



Universiteit
Leiden
The Netherlands

Computing consensus: language ideological work in LLM-assisted deliberative democracy

Connor, J.E.

Citation

Connor, J. E. (2026). Computing consensus: language ideological work in LLM-assisted deliberative democracy. *Journal Of Linguistic Anthropology*. doi:10.1111/jola.70071

Version: Publisher's Version

License: [Creative Commons CC BY-NC-ND 4.0 license](#)

Downloaded from: <https://hdl.handle.net/1887/4308916>

Note: To cite this publication please use the final published version (if applicable).

ORIGINAL ARTICLE

Computing consensus: Language ideological work in LLM-assisted deliberative democracy

Janet E. Connor 

Leiden University, Leiden, The Netherlands

Correspondence

Janet E. Connor, Leiden University,
Leiden, The Netherlands.

Email: j.e.connor@hum.leidenuniv.nl

Abstract

This article examines how the digital platform *Polis* and its experiments with large language models (LLMs) reconfigure democratic participation through a particular vision of consensus. Through analysis of media coverage of Polis and developers' public discussions of the platform, I argue that developers reframe consensus from a discursive achievement to a computationally surfaced pattern, turning deliberative democracy into a computing problem which LLMs are perfectly fit to solve. Under the guise of “saving” or “improving” democracy, Polis produces “illusions of consensus” that depoliticize purported public opinion and further abstract it away from the individuals assumed to have produced it.

KEYWORDS

consensus, democracy, languages ideologies, large language models, publics

Samenvatting

Dit artikel onderzoekt hoe het digitale platform Polis en zijn experimenten met large language models (LLM's) democratische participatie herconfigureren via een specifieke consensusvisie. Op basis van een analyse van mediaberichtgeving over Polis en publieke ontwikkelaarsdiscussies betoog ik dat ontwikkelaars consensus herdefiniëren van discursief resultaat naar computationeel zichtbaar gemaakt patroon, waarmee deliberatieve democratie in hun ogen een rekenprobleem wordt waarvoor LLM's bij uitstek geschikt zouden zijn. Onder het mom van het “redden” of “verbeteren” van de democratie produceert

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs](https://creativecommons.org/licenses/by-nc-nd/4.0/) License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2026 The Author(s). *Journal of Linguistic Anthropology* published by Wiley Periodicals LLC on behalf of American Anthropological Association.

Polis “illusies van consensus” die de veronderstelde publieke opinie depolitiseren en verder abstraheren van de individuen die haar geacht worden te hebben voortgebracht.

Most of what we hear about democracy and AI focuses on the dangers AI poses to democracy. Yet there is also a group of technologists who claim that AI, when used in the right way, by the right people, can in fact save democracy. Yet in claiming to be able to save democracy, these technologists are at the same time producing a particular idea of what democracy is and the discursive work that constitutes it. One of perhaps the most talked about organizations in this group is the US-based Computational Democracy Project (CompDem), which created Polis (sometimes written as pol.is), an open-source software designed to facilitate democratic decision-making among large groups. A Polis deliberation begins with participants submitting short text statements about a particular topic set by the organizers. Those comments are then sent out to other participants, who can agree, disagree, or pass on them. The platform organizes those votes into a matrix, with rows corresponding to participants, and columns corresponding to comments. Through common statistical methods of dimension reduction and clustering, participants are then mapped onto an “opinion space,” where participants sharing similar opinions are grouped together. Polis then identifies the comments that received the most agree votes for each opinion group, which it calls “consensus comments” (Small et al., 2021, 5).¹ These become the group’s shared opinion. While Polis does not currently use large language models, the reliance on matrices and vectors uses the same principles of machine learning, and CompDem has been experimenting with integrating LLMs into the platform.

