

# Guiding Principles for Participatory Design-inspired Natural Language Processing

Tommaso Caselli<sup>✉</sup>, Roberto Cibin<sup>✉</sup>, Costanza Conforti<sup>✉</sup>

Enrique Encinas<sup>✉</sup>, Maurizio Teli<sup>✉</sup>

<sup>✉</sup> University of Groningen

<sup>✉</sup> Institute of Sociology of the Czech Academy of Sciences & Masaryk University

<sup>✉</sup> Rural Senses Ltd., <sup>✉</sup> University of Aalborg

<sup>✉</sup> t.caselli@rug.nl, <sup>✉</sup> roberto.cibin@soc.cas.cz,

<sup>✉</sup> stanze@ruralsenses.com, {eencinas@cs|maurizio@plan}.aau.dk

## Abstract

We introduce 9 guiding principles<sup>1</sup> to integrate Participatory Design (PD) methods in the development of Natural Language Processing (NLP) systems. The adoption of PD methods by NLP will help to alleviate issues concerning the development of more democratic, fairer, less-biased technologies to process natural language data. This short paper is the outcome of an ongoing dialogue between designers and NLP experts and adopts a non-standard format following previous work by Traum (2000); Bender (2013); Abzianidze and Bos (2019). Every section is a guiding principle. While principles 1–3 illustrate assumptions and methods that inform community-based PD practices, we used two *fictional design scenarios* (Encinas and Blythe, 2018), which build on top of situations familiar to the authors, to elicit the identification of the other 6. Principles 4–6 describes the impact of PD methods on the design of NLP systems, targeting two critical aspects: data collection & annotation, and the deployment & evaluation. Finally, principles 7–9 guide a new reflexivity of the NLP research with respect to its *context, actors and participants*, and *aims*. We hope this guide will offer inspiration and a road-map to develop a new generation of PD-inspired NLP.

## 1 PD is about consensus and conflict

PD has its origin in Scandinavia forty years ago, when it was articulated as an *offensive strategy* for the trade union movement to promote industrial democracy (Group, 1981; Ehn, 1992). PD was seen as a way to allow workers to shape the technologies they would use at the workplace (Ehn, 2016).

<sup>1</sup>The principles are guided by the authors experience, primarily focused in Europe (with the exception of one of them). However, we would defend the applicability of most of them to a wider range of contexts, with the situated effort of appropriation and transformation that is an integral part of PD.

As a form of system design performed *with and by* people (Briefs et al., 1983), PD entails a process of mutual learning among participants, among design researchers, and between design researchers and participants (Simonsen and Robertson, 2012). Traditionally, that means adopting a variety of research and design methods, from workshops (Ehn et al., 1996) to participant observations (Blomberg and Karasti, 2012a), passing through cards (Teli et al., 2017) or games (Vaajakallio and Mattelmäki, 2014), to include scenarios (Bødker, 2000), prototypes (Kannabiran and Bødker, 2020), and many others. The appropriate combination of methods and activities is determined, in a situated way, beginning with the involvement of different social groups (Bratteteig et al., 2012).

Historically, PD questions who is involved in the design process from various communities (DiSalvo et al., 2012) to specific socio-economic actors (Teli, 2015) and how. As a consequence, the design process can and should reflect on the visions for social transformation that the participants can develop (Huybrechts et al., 2020; Helgason et al., 2020), by translating those visions into alternatives to existing technologies (Korsgaard et al., 2016).

## 2 Design is an inherently disordered and unfinished process

Being based on nurturing relations between professional technology designers and members of the various social groups they interact with, PD methods and practices acknowledge that designing digital technologies with non-professionals does not follow a linear model (Callon, 2004; Cibin et al., 2020). Even when formalized (Bratteteig et al., 2012), the design process is disordered and unfinished. This character is well represented by the expressions *use-before-use* and *design-after-design* (Ehn, 2008; Fry, 2017).

