university of groningen

umcg

University of Groningen

Enabling the Digital Democratic Revival

Grossi, Davide; Hahn, Ulrike; Haret, Adrian; Mäs, Michael; Nitsche, Andreas; Behrens, Jan; Boehmer, Niclas; Brill, Markus; Endriss, Ulle; Grandi, Umberto

DOI:
10.48550/arXiv.2401.16863

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Early version, also known as pre-print

Publication date:
2024

Link to publication in University of Groningen/UMCG research database

Citation for published version (APA):
Grossi, D., Hahn, U., Haret, A., Mäs, M., Nitsche, A., Behrens, J., Boehmer, N., Brill, M., Endriss, U., Grandi, U., Heitzig, J., Janssens, N., Jonker, C. M., Kaijzer, M. A., Kistner, A., Lackner, M., Lieben, A., Mikhaylovskaya, A., Murukannaiah, P. K., ... Van de Putte, F. (2024). *Enabling the Digital Democratic Revival: A Research Program for Digital Democracy*. arXiv. https://doi.org/10.48550/arXiv.2401.16863

Copyright
Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment.

Take-down policy
If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): http://www.rug.nl/research/portal. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Download date: 13-09-2026

arXiv:2401.16863v1 [cs.CY] 30 Jan 2024

Enabling the Digital Democratic Revival:
A Research Program for Digital Democracy

Lorentz Center Workshop on "Algorithmic Technology for Democracy"
The list of authors follows at the end of the paper

January 31, 2024

Abstract

This white paper outlines a long-term scientific vision for the development of digital-democracy technology. We contend that if digital democracy is to meet the ambition of enabling a participatory renewal in our societies, then a comprehensive multi-methods research effort is required that could, over the years, support its development in a democratically principled, empirically and computationally informed way. The paper is co-authored by an international and interdisciplinary team of researchers and arose from the Lorentz Center Workshop on "Algorithmic Technology for Democracy" (Leiden, October 2022).

For comments, suggestions, improvements, criticisms please reach out to the coordinating authors listed at the end of the paper. We plan to update the document regularly.

Contents

1 Digital tools for a participatory democratic revival? 1

2 No digital democracy without digital democracy science 4
2.1 The state-of-the-art of digital democracy . . . 4
2.2 The core technological challenge: from democratic values to software . . . 5
2.3 The core research challenge: democracy requires scientific rigor . . . 6

3 Towards a research agenda 9
3.1 Equality in digital democracy . . . 9
3.2 The rule-tool-user nexus in digital democracy . . . 9
3.3 Citizens as users: preference elicitation in digital democracy . . . 11
3.4 Political representation and digital democracy . . . 11
3.5 Democratic deliberation and digital democracy . . . 12
3.6 Digital democracy, identities and grassroots digital communities . . . 14

4 Closing words 15

1 Digital tools for a participatory democratic revival?

Democratic crisis and participatory democratic revival Democracy is under pressure. Citizens' dissatisfaction with democracy has been running high for decades [FKS+20], and one of its causes is the perception of a growing disconnect between political decision makers and citizens [Gil05, FA12, WEATC+22, WF21]. Many citizens find themselves distrusting democratic institutions, feeling sidelined rather than represented by political elites perceived to pursue the needs of various interest groups over those of society. A chief route to counter these feelings and strengthen democracy is to better include citizens throughout the deliberation and decision-making processes behind policy decisions [Mat20, PL+20].

However, most existing approaches to civic engagement either include tiny numbers of citizens (as in citizen assemblies) or exclude active deliberation and limit participation to the support of petitions

(as in e-petitions). Furthermore, they typically fail to explicate a clear and transparent path from citizens' input to political action that is required for any democratic process [FSDB21]. In this position paper, we echo earlier calls for the deployment of digital tools to overcome the limitations of existing approaches to civic engagement. Digital tools are uniquely suited to address these shortcomings of current approaches, as they allow for the development of scalable and interactive methods deployable at various levels of governance: from local, to national and even global. However, we also identify remaining roadblocks to the application of digital tools in a democratically principled way, and propose a research agenda that directly addresses these challenges. We argue that an interdisciplinary effort integrating insights from social, economic, psychological, political and computational sciences, as well as the application of formal modelling and rigorous empirical testing, is essential to realize impactful digital democracy that can gain the trust of citizens.

Participation that counts Participation needs to be impactful in order to overcome the divide between citizens and decision makers. We use the term "citizens" here not in its strictly legal meaning, but rather to refer to all the people residing in a given community and having rights of participation towards the policy making of that community. The outcomes of citizen participation initiatives need to feed into the legislative process and shape political decisions. To be sure, citizen participation can inform political decision-making in various ways even without direct impact on policy making [SBBM17]. However, forms of participation that are lacking an ex ante discernible effect on policy will likely foster frustration, cynicism, and distrust in citizens and will worsen the divide between them and political decision makers. Even one of the most ambitious citizen participation events to date, the EU's Conference on the Future of Europe, has been criticized for a lack of "clarity on the follow-up" and "concrete recommendations to lead to legislative action or initiation of Treaty amendments" [Ngu21].

It is our view that a vicious circle of distrust between citizens and political decision makers constitutes the major roadblock to more impactful citizen participation. Digital tools such as communication platforms like Facebook and X (formerly Twitter) have allowed the public to express their concerns, but are also a breeding ground for opinion polarization, the spread of fake news and conspiracy theories, as well as troll attacks against decision makers, and hate speech. All of this fuels distrust on the side of political actors. While there is broad agreement on the importance of political participation, there is also growing reluctance to actually involve citizens digitally. Online communication is perceived as an unproductive black box, likely to generate digital outrage, personal harassment, even offering a point of attack for political competitors or subversion by foreign powers. In fact, before the launch of the Conference on the Future of Europe, 12 EU member states demanded that the conference "should not create legal obligations, nor should it duplicate or unduly interfere with the established legislative processes."[1] Alarmingly, this distrust reinforces citizens' feelings of disconnect with policy making: the reluctance of politicians to involve citizens fosters citizens' feelings of being excluded from decision-making, which in turn further demotivates politicians to involve (disgruntled) citizens.

Preconditions for impactful democratic participation We argue here that the dynamic of distrust outlined above can effectively be overcome if citizen participation is approached with the goal of meeting three key preconditions.

First, if participation is meant to translate into actual policy decisions, it is critical that participation happens according to a ***transparent*** process. How the process unfolds should be clear to participants. Steps need to be diligently documented in a privacy-respecting but traceable way, in order to be able to demonstrate that outcomes are not the result of arbitrariness or manipulation. Without proper documentation, citizens may not be able to gain trust in the process and political actors may more easily disregard the outcomes of a participatory process.

Second, it is critical that participatory processes are ***open*** and ***fair***. Every citizen able and interested in participating in decision-making should have the opportunity to do so as a peer. This includes: on the one hand, providing avenues to help citizens overcome digital divides, to transcend cultural and gender norms around participation, and to enable participation in multiple languages; on the other hand, to guarantee that the participatory process satisfies key democratic desiderata of equality, inclusivity and representativity. When implemented at national or even international levels,

[1] "Conference on the Future of Europe: Common approach amongst Austria, Czech Republic, Denmark, Estonia, Finland, Ireland, Latvia, Lithuania, Malta, the Netherlands, Slovakia and Sweden", 2021, https://www.staten-generaal.nl/eu/overig/20210325/non_paper_conference_on_the_future/document

this can imply that millions of citizens may want to contribute and that vulnerable, hard-to-reach, and marginalized groups get the voice they need. If participation has direct impact on the lives of citizens, then every citizen needs to be given the opportunity to participate.

Third, impactful participation requires participatory processes that are ***fit for purpose***. Such processes should be safe from undesired dynamics like the spreading of disinformation, opinion polarization, hate speech, and be secured against fraud. At the same time, they should be able to fully reap the benefits that the diverse insights emanating from wide participation can bring. It is important that malicious behavior by trolls, bad-faith political actors, and exogenous powers is kept at bay, while the contributions of good-faith participants are capitalized on. Democracy is already experiencing a crisis. Well-intended but not well-designed participatory processes have the potential to deepen this crisis rather than strengthen trust in democracy. The public and scholarly debate about echo chambers on online social networks, about social bots, fake news propagation, and foreign interference in democratic elections illustrates how digital communication technology can contribute to undemocratic outcomes of public debates. Participation platforms need to be developed in a way that enables them to prevent such undesired dynamics *by design*.

The untapped potential of digital democracy for impactful participation The current state of democratic participation is worrisome. While there are clear paths for companies and interest groups to lobby their ideas into democratic decision-making processes, citizens have very limited opportunities to participate in deliberation and decision-making. Compounding this problem, current attempts at enhancing democratic participation still fall short of delivering the means for impactful citizen participation. And this holds even for the most high-profile efforts, such as the Conference on the Future of Europe we mentioned above. Digital technology holds the promise to enable democratic participation that is transparent, open, and fit for purpose. However, while public authorities from all levels of government increasingly turn to digital democracy tools to improve citizens' participation [OEC20, SBBM17], the true potential of digital approaches is hardly exploited.

First, digital communication fosters transparency as it allows to document who was when contributing what to a given debate or decision and to make this information publicly available. Existing approaches to digital citizen participation are often criticized for a lack of transparency as to how the input of citizens was ultimately incorporated in the final decision-making and for including multiple filters between citizens' input and final outcome.[2]

Second, digital citizen participation fosters accessibility. Online social networks demonstrate that digital communication is possible even on national and international scales. What is more, digital platforms allow users to interact asynchronously. That is, digital communication does not require that participants meet at a given time and place, alleviating one substantial burden of participation. Existing digital citizen-participation approaches, however, often limit participation to the formulation and support of petitions and, thus, fail to support a process of genuine deliberation. Likewise, approaches that do include deliberation are often restricted to a few hundred participants usually selected at random by sortition. While this approach has proven effective in eliciting fruitful and more representative input for decision makers [VR16], its legitimacy as a tool to determine political decisions is open to objections because sortition is, in and of itself, a barrier to participation.