Media coverage of Polis praises it for its ability to “f[ind] consensus where none seemed to exist” (Miller, 2019). Polis is primarily known for its frequent use by the Taiwanese government, but it has also been used by national and municipal governments, as well as civic associations around the world. I first learned about Polis during exploratory research for a project on participatory democracy in the Netherlands. I was curious how Dutch conceptions of democracy came together with the emergence of new technologies designed to increase democratic participation. An employee at a Dutch pro-democracy NGO told me to contact Rogier de la Rive Box, the leader of a civic association in a small city near Amsterdam, about his experiences using the tool. On first look, it would seem like Polis is the perfect fit for Dutch organizations: while Polis’s main aim is to “realize consensus around...controversial question[s],”² narratives of Dutch democracy also center on consensus. Key to foundational myths of Dutch democracy is the consensus-oriented “polder model,” which goes back to the Middle Ages, when individual farmers and cities had to cooperate to maintain dikes so that polders, areas of land reclaimed from the sea, would remain above water. Today, the term “polder model” is used more metaphorically to speak about governance based in broad coalitions, although it tends to be more of an ideal than a reality as Dutch politics, like much of the world, has become more polarized. However, I learned that Polis was far from a perfect fit. As de la Rive Box explained to me over email, his organization had decided after one experiment not to continue using Polis because what it called “consensus” was not recognizable to him:

It is highly questionable whether consensus is the appropriate word here [to describe what Polis produces]. In our experiment, there were no substantive connections between the statements. Yet, consensus relationships were identified.

When the comments are not related, it seems to us that the word consensus is not applicable.³

De la Rive Box was referring to the final stage of a Polis deliberation, when the algorithm identifies the individually submitted comments that received the most votes for each opinion group. In the deliberation he had organized, de la Rive Box's organization could not see what the relationship between these different comments was, or what they could take away as group consensus. This disconnect occurred because, while de la Rive Box and most humans participating in deliberative democracy see language as referential, Polis works based on analyzing participants' voting patterns with machine learning, not on the denotational content of the statements being voted on. While the humans could see no denotational coherence, or "no substantive connections" between the statements people had voted on, Polis groups participants together on shared voting patterns alone, regardless of the denotational meaning of those statements.

De la Rive Box went on in his email to me to write that he sees consensus as something that can only be achieved through "physical discussion," and that his organization had decided not to continue using Polis and was instead working on building their own alternative.⁴ However, judging from the large amount of positive media attention CompDem has received, many people do see Polis to be helpful (e.g., Applebaum & Pomerantsev, 2021; Horton, 2018; Miller, 2019; Tang, 2019). While this coverage sometimes mentions problems like ease of use, I have not found any other public critique like de la Rive Box's about whether what Polis produces is really consensus, although some supporters do agree that Polis can only supplement discussions, not replace them (Landemore, 2024). At the same time, Polis's creators are moving even further from the position that consensus requires physical discussion to instead experiment with how LLMs could be integrated into the platform to help it identify consensus more quickly and with less human labor. These developers see LLMs as allowing for deliberative democracy to scale up, beyond the size where physical discussions are possible. This dissonance between de la Rive Box's experience, positive media coverage, and Polis's move to incorporate LLMs leads to this article's central questions: What is Polis producing, and why do supporters call it "consensus"? What assumptions about language and democracy does that framing rely on?

I argue that Polis is, to some extent, a continuation of an Enlightenment project, as outlined in Jurgen Habermas's *Structural Transformation of the Public Sphere* (1989), which saw "public opinion" as something that emerged through the circulation of rational, decontextualized discourse. At the same time, it tries to address longstanding critiques of the Habermasian public sphere's exclusivity, rethinking the communicative forms through which "the public" emerges. While, in the bourgeois public sphere, consensus was a discursive achievement, technologists instead see it as a computationally surfaced pattern. In so doing, they turn deliberative democracy into a computing problem which Large Language Models are uniquely fit to solve. Yet their "solutions" raise new questions and contradictions around what the public is, how it circulates, and the nature of consensus. Consensus arrived at through machine learning comes with a kind of authority based in probability and computation, not in "the people," adding an extra layer of abstraction away from the individual subjects who are assumed to have created it (cf. Warner, 2002). As I will show, similarly to how Messeri and Crockett (2024) have claimed that artificial intelligence risks creating "illusions of understanding" in scientific research, LLM-assisted deliberations would seem to create "illusions of consensus," producing widely acceptable, generic statements that may produce an appearance of agreement, or an idea that the LLM at least can see consensus. Yet these remain "illusions" for several reasons. First, they start from voting on decontextualized statements, assuming that these are representative of public opinion by themselves and ignoring the socially situated, relational way that opinion is formed (Bourdieu, 1979; see