PD principles	1. <i>PD is about consensus and conflict</i>	<ul style="list-style-type: none"> <li>• PD entails a process of mutual learning between researchers and community</li> <li>• PD adopts a variety of research and design methods (workshops, participants observation, cards, ...)</li> </ul>
	2. <i>Design is an inherently disordered and unfinished process</i>	<ul style="list-style-type: none"> <li>• <i>Use-before-use</i>: tool’s use is envisioned <i>before</i> the tool is actually implemented</li> <li>• <i>Design-after-design</i>: tool’s design isn’t exhausted with delivery, but will be modified by the users’ appropriation, use, and feedback</li> </ul>
	3. <i>Communities are often not completely determined a priori</i>	<ul style="list-style-type: none"> <li>• Communities are not a unitary whole, but can get formed within and through the design process</li> </ul>
NLP tools	4. <i>Data and communities are not separate things</i>	<ul style="list-style-type: none"> <li>• The shift from <i>language as data</i> to <i>language as people</i>: language data are produced by human speakers</li> <li>• Communities should be involved in the different stages of the NLP pipeline</li> </ul>
	5. <i>Community involvement is not scraping</i>	<ul style="list-style-type: none"> <li>• Collaboration with a community should imply ethical engagement practices based on respect, equity and reciprocity</li> <li>• Researchers should communicate to the community the usage of the collected data in a transparent and appropriate way</li> </ul>
	6. <i>Never stop designing</i>	<ul style="list-style-type: none"> <li>• Community adaptation should be treated as a feature of an NLP system at the design stage</li> </ul>
Researchers’ reflexivity	7. <i>Text is a means rather than an end</i>	<ul style="list-style-type: none"> <li>• The linguistic output of NLP systems should serve people’s needs rather than imitate people’s production of language.</li> </ul>
	8. <i>The thin red line between consent and intrusion</i>	<ul style="list-style-type: none"> <li>• Do not assume that community members are technology experts nor technologically illiterate</li> <li>• A community’s refusal to collaboration is a risk that must be accepted</li> </ul>
	9. <i>The need to combine research goals, funding, and concrete social political dynamics</i>	<ul style="list-style-type: none"> <li>• Designers and researchers as intermediaries between the interests of the different actors involved (project beneficiaries, investors, funding agencies, and other stakeholders’ goals)</li> </ul>

Table 1: Summary of guiding principles for developing PD-inspired NLP tools

Use-before-use addresses the common practice to build an image of the use of a product by people *before* use actually take place. The methods employed to favor people determination of use-before-use (e.g., workshops, design games, fictional scenarios, and prototyping) can become part of forms of *participation washing* (Sloan et al., 2020), that is the use of methods belonging to PD in processes in which participants do not have a significant influence on the outcome. When done properly, the keys in PD process are the articulation of transformative visions (Huybrechts et al., 2020), the ethnographic approach to design (Blomberg and Karasti, 2012b), and the reflexive discussion on the position of designers, communities, and institutions (Lyle et al., 2018; Teli et al., 2020).

Design-after-design addresses the possibility of people’s manipulation of “finished” products. Design-after-design needs to be investigated and favored through concepts like *infrastructuring* (Karasti, 2014) or by looking at the connections between specific digital artifacts and wider artifacts ecologies (Bødker and Klokmoose, 2012).

### 3 Communities are often not determined *a priori*

The last 20 years have seen a change in the subjects involved in PD, with the notion of *community* becoming one of the most relevant to describe the participants to PD projects (Dittrich et al., 2002; DiSalvo et al., 2012; Light and Miskelly, 2019).

The notion of community is complex and multifaceted. Long lasting criteria such as the sharing a place, an interest, or a condition have proven to be limited (Mosconi et al., 2017; Thinyane et al., 2018; Cibin et al., 2019; Teli et al., 2020). This paper defines a community as the presence of dense social relations and of, at least, an element - being it geography, interests, specific conditions, or structural position in society in terms of power - tying together its members. Each of these dimensions represents a challenge to current practices of design and realization of NLP systems.

Although the definition of community recalls an idea of a *unitary whole*, the ensemble of the participants to a project is not always completely determined *a priori* but it could get formed *within and through* the design process (Le Dantec and DiSalvo, 2013), which current sampling methods in NLP mostly fail to capture.

A consolidated tendency is to look at PD practices in terms of empowerment of marginalized groups (Ertner et al., 2010; Racadio et al., 2014). Their adoption and integration in the NLP pipeline can help to address underexposure of both language varieties and linguistic phenomena.

*Mario is a scholar in Human-Computer Interaction and technology design. He works on a project to support the development of community radio stations by rural and isolated communities. One of the communities involved belongs to a village of about 600 inhabitants located between a river delta and the Black Sea in Romania. The inhabitants are mainly descendants of a group of Ukrainian Cossacks who immigrated there in the 18th century. In addition to speaking Romanian the residents speak a Ukrainian dialect. Together with a Romanian NGO specialising in human rights and media democracy, Mario works to involve the inhabitants as volunteers to run the radio station and create content for the programs. However, the Romanian broadcasting license obliges stations to transmit 24 hours a day, and the volunteers struggle to create enough content. Mario proposes to use a new and advanced natural language generation system, GPT-3, to generate content. Besides the fact that the machine does not “speak” the community’s dialect and requires English translations, GPT-3 produces output with prejudices and negative stereotypes against the community.*

#### 4 Data and communities are not separate things

As we saw in the first three points, communities represent the core element of PD. One might expect that communities have a prominent role in the development of NLP systems. Indeed, communities are the producers of the oil that runs NLP research: language data.