Third, digital citizen participation is mediated by algorithms supporting deliberation and decision making. Precisely this algorithmic layer of digital tools, once properly leveraged, makes it possible to carefully craft participatory processes that are truly fit for purpose. Existing open-source platforms such as LiquidFeedback [BKNS14] or Polis [SBE+21] (see Section 2 for an overview of the state of the art) show this is possible. In 2018, the government of Taiwan crowd-sourced legislation, using Polis, that enabled over 4000 citizens to deliberate for four weeks on how to regulate Uber's operations in Taiwan. The results of this large-scale deliberation were adopted by the country's legislators and became law.

While these attempts are encouraging, we need to move on from the current trial-and-error approach to the development of robust digital democracy technologies that shall serve us well in years to come. Errors in the current phase involve the failure of participatory processes: not something we can allow ourselves, given the already dire state of citizen participation. A more principled approach to technology development is needed. Crucially, such an approach is within reach. Scientific disciplines as diverse as philosophy, economics, computer science, political science, sociology, and social psychology

[2] Criticisms of precisely this nature have again been aired against the Conference for the Future of Europe [You22].

already provide a rich theoretical and empirical toolbox to tackle the systematic design of platforms that are fit for enhancing democratic participation.

2 No digital democracy without digital democracy science

For the reasons outlined above, our aim with this position paper is not to contribute to the philosophical and normative debate about what democracy is. We take a much humbler, and practice-oriented perspective and aim to contribute to the development of democratic institutions from a 'procedural' view-point. We view digital democracy as a library of computational tools that support such practices at various levels in our societies, from the local, to the national and the trans-national. The challenge is to identify the combination of tools that best meets the standards of transparency, openness, and fitness for purpose in the settings that are relevant for the institutions or grassroots organizations that intend to deploy such tools.

2.1 The state-of-the-art of digital democracy

Digital participatory tools: a short history Technology has been part of democratic practice since at least Athenian democracy. In ancient Greece, for instance, the so-called *kleroterion* was used to randomly select citizens into councils and public offices. But it was the advent of digital technology that brought the relation between technology and democracy to the forefront. Starting in the 1960s, scholars such as Gordon Tullock envisioned how emerging information technologies could be used to enable mass direct participation to the legislative process on a national level [Tul67, Tul92]. Similar perspectives have been echoed for decades by democratic theorists, from Robert Dahl [Dah08] to, more recently, Hélène Landemore [Lan20, BLR21]. Currently, we are faced with a wealth of platforms for democratic participation. These have been used for a variety of purposes: from citizens' assemblies [Par19], to participatory budgeting [MR09], to policy-drafting [RLL22], to full-fledged decision-making [SBBM17]. Prominent examples are: Polis [SBE+21],[3] which was used in 2018 by the Taiwanese government to crowdsource high-profile state-wide legislation on controversial matters such as the regulation of ride-sharing services; Consul,[4] which was developed by the Municipality of Madrid and has been used by several cities worldwide and received the United Nations Public Service Award; YourPriorities,[5] which has been used for crowd-legislation and policy-making in Iceland since 2010; Decidim,[6] which has been used by the city councils of Barcelona and Helsinki; LiquidFeedback [BKNS14],[7] which was developed in the early 2010s and has been used by various political and civic organizations across Europe since. Many more platforms exist, as demonstrated by the online databases of Democracy Technologies[8] and The Democracy Foundation[9]. Beyond the anecdotal successes in the deployment of participatory tools such as the ones above, evidence is slowly accumulating that technology-mediated citizen participation can improve perceived legitimacy and inclusiveness, as well as the quality of policies and government decision-making [Fun15]. It is high time to fully tap into this potential.

The Conference on the Future of Europe The most ambitious digital participation initiative in Europe so far has arguably been the "Conference on the Future of Europe", an event organized by the European Union that reached, digitally, more than 650,000 European citizens and generated a joint declaration that was endorsed by the European Parliament (resolution 2022/2648(RSP)).[10] One of the central instruments was the software *Decidim*, which allowed citizens to contribute in the 24 official EU languages and supported more than 52,000 active participants to debates and discussions. These contributions informed four citizen panels consisting of 200 randomly selected European volunteers. One of these panels was concerned with the "perceived distance between the people and their elected representations" [Eur22, p. 20], the same observation that motivates the present position paper. Recommendations formulated by these and other panels were repeatedly debated by a plenary

[3] https://pol.is
[4] https://consulproject.org
[5] https://citizens.is/your-priorities-features-overview
[6] https://decidim.org
[7] https://liquidfeedback.com
[8] https://democracy-technologies.org/database/
[9] https://democracy.foundation/similar-projects/
[10] https://www.europarl.europa.eu/resources/library/media/20220509RES29121/20220509RES29121.pdf

composed of citizens and politicians. The outcome of this process was a final report listing 49 proposals for reformation of regional, national, and European policy. One of the central proposals developed by this large-scale citizen participation was that citizens demanded regular direct participation on regional, national, and the European levels. They also explicitly requested digital solutions. Similar proposals were formulated in the Charter on Youth and Democracy, another European participation initiative.[11] In this charter, young Europeans demand more direct participation in regional, national, and international legislative processes and explicitly call for digital tools such as e-voting.

However, despite the great ambition and the impressive deployment of digital tools, the Conference on the Future of Europe has also been criticized. Most importantly, it remained unclear what the impact of the conference was. Citizens formulated their demands and the European Parliament endorsed the final report, but no clear legislative response followed. Markedly, the three aforementioned conditions of transparency, openness, and fitness for purpose were not met. There was a notable lack of transparency in the process transforming citizens' input into the final set of proposals. Second, the most significant part of the deliberative process and the actual decision-making were restricted to a very small group of citizens (panels and the plenum). For most citizens, participation was limited to writing and endorsing comments. Third, it is unclear to what degree the participation process was actually fit for purpose. Can we consider the outcome to really represent the demands of European citizens? Have the main issues citizens consider important been covered? To what degree has the outcome been shaped by the politicians and experts who have been involved in the process? These weaknesses limit the impact of this ambitious and visionary citizen participation process, precisely because it is all too easy to disregard its outcomes as failing to truly reflect the interests of citizens.

In a nutshell, the Conference on the Future of Europe serves as an ideal-typical illustration of the currently prevailing approach to digital participation. On the one hand, the conference demonstrates that large-scale digital participation is technically possible and that there is a strong demand on the side of citizens. On the other hand, it also illustrates that designing digital participation platforms is challenging and that problematic design decisions may generate frustration and potentially even intensify the sidelining of citizens.

2.2 The core technological challenge: from democratic values to software

Democratic values in digital democracy We documented above a strong demand for more citizen participation in general, and for digital means to enable such participation in particular. Any digital tool for democratic decision-making and participation operates in a context that is necessarily value-laden. There are general fundamental values that both conceptually and legally define democracies (participation, autonomy, rule of law) and fundamental rights and goods (free speech, various freedoms from intervention and harm, protection of minorities, as captured, for example, in fundamental human rights). These values, in practice, are subject to trade-offs, both at the abstract level and in the individual case (e.g., effectively protecting individuals from harm may require prohibiting certain kinds of speech). Those tensions and trade-offs are fundamentally not ours to resolve, last but not least because they will be resolved differently across countries and across contexts. At the same time, it is important to acknowledge that technology is not value-neutral [VdPR11], and this is even more so the case in the democratic context we want to address.

From values to software Concrete system design decisions will impact features such as effective participation and will likely do so in multiple complex ways. Importantly, designers will also have to be vigilant to prevent discriminatory biases and include the necessary safeguards. This means that we must necessarily adopt an *instrumental perspective*: *If* we want democratic tools to satisfy (or even promote) a particular system behavior, *that* is how the software should be designed. Even with that framing, there will be multiple limitations. Once understood what features promote impactful participation, there will still remain instrumental questions about how manipulating those features in a digital system promote the wider democratic goals. Such a question may sound simple, but the answer will likely be the result of multiple independent, or semi-independent components, design features and their interactions. Coming to understand these components and design features, and how they jointly produce particular outcomes is what we believe is currently missing for the development of fit-for-purpose digital democracy solutions. *Understanding how specific design choices affect the democratic*

[11] https://cor.europa.eu/en/our-work/Pages/charter-youth-democracy.aspx

quality of digital democracy systems is the focus of the research efforts we are promoting.

Digital democracy as a complex socio-technical system We see it as a likely consequence of the nature of the type of systems under study—complex systems of multiple interacting individual agents with the possibility of feedback dynamics—that there will be no single unique answer of how to design for effective participation. Instead, the likely outcome of research in this area will involve formulating a space of possible designs along with evidence-based guidance on what designs might be more or less appropriate with respect to particular value-based choices in a given context. Despite the huge interest and demand in digital democracy, the interdisciplinary research endeavour aimed at these tools is presently still rather far from such a deep instrumental understanding of the design space of tools. Progressing toward that goal, given its intrinsic complexity, will consequently have to resort to a rich combination of methods.

2.3 The core research challenge: democracy requires scientific rigor

As argued in the previous section, design decisions have an impact on whether and to what degree a given software is able to realize democratic values. Thus, it is critical that such decisions are made explicit and that they are explained. That is, developers need to be able to demonstrate that a given design contributes to reaching a given democratic value. This level of confidence in digital tools cannot possibly be achieved by trial-and-error. It requires a valid theoretical argument about why a given design decision has a desired consequence, as well as the empirical research necessary to empirically support the theoretical argument.

Fortunately, various scientific disciplines provide a rich toolbox of theoretical models, formal methods, and empirical approaches that have proven effective in rigorously developing scientific understanding of complex social systems and computational artifacts. We argue that these tools need to be applied to the development of digital democracy in order to rigorously anticipate the consequences of design decisions and to best implement democratic values in software. Below, we outline the three methodological pillars of the scientific approach to digital democracy we argue for in this position paper: modeling, behavioral experiments, observational studies.

Why formal and computational models? First, we argue that mathematical and computational models are key for the development of digital tools for digital democratic participation. Rather than developing theoretical arguments based on intuition, formal approaches use logic, mathematics, and computational methods to develop valid theory. In particular when the system under investigation is complex, formal modeling allows one to test the logical validity of one's intuition and to rigorously demonstrate that a given theoretical argument is valid [LMXF22, Sch06, Eps08]. This theoretical transparency is a value in and of itself, especially in the context of the development of democratic tools. Furthermore, mathematical analysis allows one to explore in detail the extent to which results rest on any one of their assumptions. In this way, formal methods allow one to move beyond intuition to a fine-grained understanding of a problem space. This is especially useful for the analysis of the algorithms powering digital democracy systems [BCE+16]. But besides these general methodological points, there are three more concrete functions of formal methods that specifically match our purposes.