also Cody, 2011a; Urban & Silverstein, 1996; Warner, 2002). Additionally, as others have already noted, while LLMs are skilled at reproducing textual genres (Gershon, 2023), they have no access to the world (Castelle, 2024; Handman, 2025; Keane, 2024). LLM-supported consensus lacks both the interactional grounding that linguistic anthropologists have shown to be central to how publics are created and any denotational reference. As a result, LLMs not only fail to create a more inclusive public sphere, but they also further abstract a purported public opinion away from material conflicts.

My argument builds on critiques in political science that AI engineers have misunderstood how democratic publics actually work and tried to remove the antagonism and power struggles of democracies to focus on deliberation and sortition (Farrell & Han, 2025). That resonates with longstanding critiques of Habermas (1989), that he assumes “public opinion” to be something that not only can, but ought to, be abstracted outside of the agonism of materialist relations (Calhoun, 1992; Fraser, 1992; Hirschkind, 2006; Warner, 1992, 2002). This criticism has drawn attention to the ways a purportedly universal public is built on exclusion; while a liberal public promises to abstract away from the individual bodies that make it up, some bodies, usually racialized, gendered, or otherwise marginalized, cannot be abstracted away from (Cody, 2015; Warner, 1992, 2002). Crucially for linguistic anthropologists, this abstraction relies on particular genres of language, “purified” of private passions, which came to be seen as best-suited to reason, and most capable of being decontextualized and circulating as a universal, disembodied voice (Bauman & Briggs, 2003). As I will show, technologists have taken this language ideology to an extreme, beginning from anonymous statements outside of any context, and culminating in consensus statements created through statistical probabilities rather than any grounding in the material world.

Linguistic anthropology has been central to problematizing the language ideological work that sustains the liberal public sphere. Previous work has drawn attention to the moments of interdiscursive, genred interaction through which publics are constructed and circulate (Cody, 2009, 2011b; Gal, 2006; Gal & Woolard, 2001; Graan, 2022b; Paz, 2016; Spitulnik, 1996; Warner, 2002; Yeh, 2018). Publics emerge through genred interactional events that are treated as interdiscursive with one another: the Habermasian bourgeois public sphere relied on discussions in salons and coffeehouses where strangers, who were almost exclusively wealthy, white men, came together, set aside their private, emotional selves, and rationally debated political concerns. Other publics and counterpublics rely on different communicative genres and forms of media, including circulating hearsay at the US-Mexico border (Yeh, 2018), listening to cassette recordings of preaching (Hirschkind, 2006), or silently reading a newspaper at home (Anderson, 2006; Cody, 2011a). Thinking with this work, I ask what kind of public LLM-assisted digital deliberation is creating.

In the rest of this article, I first situate Polis within a larger trend toward participatory democracy. I then discuss how Polis’s creators speak about democracy and the role of deliberation and communication within it. I show how Polis is simultaneously inscribing a particular view of democracy into its technology and framing democracy as a computational problem, which makes LLMs appear to be a natural next step in a progression toward a “better” democratic system. The final section analyzes how CompDem, in an experiment conducted with Anthropic’s LLM, Claude, imagines LLMs to be able to help Polis participants reach consensus. My discussion is based on publicly available interviews, lectures, and papers, primarily from CompDem’s founders, data scientist Christopher Small and technology designer Colin Megill. I also consulted raw data from Polis conversations available on github and media coverage about Polis. No one from CompDem responded to my attempts to contact them.