We observe, however, that this is not the case. Searching for the term “community” in the ACL Anthology<sup>2</sup> returns 100 papers. However, by manually inspecting each of them, we discovered that only 9 present some sort of engagement with a community of speakers (Garcia et al., 2008; Levin, 2009; Bird et al., 2014; Everson et al., 2019; Kemp-ton, 2017; Susarla and Challa, 2019; Conforti et al.,

<sup>2</sup>Accessed on April 30th, 2021

2020; Griscom, 2020; Le Ferrand et al., 2020). These works target endangered languages and propose technological solutions to an array of problems (e.g., archiving, documenting, or tooling). None of them presents an active and direct involvement of the communities *in the design process* of the suggested NLP solution. As pointed out by Bird (2020), people agency is absent and language is seen as data to be dug.

Compliance with PD methods requires for NLP to become more aware of the relationship between language data and the speakers who first produced. In this context, we advocate for a shift of paradigm, from *language as data* to *language as people*.

Mario’s story exemplifies the danger of forgetting the link between NLP training data and its underlying producers: by not asking himself whether the language varieties behind GPT-3 are representative of the community he is trying to help, he ends up hurting it. The application of PD methods is a viable solution to overcome part of this predicament. The next principles will address two key steps of the development of NLP systems: data collection & annotation, and evaluation & deployment.

#### 5 Community involvement is not scraping

The training of current SOTA language models (LMs) is based on large amounts of written text crawled from the Web, with no or little documentation (Bender et al., 2021). However, the attempt to calibrate a tool to the needs of a specific community demands *concrete* social interactions. This requires the development of ethical engagement practices based on respect, equity, and reciprocity to gain the trust of the gatekeepers of the community (Le Dantec and Fox, 2015; Hirmer et al., 2021; Bird, 2020). Gaining trust of communities is fundamental, especially when dealing with small groups of people. In that case, all information is sensitive and often considered a currency that can be devalued once made public (Giglietto, 2017).

Innovative, flexible and transparent approaches to data collection and annotation should be put in practice. In line with PD methods, the way this cannot be reduced to a check-list valid for each and every community: context-specificity, which affects participation practices, cannot be avoided (Sloan et al., 2020). Documenting, describing, explaining, and showing how the data a community makes available is processed by and used to create an NLP

system is an essential step. It is up to the NLP researchers to gain trust by describing as best as they can the purpose of the work and the risks and benefits for the community. Additional advantages of designing NLP systems around the needs of a community are the possibilities of challenging existing power dynamics and also reduce risks of dual use. In this context, initiatives such as the *Feminist.AI*<sup>3</sup> collective and *Indigenous data sovereignty practices* (Kukutai and Taylor, 2016; Walter and Suina, 2019) are positive and innovative examples.

## 6 Never stop designing

Mario’s scenario is a good example of a bottleneck in the deployment of NLP systems: in most cases, they will not fit the needs of a community and adapting them is a challenging task.

The adoption of Machine Learning techniques for developing NLP systems adopts a vision where statistical generalizations can be learned and applied to broader contexts (Sloan et al., 2020). Datasets are assumed to be good samples of language phenomena, but are actually deeply context-bound at different levels (e.g., time period, medium, population sample, among others). It is known that NLP tools struggle with tail phenomena (Ettinger et al., 2017) and are subject to bias (Bender and Friedman, 2018). Solutions are varied and focused on areas such as Domain Adaptation and Transfer Learning (Blitzer et al., 2006; Daumé III, 2007; Ma et al., 2014; Ganin and Lempitsky, 2015; Wu and Huang, 2016; Ruder et al., 2017; Ruder and Plank, 2017; Ramponi and Plank, 2020) or de-biasing (Gonen and Goldberg, 2019; Paul Panenghat et al., 2020; Liang et al., 2020; Zhou et al., 2021).

