning. In Langley, P. (ed.), Proceedings of the 17th International Conference on Machine Learning (ICML 2000), pp. 1207–1216, Stanford, CA, 2000.
- [43] Morgan Kaufmann. Le Dantec, C. A. and DiSalvo, C. Infrastructuring and the formation of publics in participatory design. 43(2):241–264. ISSN 0306-3127. Type: Journal Article.
- [44] Martin Jr., D., Prabhakaran, V., Kuhlberg, J., Smart, A., and Isaac, W. S. Participatory Problem Formulation for Fairer Machine Learning Through Community Based System Dynamics. arXiv:2005.07572 [cs, stat], May 2020.
- [45] Mayer-Schönberger, V. and Cukier, K. Big Data: A Revolution that Will Transform how We Live, Work, and Think. Houghton Mifflin Harcourt. ISBN 978-0-544-00269-2.
- [46] Metcalf, J., Moss, E., Singh, R., Tafese, E., and Watkins, E. A. A relationship and not a thing: A relational approach to algorithmic accountability and assessment documentation. pp. 19. URL <http://arxiv.org/abs/2203.01455>.
- [47] Mouffe, C. Pluralism, dissensus and democratic citizenship. pp. 1–10. Type: Journal Article.
- [48] Neuhauser, L. and Kreps, G. L. Participatory design and artificial intelligence: Strategies to improve health communication for diverse audiences. In AI and Health Communication – Papers from the AAAI 2011 Spring Symposium (SS-11-01), pp. 4. AAAI.
- [49] Nissenbaum, H. How computer systems embody values. 34(3):117–119. Type: Journal Article.
- [50] Noble, S. U. Algorithms of Oppression: How Search Engines Reinforce Racism. NYU Press.
- [51] O’Malley, J. Captcha if you can: how you’ve been training AI for years without realizing it | TechRadar. URL <https://bit.ly/37H1esA>.
- [52] O’Neal, C. Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy. Crown.
- [53] Peet, R. and Hartwick, E. Theories of Development: Contentions, Arguments, Alternatives. The Guildford Press.
- [54] Perez, C. C. Invisible women: Data bias in a world designed for men. Abrams, 2019.
- [55] Picketty, T. Capital in the Twenty-first Century. The Belknap Press of Harvard University Press. ISBN 978-0-67443000-6.
- [56] Roberts, S. T. Behind the Screen: Content Moderation in the Shadows of Social Media. Yale University Press. ISBN 978-0-300-23588-3. OCLC: on105263168.
- [57] Rosenblat, A. Uberland: How Algorithms are Rewriting the Rules of Work. Univ of California Press.
- [58] Saad-Sulonen, J. and Horelli, L. The value of community informatics to participatory urban planning and design a case study in helsinki. 6(2).

- [59] Sanders, E. B.-N. From user-centered to participatory design approaches. pp. 1–8. Type: Journal Article.
- [60] Schuler, D. and Namioka, A. *Participatory design : principles and practices*. L. Erlbaum Associates. ISBN 0805809511 (cloth) 080580952X (paper). Type: Book.
- [61] Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., and Vertesi, J. Fairness and abstraction in sociotechnical systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency FAT\* '19*, pp. 59–68. ACM Press. ISBN 978-1-4503-6125-5. doi: 10.1145/3287560.3287598.
- [62] Sendak, M., Elish, M. C., Gao, M., Futoma, J., Ratliff, W., Nichols, M., Bedoya, A., Balu, S., and O'Brien, C. "The human body is a black box": supporting clinical decision-making with deep learning. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pp. 99–109. ACM. ISBN 978-14503-6936-7. doi: 10.1145/3351095.3372827.
- [63] Shilton, K. Values levers: Building ethics into design. 38 (3):374–397. doi: 10.1177/0162243912436985. Type: Journal Article.
- [64] Sloane, M. Inequality is the name of the game: Thoughts on the emerging field of technology, ethics and social justice. doi: 10.34669/WI.CP/2.9. Publisher: WI Weizenbaum Institute for the Networked Society Version Number: 1.
- [65] Sloane, M. On the need for mapping design inequalities. 35(4):3–11. ISSN 0747-9360, 1531-4790. doi: 10.1162/desi.a.00559.
- [66] Sloane, M. *The shelves that won't hold: Material politics and social inequality in spatial design practice*. doi: 10.1162/desi.a.00565. Publisher: MIT Press One Rogers Street, Cambridge, MA 02142-1209 USA journals-info@mit.edu.
- [67] Spinuzzi, C. The methodology of participatory design. 52 (2):163–174. ISSN 0049-3155. Type: Journal Article.
- [68] Tarnoff, B. *The Making of the Tech Worker Movement*. Logic. Library Catalog: logicmag.io. Tsing, A. *Friction: An Ethnography of Global Connection*. Princeton University Press.
- [69] Yang, K., Qinami, K., Fei-Fei, L., Deng, J., and Russakovsky, O. Towards fairer datasets: Filtering and balancing the distribution of the people subtree in the imagenet hierarchy. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pp. 547–558, 2020.