First, considering that democracy is already under pressure, we deem it critical to deploy digital tools only when one is confident that they satisfy core democratic values and do not generate undesired effects. The public and scholarly debates about filter bubbles, social bots, and fake news on online social networks illustrate that seemingly innocent design decisions in digital tools can have highly unexpected yet impactful consequences. These issues will likely be amplified further by the ever faster inclusion of artificial intelligence components in more and more areas of life and the fact that the behavior of existing AI systems is typically not well understood, let alone formally verifiable. Formal methods that are carefully tailored to the application contexts they represent, on the other hand, enable one to demonstrate mathematically how envisioned design satisfies democratic criteria, such as fairness, equity, and representation. Furthermore, models that are carefully calibrated to the application domain they represent make it possible to gain insights into digital democracy tools before they are actually deployed in order to anticipate undesired effects. In other words, formal methods make it possible to demonstrate to citizens and decision makers that digital tools meet the democratic desiderata expected from them even before they are actually deployed.

Second, approaches to digital democracy that rely exclusively on real-life experimentation can be fundamentally limited. A central problem is that experiments with digital tools may fail and further damage citizens' trust in democratic institutions, in particular when implemented on a large scale. Unfortunately, there are also limits to the gradual up-scaling of digital participation. As a consequence of this, it is certainly wise to experiment with digital tools only on a small scale. However, "more is different" [And72], a notion that can apply to any system that consists of micro-entities exerting influence on each other. Even within the field of physics, for instance, Anderson famously argued that the "behavior of large and complex aggregates of elementary particles [...] is not to be understood in terms of a simple extrapolation of the properties of a few particles. Instead, at each level of complexity entirely new properties appear" [And72]. This echoes an argument in the social sciences developed already by Emile Durkheim, one of the founding fathers of sociology, who reasoned that "society is not the mere sum of individuals, but the system formed by their association represents a specific reality which has its own characteristics" [Dur16]. Thus, drawing conclusions about the dynamics in large populations based on observations made in smaller communities is inherently problematic. Computational models enable extensive experimentation at scale and provide insights into the behavior of digital democracy tools in large communities that would not be accessible via observations alone.

Third, and related to the above point, formal methods allow one to study counterfactuals and to systematically improve digital tools without experimenting in real life. A formal model that has been carefully calibrated to a specific empirical setting can function as a so-called "digital twin", a computational model that mimics all relevant aspects of the real system it is supposed to represent [TQ19]. NASA, for instance, developed digital twins of their space vehicles since they cannot be investigated once deployed in space. What is more, copies on Earth are of limited use, since they are not exposed to the harsh environment of space. Digital models of their vehicles, however allow one to simulate how they cope for instance with radiation. Thus, digital twins can be used to predict dynamics of a system without studying the real system. Moreover, aspects that cannot be changed in the real system can be manipulated in the digital twin without ethical, technological, or financial limits. Comparing the digital twin with these counter-factual siblings allows one to rigorously quantify differences between the two models. In the context of digital democracy tools, for instance, one could use a digital twin that has been carefully calibrated to deliberation dynamics in small groups and study dynamics in much bigger systems with millions of virtual users. AI systems, in particular large language models, can be of help here (e.g., [BCS+22, FGP+23]), but their use in scientific studies also poses novel challenges of explainability and reproducibility.

Why behavioral experiments? At the heart of research on human behavior is the ability to run behavioral experiments, that is, studies in which factors are *manipulated* across experimental conditions in a study with participants typically drawn by random sampling, so as to afford generalization of results to a wider population. Experimental manipulation of putatively relevant factors is what allows one to infer causal effects, whose presence or absence (and effect sizes) are typically ascertained through the use of inferential statistics.

Behavioral experiments vary with respect to how precisely individual 'factors' are isolated and how 'random' the sampling of the study sample is. At one end of the spectrum, laboratory studies allow careful stimulus control. Thus, in conjunction with judicious sampling of participants, laboratory studies allow one to eliminate confounding variables. However, this degree of control may come at the expense of so-called 'ecological validity', that is, the extent to which the laboratory task still matches the conditions experienced 'in the wild'. This may limit the transferability of results obtained under tightly controlled laboratory conditions to real world settings. At the other end of the spectrum, field studies conduct experimental manipulations 'in the wild', with limitations on experimental control and sampling. There is, arguably, no such thing as a 'perfect experiment' or a 'perfect experimental design', as a result. Rather, results from behavioral experiments are strongest where they can draw on convergence across methods.

Importantly, the digital environments of behavioral experiments can conveniently reflect real world 'digital environments' to a greater extent than is possible for many other domains of inquiry, simply because digital tools can be used to simulate the very environments of interest even where those environments are not actually present. At the same time, experimental conditions can be created to match most closely either the real world or a particular model chosen to study the system (more on this below). In this context, behavioral experiments serve to connect micro and macro behaviors of a

digital system.

A further feature of behavioral experiments that is utilized particularly within economics and behavioral economics is the possibility of incentives. Psychologists and economists disagree somewhat on the need for providing explicit incentives for participants in behavioral experiments [Rea05]. Regardless of how one comes down on the issue of whether behavioral experiments should be incentivized as a matter of course or not, incentives offer an important additional tool: incentives make decisions meaningful by including utilities, and this allows one to study preferences and the intensity of those preferences. Incentivized experiments allow researchers to measure and experimentally manipulate individuals' preferences with great precision. Specifically, when a behavior is made costly but is engaged in nevertheless, there is strong evidence that there is a motive. Furthermore, by varying those costs, one can determine how strong the motive is. In the context of building a science of digital democracy, this means, for example, that one can measure the degree to which individuals consider certain decision-making processes unfair. Yet, people are motivated by different things: financial incentives, status, reputation, altruism, higher meaning, religion, ideology, loyalty, and so forth. How much does habituation matter? Could children be socialized into lifelong democratic participation by engaging (and empowering) them early on in matters that affect their lives? Beyond families, could schools, as a form of civics education, involve children in context-appropriate decision-making and instruct them in idea formulation, deliberation, decision-making, and voting? Experimental designs in digital democracy should incorporate awareness of these factors.

One important difficulty, finally, for the context of an empirical science of digital democracy is that fully understanding human behavior in a collective may also require experiments with *groups*. Such experiments are vastly more costly, both from a logistical and a resource perspective, than experiments involving individual participants alone. In particular, what from the perspective of statistical analysis was an individual, is now replaced by a group. Given that the need for statistical power (which characterizes sample sizes required to detect a given level of empirical effect) will often require sample sizes that place a considerable financial burden on experimenters (for participant payments and incentives), theoretically desired experiments with groups may often not be realisable in practice. This is exacerbated by the possibility of 'more is different' phenomena highlighted above, whereby studying a small to medium-sized group may not all translate readily to the dynamics of larger groups.

Coupling simulations with experimental studies to derive suitable 'digital twins' is one way of addressing these challenges. Another path forward may lie in the future creation of shared large scale resources, a kind of CERN for the study of information environments, that would involve the creation of online platforms that are available for experimental testing [LSG+20, WS22]. Yet another alternative is to draw on the complementary strengths of non-experimental, observational, methods.

Why observational studies? Observational studies lack the defining features of experiments: an experimental manipulation. This makes them less suited to identifying causal effects, although statistical methods exist to aid causal inference from observational data [Bre22]. This limitation is compensated for by the richness that qualitative data can provide, including its typically greater external validity: observing a behavior in a real-world context does not suffer the same problems of distortion that the development of a tightly controlled experimental stimulus may entail. Consequently, observational studies may help identify and understand limitations of behavioral experiments and simulations via computational models.

Observational studies, too, come in many forms and may be either qualitative or quantitative. Among the latter, the analysis of social media data is a natural element for the study of digital democracy. Here, the burgeoning field of computational social science [FMK22] has provided many studies seeking to understand how and why information spreads in online communities (see e.g., [VRA18], but see also [BCH21] on the methodological difficulties). Furthermore, such data may connect quite readily with measures used to evaluate computational simulations. At the same time, it is important to remember the limitations inherent in the samples that are available for study by researchers, although recent legislative developments such as the EU's Digital Services Act have also sought to widen the accessibility to researchers of data obtained by large social media companies [NC22].

We end by stressing that we ultimately not only need to use all the above methods, but that we should think about how they can be most effectively deployed in concert. That is an important methodological aspect of the research agenda that needs to be developed, albeit one which is not pursued further in this paper.

3 Towards a research agenda

In what follows we outline, in some detail, six research challenges that we consider particularly salient. These challenges have been defined by working groups at the Lorentz Center Workshop on "Algorithmic Technology for Democracy", which prompted this position paper. Such challenges constitute a starting point, and by no means an exhaustive list. At the same time, they exemplify well the complexities of digital democracy research that we identified above.

3.1 Equality in digital democracy

Equality may be considered one of the *raisons d'être* of democracy. In this view, democracy is not an end in itself, but rather a means for people to have equal standing in all aspects of the political decision-making process. Accordingly, the success of a process—including a digital one—in providing such equality should be one of the key measures by which it is evaluated, critiqued, and corrected. Traditional democracies have practical limitations on granting equality throughout the entire political process, hence they typically make do by providing citizens with the right to elect their representatives every few years. Some democracies provide also citizens with due process to initiate new legislation, including amendments to the constitution, as well to annul any legislation passed by their elected representatives—the parliament. As already observed by political scientists in the 60s [Mil69, Tul67], digital democracy has no such practical limitations, and therefore may open for investigation all aspects and components of the political decision process regarding whether and how they respect equality when being conducted in the digital realm. For concreteness, we list some key such components of the decision-making process below:

Equality in agenda setting. Decisions on what proposals to vote upon, and in what order, may have tremendous impact on the final outcome of the political process. *How can the power to set the agenda and to make proposals be shared equally among all participants?* [BSST21]

Equality in deliberation and coalition formation. Similarly, processes that precede a binding vote can also have tremendous impact on the final outcome. *How can deliberative processes*, which may include negotiations on amendments to proposals to be voted upon and the formation of coalitions behind proposals or behind amendments to existing proposals, *be structured in an egalitarian way, giving all participants an equal footing in the process*? [EGST21]

Equality in decision making. How can we ensure one person – one vote in the digital realm? Fake and duplicate digital identities (aka Sybils), coercion, bribery, and fraud are all possible also in the physical realm, but their risk is amplified in the digital realm [MTSS22]. In this context, a crucial issue for the development of digital democracy is to understand how digital participatory processes should interact with standard institutions of representative democracy. One may also ask which votes, or decisions, should be taken via direct open digital processes, and which ones via the standard representative institutions: *what should be the division of labor between digital democracy and representative democracy*?