A PD-aware NLP tool should foresee this *community adaptation* feature at its design stage. This requires to overcome technical (i.e., access or manipulation of the code) and resource (financial and human) predicaments as well as the use of predatory practices of users’ involvement (i.e., recognize participation as labor). Having access to continuous and updated feedback from a community is paramount for ensuring that tool adaptation effectively addresses their evolving needs. In this context, researchers should put in place appropriate socio-technical solutions considering the peculiarities of the community (e.g., developing an API to report bugs might not be appropriate in areas

<sup>3</sup><https://share.hek.ch/en/participatory-ai-how-to-make-better-ai/>

with limited internet connection). This open-ended evaluation process challenges existing industrial paradigm based on the idea of scaling.

*Katie is a PhD candidate in Interaction Design working on a project on compliance to labor norms. She engages relatively small trade unions in understanding how the unions can communicate widely and effectively to the public, and to the large population of prospective new members. She has collected a variety of information, through interviews and workshops. During these activities, she has encountered two main challenges for her research: (i) she collected a large amount of textual data about labor conditions and used out-of-the-shelf NLP tools to run sentiment analysis on it; however, the tools provide predictions only in an aggregated, uninterpretable form, which prevents Katie from providing the unions with specific insights. She has also applied for funding to improve the tools’ interpretability but her request has been conditionally accepted subject to changes in her research topic; (ii) although she is mindful of her role as a researcher, Katie has faced frictions when engaging with the unions as some of their members feel overtly exposed when sharing their experiences.*

## 7 Text is a means rather than an end

Introducing PD methods in the design of NLP tools promotes and embraces a philosophical perspective on the interactions between humans and machines, and of Artificial Intelligence in general, as a problem-solving tool rather than as an adaptive mechanism mimicking human abilities (Winograd, 1997; Auernhammer, 2020). On the contrary, current trends in NLP are more oriented towards a rationalist perspective, attempting to develop *intelligent* systems that *understand* language (Bender and Koller, 2020).

This follows a logic of automation that attempts to ultimately remove human intervention (Crawford, 2021), reinforcing a vision of *language as data*. Language, however, is not a uniform entity but it adapts to the context where it is used. NLP systems have the potential to support the flow of meanings between contexts but in order to do so, and act as means rather than ends (Auger et al., 2017; Hanna et al., 2017), they must contend with the structural solidity of the categories on which its

algorithms are built (Bender et al., 2021). The tools Katie uses are unable to offer insightful information to her respondents because the output is uninterpretable (i.e., why a message has been labeled in such a way?). To see NLP technologies aligned with participatory methods and tasks demands a shift in the conceptualization of the outputs, or products, of NLP systems. The linguistic output of NLP systems should be material that triggers iterations or refinements to serve people's needs rather than imitate people's production of language.

## 8 The thin red line between consent and intrusion

Katie's scenario highlights how common it is to take for granted that the community always wants to be helped authorizing researchers to use any tool. Refusing collaboration is a risk that must be accepted thus preventing or interrupting the development of a proposed technical solution.

Importantly, the community's consent can be considered authentic only if it was preceded by appropriate communication. When introducing a technology or a tool to a community, researchers must avoid two unethical approaches. On one hand, using terminology with which a community is not familiar with might confuse more than explain, thus potentially resulting in uninformed consent (Tekola et al., 2009). Note, however, that researchers might also find themselves in the opposite situation. When approaching (small) communities, researchers can be misled by what is called a *deficit model* (Irwin and Wynne, 1996), i.e., taking for granted that the reference community whom one is going to collaborate with lacks of knowledge regarding science and technology. However, people are constantly immersed in an ecology of technologies (Bødker and Klokmoose, 2012) and practical knowledge to which they refer when called upon to understand something new.

To avoid misunderstandings, one must offer transparent information about the actions that will be carried out, making use of metaphors and comparisons with existing artifacts, even if the complexity of the technological architecture represent a communication challenge (Bratteteig and Verne, 2018). And always keep in mind that this dialogue can steer people's eyes in the wrong place.

## 9 The need to combine research goals, funding, and concrete social political dynamics

All the cases observed highlight how a community-based collaboration between NLP and PD is an issue where multiple dimensions continuously interact. In addition to this, Katie's fiction introduces an additional challenge: the need to obtain external funding to conduct her research and the interests (and requests) of the funding providers/agencies.

These dynamics must take into account the goals of the researchers/designers, and of the communities involved, which cannot be completely overturned by the founders. It is evident that in this context the role of the designer/researcher becomes more and more that of an intermediary capable of translating and holding together the interests of the different stakeholders involved, without risking being co-opted and involved only in a token way (Cibin et al., 2020; Teli et al., 2020).