POLIS AND PARTICIPATORY DEMOCRACY

Polis, founded in 2014, emerged within a moment of increasing interest in participatory democracy that began in the early 2010s (Polletta, 2014). Participation in this form is often traced to the 1960s, when Students for a Democratic Society released the Port Huron Statement, and “maximum feasible participation” became a central concept in US President Johnson’s War on Poverty. At that time, participation was seen as a solution to representative democracy’s failure to include minorities, specifically African American men, and a response to critiques of expertise in urban planning (Kelty, 2019). Similar concerns motivate the more recent interest in participatory democracy, which ranges from grassroots activist movements like Occupy or the Arab Spring, to municipal experiments in participatory budgeting and citizens’ assemblies. Participatory democracy today has received increasing attention and support from institutions like the Council of Europe and the United Nations, which see it as an “antidote to populism” (Congress of Local and Regional Authorities, 2024). In this view, representative democracy is an adversarial form of democracy, based on competition between conflicting interests, which has led to current polarization. Participatory democracy, meanwhile, is a kind of unitary democracy, where a public deliberates to arrive at some kind of consensus or general will (Mansbridge, 1983; see also Barber, 1984). Scholars have pointed out that much of the focus on participation is ironically still about representation, trying to reach particular demographic groups and create what political scientists call “deliberative mini-publics,” small demographically representative groups like citizens assemblies, more than it is about in fact giving people more political power (Kelty, 2019; see also Farrell & Han, 2025 for a discussion specifically about AI).⁵ Rather than a move away from representation, these projects instead aim to *better* represent the public.

Participatory democracy also tends to be associated with the local and the small scale, where strangers living near one another can deliberate face-to-face (Dewey, 1970). The deliberative side of participatory projects is particularly influenced by Habermas’s discussion of the emergence of the bourgeois public sphere (1989). While Habermas viewed this kind of rational, deliberative public as a lost ideal, he tried in his later work to lay a foundation for the kinds of communicative practices through which we might create sustainable, democratic decision making (Habermas, 1984). Habermas’s public sphere assumed that uncoerced, rational discussion among equal strangers produced an authoritative, disinterested “public opinion.” Individuals set aside their emotional, private selves, and spoke in a decontextualizable, “purified” way, an ideal we can trace back to the Enlightenment, particularly John Locke (Bauman & Briggs, 2003).

Governments and organizations focused on participatory democracy today see the face-to-face nature of deliberation as a limitation and question how participation might be scaled to the national or even global level. There has recently been a wave of scholarship across political science and computer science that advocates incorporating LLMs into deliberative democracy (Landmore, 2024; Ovadya, 2023). These new tools frequently reference Habermas: Google DeepMind’s AI system, which has similarities to Polis, is even named the “Habermas Machine” (Tessler et al., 2024). Yet their ideas of democratic deliberation are more often filtered through work on “collective intelligence,” an idea of an enhanced capacity and knowledge that people have when working together, and often when also working with machines. Collective intelligence is based loosely on a term introduced by philosopher Pierre Lévy and draws wide-ranging inspiration from Durkheimian collective consciousness to biochemist Vladimir Vernadsky’s concept of the “noosphere,” where human reason will create a new evolutionary geological layer (Mulgan, 2017). Collective intelligence has been taken up by political theorists arguing that democratic deliberation produces “smarter” decisions than other political forms (e.g.,

Landemore, 2012), but it has also become popular in computer science and machine learning. As I show in the next section, these entanglements between political theory and computer science are key to how CompDem conceptualizes democracy, and, as a result, how they can claim to save it.