3.2 The rule-tool-user nexus in digital democracy

Collective decision making is the main building block of democracy. In a process of collective decision making, a group of people (or *agents*) interact by exchanging arguments or by expressing opinions about some issues at stake, and they submit certain data to a mechanism, e.g., by filling in a ballot. A final decision is usually produced by *tallying*—transforming this information into a result via a voting *rule*, e.g. by aggregating ballots. In cases where the decision-making process is supported by a platform, agents act via their *user* profiles, while deliberation and tallying are carried out through a voting *tool*. The three main components of digital collective decisions—the *users*, the *rule*, and the *tool*—are not independent [TKGH23]. Their behavior relies on the existence of each other: they anticipate feedback and (should) respond to it. As in every *complex system*, governing forces are the feedback mechanisms and not the particular elements of the system in isolation. Importantly, although agents that form opinions and make decisions together within groups take part in a complex system as

well, social influence and collective decision making have so far been studied separately. By treating digital collective decisions as a complex system, we improve our understanding of their properties and capabilities, and enable further progress on the platforms where they take place.

User–user interaction Digital collective decisions are not simply the result of voting rules applied to some fixed ballots submitted by the users of a platform. Users often behave strategically, deliberate, and influence each other; their preferences that lead to the ballot formation are subject to complex processes of social dynamics, which are magnified in digital contexts and may severely bias the outcome. Online anonymity facilitates the expression of extreme opinions and sustains polarization, while social networks entail the risk of information bubbles that damage diversity of information diets [KM22, FGR16]. When collective decisions are at stake, imitation effects might increase the predictability of results due to easy consensus, or decrease it due to multistability or oscillatory behavior [TIY+13, LN18, GK18, FMF+17, BLT22]. Many relevant research questions arise: How should a digital voting process be designed to benefit from positive social-influence effects and mitigate negative ones? Are collective decisions vulnerable to cross-platform interaction of the users? How can digital platforms for collective decisions encourage information-seekers and disincentivize fake news? How do we identify and prevent vote manipulation, in groups with strong tendencies to either grow together or grow apart?

User–rule interaction As long as the goal of digital platforms for collective decisions is to support the open, equal, and transparent participation of everyone, the voting rules employed must ideally account for user heterogeneity. Users may differ in terms of their *strategic* or *collaborative* behavior, *motivation*, and their *perception of axioms*. This should be considered when trying to build rules that are hard to manipulate, and rules that provide incentives for participation. In social choice work to date, well-known impossibility results rely on strong assumptions of homogeneity [CK02], but in digital contexts assumptions about a group of users should be directly linked to the platform in which the decisions take place: for example, a platform that requires several steps for registration will naturally attract more motivated users, while a platform that allows a user to create multiple accounts may increase strategic behavior. Most importantly, the relation between a voting rule and a group of users is usually considered one-way: the users provide the input to the rule. But the other direction of the relation is also important: the rule can affect the users, steering them towards a specific type of ballots and preferences [VdSLB13] (for instance, a rule that asks for approval ballots will more likely incite binary preferences than a rule that asks for complete rankings). A rule may be perceived by a user through three lenses: the *axioms* it satisfies (i.e., reflection of the user requirements in the properties [Pro19]), the *procedure* it follows (i.e., cognitive simplicity and explainability of the procedure [BEN22]), or the *outcome* it produces. Examining closely the interaction between users and rule highlights several research gaps, on both the theoretical and the empirical front: What are the potentially heterogeneous attributes of users? What are the multiple values or forms they may take? Which of them could possibly be influenced by the voting context? How does the conceptual goal of a voting rule relate with its mathematical definition?

User–tool–rule interaction In the digital world, voting rules are implemented in the form of tools. At the least, these tools let users set up decision problems, list options, and fill in their secret ballots, before performing the tallying and reporting the results. A tool might offer several different ways for filling ballots, e.g., specifying an integer rank for each option or dragging options up and down to specify a ranking. For voting rules where an effective use of one's ballot depends on beliefs about others, the tool might help in forming these beliefs, e.g., by providing voters with voting data and letting them adjust their ballot once, several times, or as often as necessary until some deadline. Software tools can be designed to provide quite varying user experiences [HT06]. The choice of language and terminology and the overall visual design (colors, fonts, imagery, etc.) might induce framing effects and prime user behavior in certain ways, e.g., to be more cooperative or competitive, individual- or social-value oriented, honest or strategic, process- or outcome-oriented, etc. [CV19, GZY+13, Her86, DMP15]. Design details during voting or deliberation such as the order of questions, hints or nudges, or information on other users' behavior can influence users in wanted, unwanted, or unpredictable ways [KA87, Sun14, CJK+19]. Being the main interface between users and voting rules, voting tools need to be understood and designed properly, to assess and improve the impact of digital democracy. We locate a number of imperative research questions: How can or should a certain formal rule be turned

into a tool? Which of potentially several equivalent representations of its input, algorithm, and output shall be chosen? How can usage data from a tool be exploited to assess quantitative formal properties of a voting rule (such as welfare or satisfaction metrics, frequencies of axiom violations or strategic behavior, degree of engagement, etc.) and test related theoretical claims? How can data from tools implementing different rules be used to advance theory (e.g., by suggesting additional formal criteria, ballot designs, or more complex game-theoretical models of a decision rule in the context of certain interactions)?

3.3 Citizens as users: preference elicitation in digital democracy

Understanding citizens' preferences is arguably one of the primary goals of any digital democracy platform. Elicitation is the process of extracting or producing information, often in reaction to a query or question. When applied to preferences, such a process includes preference formation, since users may not be aware of their preferences before being asked. Preference elicitation is the point of contact between the users and the democratic procedures they engage with. The quality and quantity of the collected preference data—and of the democratic process following—heavily depend on such points of contact. One should not be surprised then to see a vast literature on elicitation in closely related fields such as voting theory and computational social choice, multi-criteria decision analysis, recommender systems, and automated negotiation [LGLA14, MP15, BNPS21, BG15, GPSO12].

Formalizing preference elicitation protocols Elicitation is supposed to track users' preferences on a set of proposals (or alternatives, or candidates), which may be described by features. Users can also have beliefs on other users' preferences, and may be not be aware of the existence of all proposals. An elicitation protocol is a sequence of elicitation actions that extracts information from users and represents it in a computationally-friendly format, designed towards the solution of a collective decision problem. Elicitation actions can take the form of explicit queries on preferences, e.g., asking users to specify which is their top-preferred proposal, approving a set of proposals, expressing pairwise comparisons. Other elicitation actions are: collecting preferences at the level of features of the proposals; eliciting awareness or belief about other users' preferences; asking whether given proposals share certain features; asking users about the expected benefit of a proposal.

However, a fundamental problem of preference elicitation via any form of communication is that an individual's actual preferences are mental states that are not directly measurable, and in most contexts where preferences are elicited, individuals might have incentives to misrepresent their preferences, either due to strategic interaction with other individuals' preferences via some form of collective choice mechanism, or due to social-psychological effects related to social image, expectations about desired behavior, fears of punishment, etc. As a consequence, stated preferences might differ systematically and in complex ways from actual preferences.

Research challenges The use of digital tools in democracy allows for collective choices over large and possibly complex sets of alternatives. Elicitation protocols thus need to take into consideration the *high computational and communication costs on the protocol side*, but also constrain the *cognitive load and the time effort on the users' side*. Another related issue is the handling of incomplete preferences, which may occur due to several reasons, such as the high cognitive load required for submitting complete preferences but also due to other factors such as general imperfect knowledge. Transfer and inference of preferences, either in time or between individuals, might be a viable solution here. Also, ethical considerations must be taken into account, such as biases in the elicitation protocol and in the inference of preferences.

3.4 Political representation and digital democracy

Representation is a core democratic value [Mil24, Man09, LP20, Rae67, Lij90, MANC23] but it is yet to be understood in a more detailed way to be operationalized in digital tools. Whatever form digital democracy takes, it will need to be able to credibly claim to represent the citizens it serves and to work along traditional, as well as novel, institutions of political representation.

Innovation in representation mechanisms There is no shortage of innovative proposals to change how representative bodies are selected around the world. For example, some propose to select representatives at random (a.k.a. *sortition*) [Bou20], to elect them through transitive delegations (a.k.a. *liquid democracy*) [BZ16, Val21], or to drastically increase the size of parliaments (see, e.g., `https://thirty-thousand.org`). Each proposed method has its benefits and drawbacks; however, we lack a systematic way to evaluate and compare them. Specifically, while there are numerous works in computer science and political science analyzing the strengths and weaknesses of specific methods, principled comparisons are rare.[12]

Research challenges We call for the development of a unified framework to formulate and compare innovative and traditional mechanisms for selecting representative bodies on a more principled basis. Having formulated different mechanisms within the same framework then offers the possibility to formulate various desiderata (that is, democratic values) in the same framework. While we believe that comparisons from different perspectives are possible and, in fact, urgently needed, we argue for the development of an *axiomatic* view on selection mechanisms, drawing inspiration from the rich social choice literature on voting rules. Notably, the focus of such an approach would not be on finding the "ideal" representation system. We rather envision building a navigator that maps selection mechanisms to axioms. We advocate for building a coherent picture of the advantages and disadvantages of competing selection proposals based on carefully crafted axioms to gear public debates towards *what kind of trade-offs societies are facing*, instead of continuing to argue for competing selection mechanisms on disconnected grounds. First steps on this line of research have been presented in [RBC+23].

3.5 Democratic deliberation and digital democracy

Digital democracy has the potential to involve large parts of a population or organization in rich processes of deliberation that go way beyond online petitions or voting. As contemporary online social media demonstrate, however, digital discussion does not necessarily promote democracy and can even foster discrimination, hate speech, and the spread of disinformation. Designing platforms that can truly foster democratic deliberation is a formidable challenge that requires substantial research efforts to succeed at scale.