## Acknowledgments

We thank the anonymous reviewers for their effort in reviewing this paper, their constructive feedback and suggestions. CC is grateful to the RuralSenses team for sharing experiences on participatory practices in sustainable project design.

## References

- Lasha Abzianidze and Johan Bos. 2019. [Thirty musts for meaning banking](#). In *Proceedings of the First International Workshop on Designing Meaning Representations*, pages 15–27, Florence, Italy. Association for Computational Linguistics.
- Jan Auernhammer. 2020. [Human-centered ai: The role of human-centered design research in the development of ai](#). In *Synergy - DRS International Conference 2020*, Online. Design Research Society.
- James Auger, Julian Hanna, and Enrique Encinas. 2017. *Reconstrained design: Confronting oblique design constraints*. *Nordes*, 7(1).
- Emily M Bender. 2013. Linguistic fundamentals for natural language processing: 100 essentials from morphology and syntax. *Synthesis lectures on human language technologies*, 6(3):1–184.
- Emily M. Bender and Batya Friedman. 2018. [Data statements for natural language processing: Toward mitigating system bias and enabling better science](#). *Transactions of the Association for Computational Linguistics*, 6:587–604.

- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the dangers of stochastic parrots: Can language models be too big?](#) In *FACCT '21: 2021 ACM Conference on Fairness, Accountability, and Transparency, Virtual Event / Toronto, Canada, March 3-10, 2021*, pages 610–623. ACM.
- Emily M. Bender and Alexander Koller. 2020. [Climbing towards NLU: On meaning, form, and understanding in the age of data.](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5185–5198, Online. Association for Computational Linguistics.
- Steven Bird. 2020. [Decolonising speech and language technology.](#) In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3504–3519, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Steven Bird, Lauren Gawne, Katie Gelbart, and Isaac McAlister. 2014. Collecting bilingual audio in remote indigenous communities. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 1015–1024.
- John Blitzer, Ryan McDonald, and Fernando Pereira. 2006. Domain adaptation with structural correspondence learning. In *Proceedings of the 2006 conference on empirical methods in natural language processing*, pages 120–128. Association for Computational Linguistics.
- Jeanette Blomberg and Helena Karasti. 2012a. Ethnography: Positioning ethnography within participatory design. In Jesper Simonsen and Toni Robertson, editors, *Routledge International Handbook of Participatory Design*, pages 106–136. Routledge, New York, NY.
- Jeanette Blomberg and Helena Karasti. 2012b. Positioning ethnography within participatory design. In Jesper Simonsen and Toni Robertson, editors, *Routledge international handbook of participatory design*, pages 86–116. Routledge. Publisher: Routledge London.
- Tone Bratteteig, Keld Bødker, Yvonne Dittrich, Preben Holst Mogensen, and Jesper Simonsen. 2012. Organising principles and general guidelines for Participatory Design Projects. In Jesper Simonsen and Toni Robertson, editors, *Routledge handbook of participatory design*, page 117. Routledge.
- Tone Bratteteig and Guri Verne. 2018. Does ai make pd obsolete? exploring challenges from artificial intelligence to participatory design. In *Proceedings of the 15th Participatory Design Conference: Short Papers, Situated Actions, Workshops and Tutorial-Volume 2*, pages 1–5.
- Ulrich Briefs, Claudio Ciborra, and Leslie Schneider. 1983. *Systems Design For, With, and by the Users: Proceedings of the Ifip Wg 9.1 Working Conference on Systems Design For, With, and by the Users, Riva Del Sole, Italy, 20-24 September 1982*. North Holland.
- S. Bødker. 2000. [Scenarios in user-centred design—setting the stage for reflection and action.](#) *Interacting with Computers*, 13(1):61–75. Publisher: Oxford Academic.
- Susanne Bødker and Clemens Nylandsted Klokmoose. 2012. [Dynamics in artifact ecologies.](#) In *Proceedings of the 7th Nordic Conference on Human-Computer Interaction: Making Sense Through Design, NordiCHI '12*, pages 448–457, Copenhagen, Denmark. Association for Computing Machinery.
- Michel Callon. 2004. The role of hybrid communities and socio-technical arrangements in the participatory design. *Journal of the center for information studies*, 5(3):3–10.
- Roberto Cibin, Sarah Robinson, Maurizio Teli, Conor Linehan, Laura Maye, and Christopher Csíkszentmihályi. 2020. [Shaping social innovation in local communities: The contribution of intermediaries.](#) In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society, NordiCHI '20*, New York, NY, USA. Association for Computing Machinery.
- Roberto Cibin, Maurizio Teli, and Sarah Robinson. 2019. [Institutioning and community radio. a comparative perspective.](#) In *Proceedings of the 9th International Conference on Communities and Technologies - Transforming Communities, Camp;T '19*, page 143–154, New York, NY, USA. Association for Computing Machinery.
- Costanza Conforti, Stephanie Hirmer, Dai Morgan, Marco Basaldella, and Yau Ben Or. 2020. [Natural language processing for achieving sustainable development: the case of neural labelling to enhance community profiling.](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8427–8444, Online. Association for Computational Linguistics.
- Kate Crawford. 2021. *The Atlas of AI*. Yale University Press.
- Hal Daumé III. 2007. Frustratingly easy domain adaptation. *ACL 2007*, page 256.
- Carl DiSalvo, Andrew Clement, and Volkmar Pipek. 2012. Participatory design for, with, and by communities. In Jesper Simonsen and Toni Robertson, editors, *International Handbook of Participatory Design*, pages 182–209. Routledge, Oxford.
- Yvonne Dittrich, Sara Eriksén, and Christina Hansson. 2002. [PD in the Wild; Evolving Practices of Design in Use.](#) In *Participatory Design Conference*.