“INPUT CROWD, OUTPUT MEANING”

CompDem's public discourses about Polis show a circularity, where they are inscribing a particular vision of democracy into the technology, at the same time as the creators' backgrounds in computation shape how they understand democracy. What Polis has taken from Habermas is the assumption that a public is organized by the discussion and circulation of texts. Yet how circulation works, and especially what counts as speech, is very different for Polis compared to analog publics. A Polis comment is limited to 140 characters, similar to a Tweet at the time the platform was created. Yet Polis importantly differs from social media in that you cannot directly reply to statements and can instead only vote agree, disagree, or pass. Unintuitively, the lack of replies and threads is precisely why Polis's developers call these interactions “conversations.” One of CompDem's co-founders, Colin Megill, explained this in a discussion with journalist and political scientist Steven Mazie at a conference on democracy and social media organized at the Hannah Arendt Center for Politics and Humanities at Bard College in 2022 (Table 1).⁶

Mazie begins with the assumption that a conversation must include a response in words (line 2). A chat thread could therefore be called a conversation, but Polis's voting system could not. He then shifts to ask if the larger issue is that more standard kinds of rational discussion associated with the public sphere are not possible today because “everyone would just be jumping down each other's throats” if they were allowed to reply to comments (lines 4–5). Mazie here is invoking an old distinction between a rational, disinterested public, and the emotional, unruly, and potentially violent crowd (Le Bon, 1897). What distinguishes crowds and publics is that publics can be abstracted away from individual bodies, what Warner has called “utopias of self-abstraction” (Warner, 1992). While the 20th century crowd was typically raced, classed, and gendered, anyone who was “excessively embodied,” (Cody, 2015, 50), the crowd Mazie suggests here consists of people fighting on Twitter. Mazie wonders if Megill believes that conversations, in the sense of rational discussions, are no longer possible today. Megill responds by pointing out that the problem is not that most people are too irrational to participate in a discussion, but instead that most people on social media only read other people's posts and do not write their own. Megill's dismissal of reading as “lurking” (line 12) is a move away from other understandings of public sphere formation, where the fact that most people were “only” reading was not seen as a limitation, from Benedict Anderson (2006)'s description of the reading public key to the emergence of nationalism, to ethnographic studies (e.g., Cody, 2011a). This work has shown how the act of reading, whether silently to oneself or in a group, is key to the production of publics because it allows readers to imagine themselves as part of a larger group of strangers taking part in the same activity and the same interdiscursive interactions. Yet for Megill, reading is a problem because it means that most of the content that people are reading has been written by the few who are passionate enough to take the effort to write. We can also see similarities to Euro-American understandings of deliberation, which see speaking, as opposed to listening, as the ideal form of political communication because speaking is understood to be active, while listening is often seen as mere passive reception (Kramer, 2013; Slotta, 2023). Yet Megill does not see engaging 100% of people meaning everyone will be writing. While he thinks that some people will be able to write opinion statements, that is only his secondary

TABLE 1 Steven Mazie Interview with Colin Megill (2022).

Line	Speaker	Utterance
1	Steven Mazie:	It's called a "conversation" (makes air quotes) on Polis, but is it a
2		conversation if no one is actually responding in words to each other's
3		comments? And second, if the premise is correct that adding a reply
4		button would ruin the whole project because everyone would just be
5		jumping down each other's throats and it would turn into Twitter,
6		is that, uh (laughs), is that a sobering lesson about maybe the, the,
7		the impossibility of reasoned discussion without it devolving into
8		terrible forms of discussion?
9	Colin Megill:	Well I think one of the interesting bits perhaps in mechanism design,
10		which is to say, building a website for people to talk to each other,
11		when I say mechanism design, I'll just say that too. The, the number
12		of lurkers, that is to say, people who show up on a social media
13		platform and READ, and don't write, they might like, or they might
14		retweet, that's a different interaction they can have.
15		MOST people do not write.
16		MOST of the content is created by very few of the people.
17		And so when we discuss what is and is not a conversation on Twitter,
18		that, that decision, that the primary way of driving a conversation is
19		by writing something, means that you have to either have an
20		anonymous account if you, first you'd have to think of a statement,
21		that's actually a high-tax activity. You have to summarize what came
22		before, add a new thought. That's that's that's a lot of calories.
23		Liking is almost no calories, right? Uh, in glucose spent.
24		Uh, and so, you know, the the uh, the idea—
25	Mazie:	But the reward is high if someone likes your post.
26	Megill:	That's right, that's right, plenty of dopamine on that end of it, yeah.
27		So, the, the idea with restricting replies also is to say lurking is,
28		the primary action is to engage 100% of the people that are here.
29		The secondary action is the high-tax action.
30		And so, from a design standpoint, it, it, it's intended to capture,
31		as we saw, you know, potentially tens of thousands or millions
32		of votes from participants.

aim because it is a “high-tax activity” (lines 21 and 29). He instead sees clicking agree/disagree/pass as a sufficient amount of activity to count as engagement.