From e-petitions to online deliberation and decision-making Digital democracy can grow to include rich dimensions of deliberative interaction: from agenda setting to opinion formation to decision-making. Many platforms aim to facilitate deliberation, while others take a voting only approach [SBBM17]. Deliberation fosters an understanding of the complex nature of a given issue, and contributes to informed decision-making. However, all too often, the opportunity to reach a decision is missed, and deliberation remains inconsequential. That's why some platforms opt for a combination of deliberation and voting. It should be noted that the result of the vote need not be a binding decision, but may also have an impact as a collective recommendation of the electorate. This requires research on how to legitimately embed credible citizen participation into the representative system. Similar to the voting systems discussed in the previous section on political representation, different voting systems can be deployed to realize decision-making on the issues at stake: from single-winner (selecting one alternative out of many) to multi-winner (selecting some alternatives out of many [FSST17]). When multiple winners are selected, quite different goals may be pursued. Lackner and Skowron [FSST17] mention three principles: 'individual excellence', 'proportionality', and 'diversity'. Voting systems and preference aggregation algorithms in general have a variety of additional properties that may make them especially suitable for use in specific contexts, but less so in others. A principled understanding of how to map the properties of these voting systems to their envisioned context of use is an ongoing challenge of great relevance for the development of digital democracy platforms that can support large scale deliberation and decision-making. These issues become even more pressing when considering decision making in richer settings such as, e.g., participatory budgeting.

[12] The few existing exceptions mostly focus on epistemic aspects, the robustness of representation, and majority agreement [AHALT22, GA15, GGIM22, AM19].

Algorithmically-supported deliberation Scaling up deliberation to large groups interacting online and asynchronously requires algorithmic solutions to several problems [VP17, BKNS22, SIM+22a, Lan22, Mik24]. We mention two such problems to illustrate the point. First, views from participants, and reactions to other participants' views, should be collected. Many examples exist. For example, the Deliberatorium [Kle11] is a web-based system that allows people to share ideas, to add supporting and attacking arguments for ideas, and to vote on these. It combines ideas from argumentation theory and social computing to help large numbers of people, distributed in space and time, combine their insights to find well-founded solutions for complex problems. Similarly, one may combine ideation and negotiation [FIK17]. For an overview, see [SIM+22b]. However, in a large deliberation, no user can possibly comment on all views expressed by others, so the problem arises of what views to present to each user (sometimes called 'opinion routing' problem [SBE+21]). Second, users need to be presented with some form of real-time overview of the state of the deliberation. Here, voting algorithms with diversity or proportionality guarantees can again play an important role as equitable mechanisms for information-processing [BKNS14, SLB+17, HKP+23]. For both the above problems the question arises of *how participants can be certain that the algorithms deployed in the process respond to democratic desiderata* [VdP13]: are all views and participants given equitable representations? Similarly, transitive proxies, i.e., liquid democracy, can be used as a mechanism to support large-scale democratic deliberations. LiquidFeedback, for example, uses transitive proxies as an empowerment for debate during the structured deliberation, to faciliate collective moderation, to determine the viable voting options, and for the final decision-making.

Modeling and measuring deliberation Work in deliberative democracy initiated in philosophy and the political sciences in the 1970s (e.g., [Raw71, Hab04]) and strengthened in the 1990s (e.g., [Dry00]). But it is in recent years that it found its way into practice by inspiring several democratic experiments such as citizens assemblies, which we already mentioned earlier [VR16, FSDB21]. Unlike for voting, however, we still lack a principled understanding of how deliberative processes should be designed [KJ97, Piv19], especially online. We even lack a basic toolbox of how to approach this kind of design problem, because we still lack suitable models of deliberative processes. Models exist that address specific features of deliberative processes, e.g.,: opinion change, consensus, meta-consensus [LLFM13, EGST21]; issue identification and clarification [LLFM13, Lis18]; decision-making improvement [Piv17, DP21]; agenda setting [REdH21]; information exchange and persuasion [GR08, CH23]. However, *no sufficiently comprehensive model exists to date that could provide a basis for design or validation efforts.* Similarly, *no consolidated method exists to date for assessing the quality of deliberation* [BDMW18]. Some metrics have been proposed both in the deliberative democracy (e.g., Discourse Quality Index [SBSS03] or the Deliberative Reasons Index [NVE+22]) and computer science literature [SCG13], but they are still either too abstract or hard to deploy in practice at scale.

Social psychology of deliberating groups Cooperation among individuals and a certain level of adherence to norms are vital for a group to work with mutual advantage of their members and, ultimately, for a society to thrive. The convention by which everybody should drive on the same side of the road is beneficial for everybody (although neither right-hand nor left-hand drive is intrinsically good or bad). Very often, spontaneous coordination emerges even in absence of written rules or external constraints [HHP17], and may lead to results that outclass individual efforts, as witnessed by many instances of the so-called *wisdom of crowds* phenomenon [Sur05]. At the same time, social sciences witness a host of examples of suboptimal or even catastrophic coordination, where failure is to be ascribed, quite to the contrary, to the presence of the group. Students may avoid asking a question in public not to be perceived as the one who did not understand, while clarifications would be beneficial for everybody [KAJ31]. A desire for acceptance, belonging, or to prove one's loyalty may equally undermine individually rational responses. Such a collective blackout, also labelled as *pluralistic ignorance*, is in fact omnipresent in daily life [JO86]. People often adapt to the behavior of other group members in a way that results in so-called *informational cascades*, where people follow each other in line with the first one heading to disaster. This type of behavior may have effects at least as big as the bursting of market bubbles [EK10].

Folklore explanations for such problematic group dynamics are that humans are inherently irrational or at least very bounded. However, formal approaches to the study of these phenomena have shown that bad informational cascades may occur even among individuals that reason in the best possible

way given the circumstances [BHW92]. The same holds as well for pluralistic ignorance [PO14] and many other detrimental group behaviors, the occurrence of which is in fact consistent with assuming that individuals are rational. Cooperation and cohesion among certain groups in diverse societies is not a given [EG18]. Political, ethnic, and racial group-based hierarchies coupled with discriminatory attitudes and a preference for excluding other groups from power leads to decreased cooperation and hampers cooperation [EG18, HHPW07].

All detrimental group phenomena mentioned above are in some measure due to a lack of respect, trust, and communication among individuals in specific situations. As a consequence, one may think that this is the cause of trouble: once barriers are down, then everything should work. Unfortunately, this is not quite the case. Indeed, collective disasters may occur also because of abundance of information exchanged among individuals. This holds for many instances of groupthink [Jan83] and group polarization [Sto61, MZ69]. In such cases, it is highly possible that initial errors get magnified by a dynamics similar to informational cascades, but this is certainly not due to lack of information circulating. In general, the *quantity* of information circulating is a relevant factor in the emergence of both good and bad opinion dynamics. Yet, how the two things are related is highly context-dependent, and easy generalizations such as "lack of exchange = bad dynamics" cannot be easily made.

Social psychology for digital deliberation design Considerations like the above are even more relevant in the context of designing tools for digital democracy. Here, the question is not simply about designing a performance-optimizing architecture: an additional issue is that the design of the platform should respect specific desiderata such as *transparency* and *openness*, which we mentioned earlier.

Furthermore, a large body of knowledge provided by psychologists, social scientists and behavioral economists needs to be taken into account and framed against a specific normative constraint for design. This holds not only for the case of openness and the quantity of information, but as well for other desiderata and dimensions. The *structure of the network* of communication in which individuals interact may heavily influence opinion dynamics [Bar16, New18]. In many cases limited connectedness among agents may even enhance the performance of a group and limit the effect of detrimental dynamics. Again, this fact needs to be taken into account when planning a mechanism where *diversity* is a normative constraint, i.e., the possibility for everyone to access alternative views. Social psychology also witnesses many instances of *overthinking* and the *paradox of choice* [Sch04], where evidence shows that choices among too many alternatives may affect the consumers' preferences and induce suboptimal choices [Che03]. In the context of discussion about digital democracy, this suggests that also the normative ideal of *representativity* has problematic aspects that need to be taken into account. In general, the message is that design guided by naïve optimism can have strong side effects and that normative ideals need to be carefully balanced against psychological realism.

3.6 Digital democracy, identities and grassroots digital communities

In this last subsection we point to a key implementation issue of digital democracy: digital identities. The more our democracies rely on digital tools, the more crucial the provision and management of the digital identities of citizens become. When one has a unique digital identifier, how much one wants to reveal about oneself might depend highly on a given context and particular goals of the digital democracy applications one engages with [Prz17]. There is, however, to this date only very little research on the very notion of digital identities and on the tools that could facilitate and maintain them in a democratic way. Therefore, the research question that this problem raises is: How could we create and manage digital identities in a democratic way? What are technological hurdles in doing so, and where can technology enable such identity creation?

The core underlying problem lies in the nature of digital identities and their proneness to being copied or even invented from scratch. If, depending on one's computational resources, one can duplicate one's opinions or votes at will, then there always needs to be some external authority or institution that validates identities, checks for duplicates and filters them out. State-supported digital identities to access public services are by now supported and used worldwide and have developed further in response to the challenges posed by the COVID-19 pandemic. Whether and how such identifiers could or should be used in the democratic setting (e.g., even just for electronic voting) is an important matter of study as the legitimacy of any decision that relies on digital identities will necessarily depend on the trustworthiness of those identities.

Alternatively to the institution-based approach to the validation of digital identities, research has focused also on how to provide trustworthy digital identities in a decentralized fashion and for grassroots digital communities: a bottom-up approach to the provision of digital identities that would seem to be impossible from the get-go. In particular, research on sybil-resilient social choice [SST19, TMLS09] and trust-based communities [PSST19, TMLS09] suggests that digital identities in grassroots communities could be based on a system of trust among community's participants. For example, existing members of a digital community could personally vouch for incoming participants, confirming their identity and the fact that the person is "real". Said approach would ensure that acceptance of new members is based on mutual trust without rules that need to be enforced by a centralized authority. This would, in turn, facilitate the self-organization of communities [PSST21] and support egalitarian decision making including decisions concerning community management [SST20] and its fundamental rules of conduct, or constitution [AST21]. This type of digital identities blueprint would thus be especially suitable for digital democracy applications implemented by grassroots civic communities. Such a bottom-up, trust-based approach may resonate with otherwise disenfranchised citizens by facilitating the self-organization of digital civic participation.