- Pelle Ehn. 1992. Scandinavian design: On participation and skill. In Paul S. Adler and Terry A. Winograd, editors, *Usability*, pages 96–132. Oxford University Press, Inc.
- Pelle Ehn. 2008. Participation in design things. In *Proceedings of the Tenth Anniversary Conference on Participatory Design 2008*, PDC '08, pages 92–101, USA. Indiana University.
- Pelle Ehn. 2016. [Design, Democracy and Work: Exploring the Scandinavian Participatory Design Tradition](#). *Critical Design / Critical Futures*.
- Pelle Ehn, Bengt Brattgård, E. Dalholm, R. C. Davies, Ann Hägerfors, Birgitta Mitchell, and Jörn Nilsson. 1996. The envisionment workshop—from visions to practice. In *Proceedings of the Participatory Design conference*, pages 141–152. MIT Boston.
- Enrique Encinas and Mark Blythe. 2018. Research fiction and thought experiments in design. *Foundations and Trends in Human-Computer Interaction*, 12(1):1–105.
- Marie Ertner, Anne Mie Kragelund, and Lone Malmberg. 2010. [Five Enunciations of Empowerment in Participatory Design](#). In *Proceedings of the 11th Biennial Participatory Design Conference*, PDC '10, pages 191–194, New York, NY, USA. ACM.
- Allyson Ettinger, Sudha Rao, Hal Daumé III, and Emily M Bender. 2017. Towards linguistically generalizable nlp systems: A workshop and shared task. In *Proceedings of the First Workshop on Building Linguistically Generalizable NLP Systems*, pages 1–10.
- Rebecca Everson, Wolf Honoré, and Scott Grimm. 2019. An online platform for community-based language description and documentation. In *Proceedings of the Workshop on Computational Methods for Endangered Languages*, volume 1.
- Tony Fry. 2017. [Design after design](#). *Design Philosophy Papers*, 15(2):99–102. Publisher: Routledge eprint: <https://doi.org/10.1080/14487136.2017.1392093>.
- Yaroslav Ganin and Victor Lempitsky. 2015. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189. PMLR.
- Krissanne Kaye Garcia, Ma Angelica Lumain, Jose Antonio Wong, Jhovee Gerard Yap, and Charibeth Cheng. 2008. Natural language database interface for the community based monitoring system. In *Proceedings of the 22nd Pacific Asia Conference on Language, Information and Computation*, pages 384–390.
- Danilo Giglitto. 2017. Community empowerment through the management of intangible cultural heritage in the isle of jura, scotland. *Imperial Journal of Interdisciplinary Research*, 3(5):567–578.
- Hila Gonen and Yoav Goldberg. 2019. [Lipstick on a pig: Debiasing methods cover up systematic gender biases in word embeddings but do not remove them](#). In *Proceedings of the 2019 Workshop on Widening NLP*, pages 60–63, Florence, Italy. Association for Computational Linguistics.
- Richard Griscom. 2020. Mobilizing metadata: Open data kit (odk) for language resource development in east africa. In *Proceedings of the first workshop on Resources for African Indigenous Languages*, pages 31–35.
- UTOPIA Project Group. 1981. The UTOPIA Project. On Training, Technology and Products Viewed from the Quality of Work Perspective.
- Julian Hanna, James Auger, and Enrique Encinas. 2017. [Reconstrained design: a manifesto](#).
- Ingi Helgason, Michael Smyth, Enrique Encinas, and Ivica Mitrović. 2020. Speculative and critical design in education: Practice and perspectives. In *Companion Publication of the 2020 ACM Designing Interactive Systems Conference*, pages 385–388.
- Stephanie Hirmer, Alycia Leonard, Josephine Tumweise, and Costanza Conforti. 2021. [Building representative corpora from illiterate communities: A review of challenges and mitigation strategies for developing countries](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2176–2189, Online. Association for Computational Linguistics.
- Dirk Hovy and Shannon L. Spruit. 2016. [The social impact of natural language processing](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 591–598, Berlin, Germany. Association for Computational Linguistics.
- Liesbeth Huybrechts, Maurizio Teli, Mela Zuljevic, and Mela Bettega. 2020. [Visions that change. Articulating the politics of participatory design](#). *CoDesign*, 16(1):3–16.
- Alan Irwin and Brian Wynne, editors. 1996. [Misunderstanding Science?: The Public Reconstruction of Science and Technology](#). Cambridge University Press, Cambridge.
- Gopinaath Kannabiran and Susanne Bødker. 2020. [Prototypes as Objects of Desire](#). In *Proceedings of the 2020 ACM Designing Interactive Systems Conference*, pages 1619–1631. Association for Computing Machinery, New York, NY, USA.
- Helena Karasti. 2014. [Infrastructuring in participatory design](#). In *Proceedings of the 13th Participatory Design Conference: Research Papers - Volume 1*, PDC '14, pages 141–150, Windhoek, Namibia. Association for Computing Machinery.