Megill then makes a register shift, away from Mazie’s question about deliberation and reasoned discussion, to instead use biological and psychological register shibboleths like calories (lines 22 and 23), glucose (line 23), and dopamine (line 26). This shift gives Megill authority by voicing official, measurable facts. It also flips the distinction between self-abstracted public and excessively embodied crowd to instead suggest that people who post on Twitter are the ones who are excessively embodied, who use “a lot of calories” (line 22). Polis’s design is not about limiting replies but instead “engag[ing] 100% of the people that are here” (line 28) in a way that uses “almost no calories” (line 23). The problem then, for

Megill, is twofold: on the one hand, most people are only observing, but not participating in a Twitter discussion, while those who are participating are only those who are exerting too much energy, in the form of too many calories, to be able to abstract away from their own bodies.

Yet this move is not about creating a self-aware, yet more inclusive, liberal public. Polis's tagline, "Input Crowd, Output Meaning," not only does not mention publics, but it includes what is often seen as the opposite of a public, the crowd, although here more in the tradition of "the wisdom of the crowd," the idea that a large number of people will probabilistically be more accurate than a single individual. In order to increase democratic participation, Megill is not advocating for disciplining individuals and creating a rational public, for example through improved education (Dewey, 1970). Instead, he is saying that how we think about what counts as democratic participation has to change, from a model where only certain bodies are able to be abstracted away from, to one where all bodies face the same biological limitations. Self-abstraction is no longer something that humans are themselves expected to achieve through deliberation in decontextualizable language, but something that an algorithm can take care of instead, turning low-calorie clicks into some kind of general will. The kind of collective Polis is creating grounds itself in the kind of biological metrics more commonly seen in tech discourses of gamification than in democratic theory. Under the guise of saving, and even improving, democracy, Polis inscribes a particular vision of what democracy is into the technology itself (Latour, 1992). In doing so, Polis fundamentally reimagines what democratic deliberation is, removing any sort of debate or agonism, and instead replacing that with a low-energy behavior more akin to doomscrolling.

Yet when speaking to different audiences, Megill uses a different justification for the same design decisions. When the only way to respond to a comment is by agree, disagree, or pass, you can easily turn those responses into a matrix, precisely the form machine learning algorithms are able to analyze. In a 2016 post on the Polis blog, Megill had explained the reason behind not allowing replies quite differently from what he later said to Mazie:

Doing away with replies gets you something very special. It gets you a matrix. You can have a look at what the data inside pol.is is like; it's not complicated. It's every participant, and what they thought about every comment. You can look for a participant and see how they agree, they agree, if they disagree, they disagree, they pass; or for a comment, you can see, this person agreed, this person disagreed, this person agreed, this person disagreed.

Humans aren't very good at analyzing this; but this scales up, and machines are really good at analyzing this (Megill, 2016).

While Megill justifies the decision to remove replies as a way to be more inclusive and low-energy when speaking to Mazie at Bard, it is also a way to turn the messiness of human interactions into a form data scientists are more comfortable with. The Polis conversation matrix is made up of rows for each participant, and columns for each comment. An agree vote is given a value of 1, a disagree -1, and a pass 0. Polis uses principal component analysis (PCA), a common data reduction method, to plot a participant's overall opinion. It then uses a clustering algorithm to group participants into opinion groups, where those who have voted in similar ways are grouped together. While Megill may call the commenting and voting process a "conversation," the kind of collective it creates is much closer to a different tradition in collective decision making that is based, not in deliberative discussion, but instead on forms of aggregation that rely on the development of probability theory (Landemore, 2012, 54). This tradition dates to Condorcet's Jury Theorem and has continued to develop in relation to probabilistic modeling.