It should be clear that at the moment, we are still far from a democratic ideal of digital participation, be it institution-based or grassroots. The predominant contemporary vehicle for online democratic participation are social media platforms whose workings are inherently autocratic: they are closed-source and are controlled, in one way or another, by a central authority (the company owning the system) that does not respond to any democratic oversight from its users. Any form of digital democracy mediated by such platforms makes citizens inevitably and strongly dependent on whoever controls those systems [BLR21, For21, MT18].

4 Closing words

Digital democracy is increasingly perceived by citizens and public authorities as a technology that could offer a viable vehicle for more citizen-responsive policy making. As such, it holds the promise of readdressing—at least in part—the growing societal discontent towards the way democracy is being practiced today. Yet, we know too little still about how social interaction and decision-making mediated by digital technology actually works: what are its pitfalls and how digital technology could be leveraged to foster democratic ideals in the decision-making processes of our societies.

If digital democracy is to meet its promise, we need a much more robust understanding of it as a socio-technical system. In this paper we motivated and outlined a long-term research program of interdisciplinary science that can provide the missing foundations for the responsible development of a transparent, open and fit-for-purpose digital democratic infrastructure for our societies.

Authors

Coordinating authors (alphabetical order):

- Davide Grossi (University of Groningen and University of Amsterdam, The Netherlands)
- Ulrike Hahn (Birkbeck College London, United Kingdom)
- Michael Mäs (Karlsruhe Institute of Technology, Germany)
- Andreas Nitsche (Interaktive Demokratie, Germany)

Contributing authors (alphabetical order):

- Jan Behrens (Interaktive Demokratie, Germany)
- Niclas Boehmer (Harvard University, USA)
- Markus Brill (University of Warwick, UK)
- Ulle Endriss (ILLC, University of Amsterdam, The Netherlands)

- Umberto Grandi (IRIT, Université Toulouse Capitole, France)
- Adrian Haret (MCMP, LMU Munich, Germany)
- Jobst Heitzig (Potsdam Institute for Climate Impact Research, Germany)
- Nicolien Janssens (EIPE, Erasmus University Rotterdam, The Netherlands)
- Catholijn M. Jonker (Technical University Delft & University Leiden, The Netherlands)
- Marijn A. Keijzer (Institute for Advanced Study in Toulouse, France)
- Axel Kistner (Interaktive Demokratie, Germany)
- Martin Lackner (TU Wien, Vienna, Austria)
- Alexandra Lieben (University of California, Los Angeles, USA)
- Anna Mikhaylovskaya (University of Groningen, The Netherlands)
- Pradeep K. Murukannaiah (TU Delft, The Netherlands)
- Carlo Proietti (CNR, Italy)
- Manon Revel (Harvard University, USA)
- Élise Rouméas (University of Groningen, The Netherlands)
- Ehud Shapiro (Weizmann Institute of Science, Israel)
- Gogulapati Sreedurga (University of Edinburgh, United Kingdom)
- Björn Swierczek (Interaktive Demokratie, Germany)
- Nimrod Talmon (Ben-Gurion University, Israel)
- Paolo Turrini (University of Warwick, United Kingdom)
- Zoi Terzopoulou (GATE Lyon-Saint-Etienne, Jean-Monnet University, France)
- Frederik Van De Putte (Erasmus University Rotterdam, The Netherlands)

Acknowledgments

The authors would like to acknowledge the generous support of the Netherlands Institute for Advanced Studies in the Humanities and Social Sciences (NIAS) and the Lorentz Center.

Paolo Turrini acknowledges the support of the Leverhulme Trust for the Research Grant RPG-2023-050 titled "Promoting Social Good Using Social Networks". Umberto Grandi acknowledges the support of the ANR JCJC project "Social Choice on Social Networks" (ANR 18-CE23-0009-01). Martin Lackner was supported by the Austrian Science Fund (FWF): P31890. Davide Grossi, Catholijn Jonker and Pradeep Murukannaiah acknowledge support by the Hybrid Intelligence Center, a 10-year program funded by the Dutch Ministry of Education, Culture and Science through the Netherlands Organisation for Scientific Research (NWO).

References

[AHALT22] Shiri Alouf-Heffetz, Ben Armstrong, Kate Larson, and Nimrod Talmon. How should we vote? A comparison of voting systems within social networks. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 31–38, 2022.

[AM19] Ben Abramowitz and Nicholas Mattei. Flexible representative democracy: An introduction with binary issues. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 3–10, 2019.

[And72] Philip W Anderson. More is different: broken symmetry and the nature of the hierarchical structure of science. *Science*, 177(4047):393–396, 1972.

[AST21] Ben Abramowitz, Ehud Shapiro, and Nimrod Talmon. In the beginning there were n agents: Founding and amending a constitution. In *Algorithmic Decision Theory: 7th International Conference, ADT 2021, Toulouse, France, November 3–5, 2021, Proceedings 7*, pages 119–131. Springer, 2021.

[Bar16] Albert-László Barabási. *Network Science*. Cambridge University Press, 2016.

[BCE+16] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

[BCH21] Jason W. Burton, Nicole Cruz, and Ulrike Hahn. Reconsidering evidence of moral contagion in online social networks. *Nature Human Behaviour*, 5(12):1629–1635, 2021.

[BCS+22] Michiel A. Bakker, Martin J. Chadwick, Hannah R. Sheahan, Michael Henry Tessler, Lucy Campbell-Gillingham, Jan Balaguer, Nat McAleese, Amelia Glaese, John Aslanides, Matthew M. Botvinick, and Christopher Summerfield. Fine-tuning language models to find agreement among humans with diverse preferences. In *Proceedings of the 36th Conference on Neural Information Processing Systems (NeurIPS)*, 2022.

[BDMW18] Andre Bächtiger, John S. Dryzek, Jane Mansbridge, and Mark Warren. *The Oxford Handbook of Deliberative Democracy*. Oxford University Press, 2018.

[BEN22] Arthur Boixel, Ulle Endriss, and Oliviero Nardi. Displaying justifications for collective decisions. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, 2022. Demo Paper.

[BG15] Tim Baarslag and Enrico H. Gerding. Optimal incremental preference elicitation during negotiation. In *Proceedings of the 24th International Conference on Artificial Intelligence (IJCAI)*, 2015.

[BHW92] Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of political Economy*, 100(5):992–1026, 1992.

[BKNS14] Jan Behrens, Axel Kistner, Andreas Nitsche, and Björn Swierczek. *The Principles of LiquidFeedback*. Interaktive Demokratie, 2014.

[BKNS22] Jan Behrens, Axel Kistner, Andreas Nitsche, and Björn Swierczek. Implementation issues of collective decision systems in practice. *The Liquid Democracy Journal on Electronic Participation, Collective Moderation, and Voting Systems*, 7, 2022.

[BLR21] Lucy Bernholz, Hélène Landemore, and Rob Reich. *Digital Technology and Democratic Theory*. University of Chicago Press, 2021.

[BLT22] Jacques Bara, Omer Lev, and Paolo Turrini. Predicting voting outcomes in the presence of communities, echo chambers and multiple parties. *Artificial Intelligence*, 312:103773, 2022.

[BNPS21] Gerdus Benade, Swaprava Nath, Ariel D Procaccia, and Nisarg Shah. Preference elicitation for participatory budgeting. *Management Science*, 67(5):2813–2827, 2021.

[Bou20] Terrill G Bouricius. Democracy through multi-body sortition: Athenian lessons for the modern day. *Journal of Deliberative Democracy*, 9(1), 2020.

[Bre22] Richard Breen. Causal inference with observational data. In *Handbook of Sociological Science*, pages 272–286. Edward Elgar Publishing, 2022.

[BSST21] Laurent Bulteau, Gal Shahaf, Ehud Shapiro, and Nimrod Talmon. Aggregation over metric spaces: Proposing and voting in elections, budgeting, and legislation. *Journal of Artificial Intelligence Research*, 70:1413–1439, 2021.

[BZ16] Christian Blum and Christina Isabel Zuber. Liquid democracy: Potentials, problems, and perspectives. *Journal of Political Philosophy*, 24(2):162–182, 2016.

[CH23] Julian Chingoma and Adrian Haret. Deliberation as evidence disclosure: A tale of two protocol types. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2589–2597, 2023.

[Che03] Alexander Chernev. When more is less and less is more: The role of ideal point availability and assortment in consumer choice. *Journal of Consumer Research*, 30(2):170–183, 2003.

[CJK+19] Valerio Capraro, Glorianna Jagfeld, Rana Klein, Mathijs Mul, and Iris van de Pol. Increasing altruistic and cooperative behaviour with simple moral nudges. *Scientific Reports*, 9(1):1–11, 2019.

[CK02] Donald E. Campbell and Jerry S. Kelly. Impossibility theorems in the Arrovian framework. In Kenneth J. Arrow, Amartya K. Sen, and Kotaro Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 1, pages 35–94. Elsevier, 2002.

[CV19] Valerio Capraro and Andrea Vanzo. The power of moral words: Loaded language generates framing effects in the extreme dictator game. *Judgment and Decision Making*, 14(3):309–317, 2019.

[Dah08] Robert A. Dahl. *Democracy and its Critics*. Yale University Press, 2008.

[DMP15] Michalis Drouvelis, Robert Metcalfe, and Nattavudh Powdthavee. Can priming cooperation increase public good contributions? *Theory and Decision*, 79(3):479–492, 2015.

[DP21] Huihui Ding and Marcus Pivato. Deliberation and epistemic democracy. *Journal of Economic Behavior & Organization*, 185:138–167, 2021.

[Dry00] John S. Dryzek. *Deliberative Democracy and Beyond*. Oxford University Press, 2000.

[Dur16] Emile Durkheim. The rules of sociological method. In *Social Theory Re-Wired*, pages 10–15. Routledge, 2016.

[EG18] Ryan D. Enos and Noam Gidron. Exclusion and cooperation in diverse societies: Experimental evidence from Israel. *American Political Science Review*, 112(4):742–757, 2018.