- Timothy Kempton. 2017. Cross-language forced alignment to assist community-based linguistics for low resource languages. In *Proceedings of the 2nd Workshop on the Use of Computational Methods in the Study of Endangered Languages*, pages 165–169.
- Henrik Korsaard, Clemens Nylandsted Klokmose, and Susanne Bødker. 2016. **Computational alternatives in participatory design: putting the t back in socio-technical research**. In *Proceedings of the 14th Participatory Design Conference: Full papers - Volume 1*, PDC '16, pages 71–79, New York, NY, USA. Association for Computing Machinery.
- Tahu Kukutai and John Taylor. 2016. *Indigenous data sovereignty: Toward an agenda*. Anu Press.
- Christopher A. Le Dantec and Carl DiSalvo. 2013. **Infrastructuring and the formation of publics in participatory design**. *Social Studies of Science*, 43(2):241–264.
- Christopher A. Le Dantec and Sarah Fox. 2015. **Strangers at the Gate: Gaining Access, Building Rapport, and Co-Constructing Community-Based Research**. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, CSCW '15, pages 1348–1358, Vancouver, BC, Canada. Association for Computing Machinery.
- Eric Le Ferrand, Steven Bird, and Laurent Besacier. 2020. **Enabling interactive transcription in an indigenous community**. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3422–3428, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Lori Levin. 2009. Adaptable, community-controlled, language technologies for language maintenance. In *13th Annual Conference of the European Association for Machine Translation*, page 8. Citeseer.
- Paul Pu Liang, Irene Mengze Li, Emily Zheng, Yao Chong Lim, Ruslan Salakhutdinov, and Louis-Philippe Morency. 2020. **Towards debiasing sentence representations**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5502–5515, Online. Association for Computational Linguistics.
- Ann Light and Clodagh Miskelly. 2019. **Platforms, Scales and Networks: Meshing a Local Sustainable Sharing Economy**. *Computer Supported Cooperative Work (CSCW)*, 28(3):591–626.
- Peter Lyle, Mariacristina Sciannamblo, and Maurizio Teli. 2018. **Fostering Commonfare. Infrastructuring Autonomous Social Collaboration**. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, pages 452:1–452:12, New York, NY, USA. ACM.
- Ji Ma, Yue Zhang, and Jingbo Zhu. 2014. Tagging the web: Building a robust web tagger with neural network. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 144–154.
- Gaia Mosconi, Matthias Korn, Christian Reuter, Peter Tolmie, Maurizio Teli, and Volkmar Pipek. 2017. **From Facebook to the Neighbourhood: Infrastructuring of Hybrid Community Engagement**. *Computer Supported Cooperative Work (CSCW)*, 26(4-6):959–1003.
- Mithun Paul Panenghat, Sandeep Sunthal, Faiz Rafique, Rebecca Sharp, and Mihai Surdeanu. 2020. **Towards the necessity for debiasing natural language inference datasets**. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 6883–6888, Marseille, France. European Language Resources Association.
- Robert Racadio, Emma J. Rose, and Beth E. Kolko. 2014. **Research at the margin: participatory design and community based participatory research**. In *Proceedings of the 13th Participatory Design Conference on Short Papers, Industry Cases, Workshop Descriptions, Doctoral Consortium papers, and Keynote abstracts - PDC '14 - volume 2*, pages 49–52, Windhoek, Namibia. ACM Press.
- Alan Ramponi and Barbara Plank. 2020. **Neural unsupervised domain adaptation in NLP—A survey**. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6838–6855, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Sebastian Ruder, Parsa Ghaffari, and John G Breslin. 2017. Data selection strategies for multi-domain sentiment analysis. *arXiv preprint arXiv:1702.02426*.
- Sebastian Ruder and Barbara Plank. 2017. **Learning to select data for transfer learning with bayesian optimization**. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 372–382, Copenhagen, Denmark. Association for Computational Linguistics.
- Jesper Simonsen and Toni Robertson. 2012. *Routledge International Handbook of Participatory Design*. Routledge, New York, NY.
- Mona Sloan, Emanuel Moss, Olaitan Awomolo, and Laura Forlano. 2020. Participation is not a design fix for machine learning (pp. 1–7). In *Proceedings of the International Conference on Machine Learning, Vienna, Austria*.
- Sai Susarla and Damodar Reddy Challa. 2019. **A platform for community-sourced indic knowledge processing at scale**. In *Proceedings of the 6th International Sanskrit Computational Linguistics Symposium*, pages 68–82, IIT Kharagpur, India. Association for Computational Linguistics.
- Fasil Tekola, Susan J Bull, Bobbie Farsides, Melanie J Newport, Adebawale Adeyemo, Charles N Rotimi,



- and Gail Davey. 2009. Tailoring consent to context: designing an appropriate consent process for a biomedical study in a low income setting. *PLoS Negl Trop Dis*, 3(7):e482.
- Maurizio Teli. 2015. [Computing and the Common. Hints of a new utopia in Participatory Design.](#) *Aarhus Series on Human Centered Computing*, 1(1):4.
- Maurizio Teli, Antonella De Angeli, and Maria Menéndez-Blanco. 2017. [The positioning cards: on affect, public design, and the common.](#) *AI & SOCIETY*, pages 1–8.
- Maurizio Teli, Marcus Foth, Mariacristina Sciannabulo, Irina Anastasiu, and Peter Lyle. 2020. [Tales of Institutioning and Commoning: Participatory Design Processes with a Strategic and Tactical Perspective.](#) In *Proceedings of the 16th Participatory Design Conference 2020 - Participation(s) Otherwise - Volume 1*, PDC '20, pages 159–171, Manizales, Colombia. Association for Computing Machinery.
- Mamello Thinyane, Karthik Bhat, Lauri Goldkind, and Vikram Kamath Cannanure. 2018. [Critical Participatory Design: Reflections on Engagement and Empowerment in a Case of a Community Based Organization.](#) In *Proceedings of the 15th Participatory Design Conference: Full Papers - Volume 1*, PDC '18, pages 2:1–2:10, New York, NY, USA. ACM.
- David R Traum. 2000. 20 questions on dialogue act taxonomies. *Journal of semantics*, 17(1):7–30.
- Kirsikka Vaajakallio and Tuuli Mattelmäki. 2014. [Design games in codesign: as a tool, a mindset and a structure.](#) *CoDesign*, 10(1):63–77. Publisher: Taylor & Francis Ltd.
- Maggie Walter and Michele Suina. 2019. Indigenous data, indigenous methodologies and indigenous data sovereignty. *International Journal of Social Research Methodology*, 22(3):233–243.
- Terry Winograd. 1997. From Computing Machinery to Interaction Design. In Peter Denning and Robert Metcalfe, editors, *Beyond Calculation: The Next Fifty Years of Computing*, pages 149–162. Springer-Verlag.
- Fangzhao Wu and Yongfeng Huang. 2016. Sentiment domain adaptation with multiple sources. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 301–310.
- Xuhui Zhou, Maarten Sap, Swabha Swayamdipta, Yejin Choi, and Noah Smith. 2021. [Challenges in automated debiasing for toxic language detection.](#) In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3143–3155, Online. Association for Computational Linguistics.