[EGST21] Edith Elkind, Davide Grossi, Ehud Shapiro, and Nimrod Talmon. United for change: deliberative coalition formation to change the status quo. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence*, pages 5339–5346, 2021.

[EK10] David Easley and Jon Kleinberg. *Networks, crowds, and markets: Reasoning about a highly connected world*. Cambridge university press, 2010.

[Eps08] Joshua M. Epstein. Why model? *Journal of artificial societies and social simulation*, 11(4):12, 2008.

[Eur22] European Union. Conference on the Future of Europe: Report on the final outcome, 2022.

[FA12] Morris P. Fiorina and Samuel J. Abrams. *Disconnect: The breakdown of representation in American politics*, volume 11. University of Oklahoma Press, 2012.

[FGP+23] Sara Fish, Paul Gölz, David C. Parkes, Ariel D. Procaccia, Gili Rusak, Itai Shapira, and Manuel Wüthrich. Generative social choice, 2023. arXiv:2309.01291.

[FGR16] Seth Flaxman, Sharad Goel, and Justin M Rao. Filter bubbles, echo chambers, and online news consumption. *Public Opinion Quarterly*, 80(S1):298–320, 2016.

[FIK17] Katsuhide Fujita, Takayuki Ito, and Mark Klein. Enabling large scale deliberation using ideation and negotiation-support agents. In *2017 IEEE 37th International Conference on Distributed Computing Systems Workshops (ICDCSW)*, pages 360–363. IEEE, 2017.

[FKS+20] Roberto Stefan Foa, Andrew Klassen, Micheal Slade, Alex Rand, and Rosie Collins. The global satisfaction with democracy report 2020. Centre for the Future of Democracy, University of Cambridge, 2020.

[FMF+17] Andreas Flache, Michael Mäs, Thomas Feliciani, Edmund Chattoe-Brown, Guillaume Deffuant, Sylvie Huet, and Jan Lorenz. Models of social influence: Towards the next frontiers. *Journal of Artificial Societies and Social Simulation*, 20(4), 2017.

[FMK22] Andreas Flache, Michael Mäs, and Marijn A. Keijzer. Computational approaches in rigorous sociology: Agent-based computational sociology and computational social science. In Klarita Gërxhani, Nan-Dirk De Graaf, and Werner Raub, editors, *Handbook of sociological science. Contributions to Rigorous Sociology*, chapter 4, pages 57–72. Edward Elgar Publishing, Cheltenham, UK, 2022.

[For21] Bryan Alexander Ford. Technologizing democracy or democratizing technology? A layered-architecture perspective on potentials and challenges. In Lucy Bernholz, Hélène Landemore, and Rob Reich, editors, *Digital Technology and Democratic Theory*. University of Chicago Press, 2021.

[FSDB21] James Fishkin, Alice Siu, Larry Diamond, and Norman Bradburn. Is deliberation an antidote to extreme partisan polarization? Reflections on "America in one room". *American Political Science Review*, 115(4):1464–1481, 2021.

[FSST17] Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. Multiwinner voting: A new challenge for social choice theory. *Trends in computational social choice*, 74(2017):27–47, 2017.

[Fun15] Archon Fung. Putting the public back into governance: The challenges of citizen participation and its future. *Public administration review*, 75(4):513–522, 2015.

[GA15] James Green-Armytage. Direct voting and proxy voting. *Constitutional Political Economy*, 26(2):190–220, 2015.

[GGIM22] Lodewijk Gelauff, Ashish Goel, Sungjin Im, and Kamesh Munagala. Representational robustness in social choice. In *Proceedings of the ACM Collective Intelligence Conference (CI)*, 2022.

[Gil05] Martin Gilens. Inequality and democratic responsiveness. *Public Opinion Quarterly*, 69(5):778–796, 2005.

[GK18] Christian Ganser and Marc Keuschnigg. Social influence strengthens crowd wisdom under voting. *Advances in Complex Systems*, 21(06n07):1850013, 2018.

[GPSO12] Inma Garcia, Sergio Pajares, Laura Sebastia, and Eva Onaindia. Preference elicitation techniques for group recommender systems. *Information Sciences*, 189:155–175, 2012.

[GR08] Jacob Glazer and Ariel Rubinstein. A study in the pragmatics of persuasion: A game theoretical approach. *New Perspectives on Games and Interaction*, pages 121–140, 2008.

[GZY+13] Jingjing Gong, Yan Zhang, Zheng Yang, Yonghua Huang, Jun Feng, and Weiwei Zhang. The framing effect in medical decision-making: A review of the literature. *Psychology, Health & Medicine*, 18(6):645–653, 2013.

[Hab04] Jürgen Habermas. Discourse ethics. In *Ethics: Contemporary Readings*, pages 146–153. Routledge, 2004.

[Her86] Paul M. Herr. Consequences of priming: Judgment and behavior. *Journal of Personality and Social Psychology*, 51(6):1106, 1986.

[HHP17] George C. Homans, A. Paul Hare, and Richard Brian Polley. *The Human Group*. Routledge, 2017.

[HHPW07] James Habyarimana, Macartan Humphreys, Daniel N. Posner, and Jeremy M. Weinstein. Why does ethnic diversity undermine public goods provision? *American Political Science Review*, 101(4):709–725, 2007.

[HKP+23] Daniel Halpern, Gregory Kehne, Ariel D. Procaccia, Jamie Tucker-Foltz, and Manuel Wüthrich. Representation with incomplete votes. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence*, pages 5657–5664. AAAI Press, 2023.

[HT06] Marc Hassenzahl and Noam Tractinsky. User experience—A research agenda. *Behaviour & Information Technology*, 25(2):91–97, 2006.

[Jan83] Irving Lester Janis. *Groupthink*. Houghton Mifflin Boston, 1983.

[JO86] Hubert J. O'Gorman. The discovery of pluralistic ignorance: An ironic lesson. *Journal of the History of the Behavioral Sciences*, 22(4):333–347, 1986.

[KA87] Jon A. Krosnick and Duane F. Alwin. An evaluation of a cognitive theory of response-order effects in survey measurement. *Public Opinion Quarterly*, 51(2):201–219, 1987.

[KAJ31] Daniel Katz, Floyd Henry Allport, and Margaret Babcock Jenness. *Students' attitudes: A report of the Syracuse University reaction study*. Craftsman Press, 1931.

[KJ97] Jack Knight and James Johnson. What sort of equality does deliberative democracy require? In *Deliberative Democracy: Essays on Reason and Politics*, pages 279–320. MIT Press, 1997.

[Kle11] Mark Klein. The mit deliberatorium: Enabling large-scale deliberation about complex systemic problems. *International Conference on Agents and Artificial Intelligenc*, 01 2011.

[KM22] Marijn A. Keijzer and Michael Mäs. The complex link between filter bubbles and opinion polarization. *Data Science*, 5(2):139–166, 2022.

[Lan20] Hélène Landemore. *Open Democracy: Reinventing Popular Rule for the Twenty-First Century*. Princeton University Press, 2020.

[Lan22] Hélène Landemore. Can AI bring deliberative democracy to the masses?, 2022. URL: https://www.law.nyu.edu/sites/default/files/Helen%20Landemore%20Can%20AI%20bring%20deliberative%20democracy%20to%20the%20masses.pdf.

[LGLA14] David Lee, Ashish Goel, Hélène Landemore, and Tanja Aitamurto. Crowdsourcing for participatory democracies: Efficient elicitation of social choice functions. *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, 2, 06 2014.

[Lij90] Arend Lijphart. The political consequences of electoral laws, 1945–85. *American Political Science Review*, 84(2):481–496, 1990.

[Lis18] Christian List. Democratic deliberation and social choice. In *The Oxford Handbook of Deliberative Democracy*. Oxford University Press, 2018.

[LLFM13] Christian List, Robert C. Luskin, James S. Fishkin, and Iain McLean. Deliberation, single-peakedness, and the possibility of meaningful democracy: Evidence from deliberative polls. *The Journal of Politics*, 75(1):80–95, 2013.

[LMXF22] Shuo Liu, Michael Mas, Haoxiang Xia, and Andreas Flache. When intuition fails: the complex effects of assimilative and repulsive influence on opinion polarization. *Advances in Complex Systems*, 2022.

[LN18] Jan Lorenz and Martin Neumann. Opinion dynamics and collective decisions. *Advances in Complex Systems*, 21(06n07):1802002, 2018.

[LP20] Dimitri Landa and Ryan Pevnick. Representative democracy as defensible epistocracy. *American Political Science Review*, 114(1):1–13, 2020.

[LSG+20] Stephan Lewandowsky, Laura Smillie, David Garcia, Ralph Hertwig, Jim Weatherall, Stefanie Egidy, Ronald E. Robertson, Cailin O'Connor, Anastasia Kozyreva, Philipp Lorenz-Spreen, Yannik Blaschke, and Mark Mark Leiser. Technology and democracy: Understanding the influence of online technologies on political behaviour and decision-making. JRC Science for Policy Report, Publications Office of the European Union, 2020.

[Man09] Jane Mansbridge. A "selection model" of political representation. *Journal of Political Philosophy*, 17(4):369–398, 2009.

[MANC23] Gregory Michener, Octavio Amorim Neto, and Jamil Civitarese. The remoteness of democratic representation. *Party Politics*, 29(1):51–64, 2023.

[Mat20] John G Matsusaka. Let the people rule. In *Let the People Rule*. Princeton University Press, 2020.

[Mik24] Anna Mikhaylovskaya. Enhancing deliberation with digital democratic innovations. *Philosophy & Technology*, 37(1):3, 2024.

[Mil24] John Stuart Mill. *Representative Government*. Alex Catalogue; NetLibrary, 1924.

[Mil69] James C. Miller. A program for direct and proxy voting in the legislative process. *Public Choice*, 7(1):107–113, 1969.

[MP15] Vincent Mousseau and Marc Pirlot. Preference elicitation and learning. *EURO Journal on Decision Processes*, 3(1-2):1–3, 2015.

[MR09] Ricardo Matheus and Manuella Maia Ribeiro. Models for citizen engagement in latin american: Case studies of public digital budgeting. In *Proceedings of the 3rd international conference on Theory and practice of electronic governance*, pages 109–116, 2009.

[MT18] M. Moore and D. Tambini. *Digital Dominance: The Power of Google, Amazon, Facebook, and Apple*. Oxford University Press, 2018.

[MTSS22] Reshef Meir, Nimrod Talmon, Gal Shahaf, and Ehud Shapiro. Sybil-resilient social choice with low voter turnout. In *European Conference on Multi-Agent Systems*, pages 257–274. Springer, 2022.

[MZ69] Serge Moscovici and Marisa Zavalloni. The group as a polarizer of attitudes. *Journal of Personality and Social Psychology*, 12(2):125, 1969.

[NC22] Brandie Nonnecke and Camille Carlton. EU and US legislation seek to open up digital platform data. *Science*, 375(6581):610–612, 2022.

[New18] Mark Newman. *Networks*. Oxford university press, 2018.

[Ngu21] Thu Nguyen. Time to be honest: The future will not be decided by the Conference on the Future of Europe. Policy position paper, Hertie School: Jacques Delors Centre, 2021. URL: https://www.delorscentre.eu/en/publications/detail/publication/time-to-be-honest-the-future-will-not-be-decided-by-the-conference-on-the-future-of-eur

[NVE+22] Simon Niemeyer, Francesco Veri, Selen Ercan, Hans Asenbaum, Ricardo Mendonça, and Nicole Curato. Deliberative reason index. In *Research Methods in Deliberative Democracy*, pages 99–114. Oxford University Press, 2022.

[OEC20] OECD. Innovative citizen participation and new democratic institutions: Catching the deliberative wave. *Organization for Economic Co-operation and Development*, 2020. URL: https://www.oecd.org/gov/innovative-citizen-participation-and-new-democratic-institutions-339306da-en.htm.

[Par19] A. Parsons. Digital tools for citizens assemblies. *mySociety Research*, 2019. URL: https://research.mysociety.org/publications/digital-tools-citizens-assemblies.

[Piv17] Marcus Pivato. Epistemic democracy with correlated voters. *Journal of Mathematical Economics*, 72:51–69, 2017.

[Piv19] Marcus Pivato. Realizing epistemic democracy. In Jean-François Laslier, Hervé Moulin, M. Remzi Sanver, and William S. Zwicker, editors, *The Future of Economic Design*, pages 103–112. Springer, 2019.

[PL+20] Ismael Peña-López et al. Innovative citizen participation and new democratic institutions: Catching the deliberative wave. OECD Publishing, 2020.

[PO14] Carlo Proietti and Erik J. Olsson. A ddl approach to pluralistic ignorance and collective belief. *Journal of Philosophical Logic*, 43:499–515, 2014.

[Pro19] Ariel D Procaccia. Axioms should explain solutions. *The Future of Economic Design: The Continuing Development of a Field as Envisioned by Its Researchers*, pages 195–199, 2019.

[Prz17] Anna Przybylska. *ICT for Dialogue and Inclusive Decision-Making*. Peter Lang, 2017.

[PSST19] Ouri Poupko, Gal Shahaf, Ehud Shapiro, and Nimrod Talmon. Sybil-resilient conductance-based community growth. In René van Bevern and Gregory Kucherov, editors, *Computer Science – Theory and Applications*, pages 359–371. Springer International Publishing, 2019.

[PSST21] Ouri Poupko, Gal Shahaf, Ehud Shapiro, and Nimrod Talmon. Building a sybil-resilient digital community utilizing trust-graph connectivity. *IEEE/ACM transactions on networking*, 29(5):2215–2227, 2021.

[Rae67] Douglas W. Rae. *The political consequences of electoral laws*. New Haven: Yale University Press, 1967.

[Raw71] John Rawls. *A Theory of Justice*. Belknap Press, 1971.

[RBC+23] Manon Revel, Niclas Boehmer, Rachael Colley, Markus Brill, Piotr Faliszewski, and Edith Elkind. Selecting representative bodies: An axiomatic view. *arXiv preprint arXiv:2304.02774*, 2023.

[Rea05] Daniel Read. Monetary incentives, what are they good for? *Journal of Economic Methodology*, 12(2):265–276, 2005.

[REdH21] Simon Rey, Ulle Endriss, and Ronald de Haan. Shortlisting rules and incentives in an end-to-end model for participatory budgeting. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, 2021.

[RLL22] Tiina Randma-Liiv and Veiko Lember. *Engaging Citizens in Policy Making: E-participation Practices in Europe*. Edward Elgar Publishing, 2022.

[SBBM17] Julie Simon, Theo Bass, Victoria Boelman, and Geoff Mulgan. Digital democracy: The tools transforming political engagement. Technical report, Nesta, 2017.

[SBE+21] Christopher Small, Michael Bjorkegren, Timo Erkkilä, Lynette Shaw, and Colin Megill. Polis: Scaling deliberation by mapping high dimensional opinion spaces. *Recerca: Revista de Pensament i Anàlisi*, 26(2), 2021.

[SBSS03] Marco R Steenbergen, André Bächtiger, Markus Spörndli, and Jürg Steiner. Measuring political deliberation: A discourse quality index. *Comparative European Politics*, 1:21–48, 2003.

[SCG13] Zaher Salah, Frans Coenen, and Davide Grossi. Extracting debate graphs from parliamentary transcripts: A study directed at UK House of Commons debates. In *Proceedings of the 14th International Conference on Artificial Intelligence and Law (ICAIL)*, pages 121–130. ACM, 2013.

[Sch04] Barry Schwartz. *The Paradox of Choice: Why More is Less*. Harper Perennial, 2004.

[Sch06] Thomas C Schelling. *Micromotives and macrobehavior*. WW Norton & Company, 2006.

[SIM+22a] Ruth Shortall, Anatol Itten, Michiel van der Meer, Pradeep Murukannaiah, and Catholijn Jonker. Reason against the machine? Future directions for mass online deliberation. *Frontiers in Political Science*, 4, 2022.

[SIM+22b] Ruth Shortall, Anatol Itten, Michiel van der Meer, Pradeep Murukannaiah, and Catholijn Jonker. Reason against the machine? future directions for mass online deliberation. *Frontiers in Political Science*, 4:946589, 2022.

[SLB+17] Piotr Skowron, Martin Lackner, Markus Brill, Dominik Peters, and Edith Elkind. Proportional rankings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 409–415, 2017.

[SST19] Gal Shahaf, Ehud Shapiro, and Nimrod Talmon. Sybil-resilient reality-aware social choice. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 572–579, 2019.

[SST20] Gal Shahaf, Ehud Shapiro, and Nimrod Talmon. Genuine personal identifiers and mutual sureties for sybil-resilient community growth. In *Social Informatics: 12th International Conference, SocInfo 2020, Pisa, Italy, October 6–9, 2020, Proceedings 12*, pages 320–332. Springer, 2020.

[Sto61] James Arthur Finch Stoner. *A comparison of individual and group decisions involving risk*. PhD thesis, Massachusetts Institute of Technology, 1961.

[Sun14] Cass R. Sunstein. Nudging: A very short guide. *Journal of Consumer Policy*, 37(4):583–588, 2014.

[Sur05] James Surowiecki. *The Wisdom of Crowds*. Anchor, 2005.

[TIY+13] János Török, Gerardo Iñiguez, Taha Yasseri, Maxi San Miguel, Kimmo Kaski, and János Kertész. Opinions, conflicts, and consensus: Modeling social dynamics in a collaborative environment. *Physical Review Letters*, 110(8):088701, 2013.

[TKGH23] Zoi Terzopoulou, Marijn A. Keijzer, Sreedurga Gogulapati, and Jobst Heitzig. The rule–tool–user nexus in digital collective decisions. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2023.

[TMLS09] Nguyen Tran, Bonan Min, Jinyang Li, and Lakshminarayanan Subramanian. Sybil-resilient online content voting. In *Proceedings of the 6th USENIX Symposium on Networked Systems Design and Implementation*, NSDI'09, page 15–28, USA, 2009. USENIX Association.

[TQ19] Fei Tao and Qinglin Qi. Make more digital twins. *Nature*, 490(573):490–491, 2019.

[Tul67] Gordon Tullock. *Toward a mathematics of politics*. Ann Arbor: University of Michigan Press, 1967.

[Tul92] Gordon Tullock. Computerizing politics. *Mathematical and Computer Modelling*, 16(8-9):59–65, 1992.

[Val21] Chiara Valsangiacomo. Political representation in liquid democracy. *Frontiers in Political Science*, page 7, 2021.

[VdP13] Ibo Van de Poel. Translating values into design requirements. *Philosophy and engineering: Reflections on practice, principles and process*, pages 253–266, 2013.

[VdPR11] Ibo Van de Poel and Lambèr Royakkers. *Ethics, technology, and engineering: An introduction*. John Wiley & Sons, 2011.

[VdSLB13] Karine Van der Straeten, Jean-François Laslier, and André Blais. Vote au pluriel: How people vote when offered to vote under different rules. *Political Science & Politics*, 46(2):324–328, 2013.

[VP17] Cyril Velikanov and Alexander Prosser. Mass online deliberation in participatory policy-making—Part I: Rationale, lessons from past experiments, and requirements. *Beyond Bureaucracy: Towards Sustainable Governance Informatisation*, pages 209–234, 2017.

[VR16] David Van Reybrouck. *Against Elections: The Case for Democracy*. Random House UK, 2016.

[VRA18] Soroush Vosoughi, Deb Roy, and Sinan Aral. The spread of true and false news online. *science*, 359(6380):1146–1151, 2018.

[WEATC+22] Pascal Wagner-Egger, Jais Adam-Troian, Laurent Cordonier, Florian Cafiero, and Gérald Bronner. The yellow vests in France: Psychosocial determinants and consequences of the adherence to a social movement in a representative sample of the population. *International Review of Social Psychology*, 35(1), 2022.

[WF21] Richard Wike and Janell Fetterolf. Global public opinion in an era of democratic anxiety. *Pew Research*, 2021. URL: https://www.pewresearch.org/global/2021/12/07/global-public-opinion-in-an-era-of-democratic-anxiety/.

[WS22] A. Wanless and J.N. Shapiro. A CERN model for studying the information environment. *Carnegie Endowment for World Peace*, 2022.

[You22] Richard Youngs. EU democracy after the Conference on the Future of Europe. *Carnegie Europe*, 2022. URL: https://carnegieeurope.eu/2022/05/12/eu-democracy-after-conference-on-future-of-europe-pub-87110.