Polis's creators frame it as allowing for free participation in a larger group conversation, but, as we can see, it is an extreme form of what Graan has called "discursive engineering" (2022a). Graan uses the term to complicate ideologies of free and voluntaristic speech by pointing out how marketing practices pervade the public sphere in North Macedonia. He argues that discursive engineering "combines practices designed to enhance, but also to constrain, the public circulation of discourse" (2022a, 302). Although Polis is a very different case, discursive engineering is a helpful concept here, and even more literal than in the Macedonian case. Limiting responses to agree/disagree/pass votes may increase participation, but it also turns human discourse into something that only a machine can analyze.

This conception of consensus as probabilistic data reduction and democracy as computational technology is even more explicit when Polis's founders speak to other programmers. In a presentation to this group, Megill's co-founder, Christopher Small, began with an explanation of Athenian democracy (Small, 2018). Polis's name comes from the Greek word for city-state, and both Megill and Small frequently speak in interviews about Athens as the first place to systematize democracy. Yet CompDem's version of this story differs from most historians' accounts of democracy's Athenian origins because Small and Megill's narrative frames democracy as an information processing technology. In his presentation, Small showed a diagram of Athenian democracy, c. 500 BC, with a PowerPoint slide dominated by a complex diagram, which is also available on Wikimedia Commons.⁷ There, the creators are credited as Mathieugp, WartDark, and Laurent Henry, technologists interested in democratic forms. Small had captioned the slide, "UML, anyone?" a comparison he also made verbally during his talk, pointing out how Athenian democracy, at least as presented in the diagram, looks a lot like a Unified Modeling Language (UML) diagram, which is commonly used to represent systems in software design. In this narrative, democracy has been based on computational models from its beginnings (cf. Latour, 1992).

These comparisons run throughout CompDem's public statements. For example, Small and his colleagues describe elsewhere that modern democracies' reliance on political parties and referenda with a binary choice are examples of "clustering or reduced dimensionality" (Small et al., 2021, 17). In making these comparisons, Small and his colleagues make democracy into a computational problem that can be solved through the right technological aids. This framing comes out clearly when Small speaks about limitations to deliberative democracy, like the issue of scaling, as computational constraints. After comparing Athenian democracy to UML, he went on in his presentation to talk about cybernetics, saying:

One of the concepts that we can get from this which I think is potentially valuable is the idea of a signal problem, the challenge of signals getting enough information into a system to make good decisions. So, the first part of this here is that, right, it's difficult to gather and synthesize information from lots of people, but it's also difficult to come to an agreement when different signals are pointing in different directions. So these are sort of the two fundamental signal problems of democracy and governance.

Small sees each individual's political opinion to be made up of many different signals, often pointing in different directions, which is a challenge for democracy-as-technology. As he went on to explain, in electoral democracy, you vote for one person who most closely aligns with all of your political opinions. It is nearly impossible to vote for someone you completely agree with because humans have what CompDem calls a "high dimensional opinion space," another term borrowed from machine learning, meaning that any one individual's perspective contains a wide array of nuanced opinions on various issues. When you bring together many individuals, each with their own high-dimensional opinion space, it becomes

even more difficult for a government system to be able to incorporate all of those different signals.

High dimensional spaces are also key to how machine learning algorithms work through vector-based word embeddings (Castelle, 2024; Lamoureaux et al., 2025). In this view, the technology behind Natural Language Processing is the perfect fit for democracy today. Small asked his audience of programmers, “What would democracy look like if it were invented today?” After he has presented the history of democracy as a progression through technologies, the answer became obvious: democracy would work like an LLM, where you can input a crowd, with all their messy opinions, turned into matrices, and get some sort of political meaning, or a general will, as an output. This is markedly different from a Habermasian public centered on discussion among its individual members, who together arrive at consensus, or public opinion. Yet it does resonate with a Rousseauian general will, with the idea that computation can help to surface points of common ground present beneath seeming disagreement. While Rousseau understood the general will to be achievable in small-scale, relatively homogenous societies like Geneva, or through guidance from public authorities (Rousseau, 1762 [2003]), CompDem is instead trusting algorithms. Small makes this explicit later in his presentation, calling consensus the “secret sauce” that “brings coherence to a discussion.” Instead of something jointly created, consensus is assumed to be there underlyingly, but political antagonisms keep people from seeing it. Small presented the example of how Polis has been used in Taiwan, concluding:

When you sort of take out all of the political manoeuvrings and, uh, backroom dealings and horse trading, et cetera, um, there are actually pieces of um, of the kind of collective conscious [sic] that can bring different groups, different sides of an issue together.

For Small, political antagonism is not constitutive of democracy but antithetical to it, obscuring underlying agreements. When used to frame democracy, computational models neutralize power struggles, where disagreements are only a signal problem that can be solved with the right technology, part of a general tendency to reduce social complexity to computational logic (Columbia, 2009). In this case, while Polis claims to “output meaning,” its outputs more closely resemble what Kockelman calls information, or the “enclosure of meaning” (2013, 115). In this sense, Polis has similarities to Habermas’s conception of public opinion, as something that constrains the messiness of human interaction to make it something that is formalized and context-independent. Yet that product is even more constrained, and even further removed from human interaction.

EXPERIMENTS WITH LARGE LANGUAGE MODELS

Returning to de la Rive Box, the head of the Dutch organization mentioned in the introduction, we can now better understand his frustrations with Polis. While he saw consensus as something deliberately achieved and grounded in specific, material concerns within his city, consensus for Polis is a computationally surfaced pattern. CompDem alternates between the verbs “build,” “surface,” and “find” when it discusses consensus, but regardless of the word choice, it always assumes the consensus is already “out there,” and does not need to be produced. CompDem may argue that this better captures the complexity of human opinions, but for de la Rive Box, Polis’s “consensus” is impossible to discern without in-person discussion or denotational coherence between statements. While I have not found anyone else who publicly shared his frustrations, CompDem does acknowledge that interpreting Polis’s findings can take a lot of time. De la Rive Box’s solution was to have an in-person discussion about the Polis statements, but that would not be possible with a group

at the scale Polis's developers are imagining. Some supporters of LLMs would also say that the problem is not Polis's algorithm, but that humans cannot interpret what Polis produces. This is just one of many examples of how language models are creating new understandings of language and truth, where a faith in probabilities and pattern recognition allows people to imagine machines are capable of finding order and meaningfulness that humans cannot (see Handman, 2025). *The New York Times* journalist Peter Coy, for example, has suggested in an opinion piece about the potential of AI to save democracy that LLMs “can find surprising points of agreement among seeming antagonists” (Coy, 2023). The points of agreement that AI models find may be so surprising that humans cannot even recognize them without help. From this perspective, public opinion is already out there in the aggregate of individual opinions, but we need LLMs to surface it and put it into a referential statement that humans can understand.

Since Polis's creators see democracy as a technology best-suited to machine learning techniques, it is not surprising that they are interested in the potential of integrating LLMs. CompDem has recently run a series of experiments with Anthropic's Claude. CompDem is adamant that LLMs should not replace human participation, and distance themselves from Ion, a Romanian chatbot designed to tell the government “in real time of Romanians' proposals and wishes” (Agence France-Presse, 2023). They were instead interested in reducing the human workload of managing Polis conversations, allowing them to scale to include more participants and produce results more quickly, and in giving participants stylistic help with writing opinion statements (Small et al., 2023). CompDem's experiments used data from a 2018 conversation from Bowling Green, Kentucky, where residents had been asked, “What do you believe should change in Bowling Green/Warren County in order to make it a better place to live, work, and spend time?” While their study included several experiments, I focus in the rest of this article on their attempt to prompt Claude to write consensus statements. CompDem imagines LLMs to be the perfect tool for democratic deliberation, but the assumptions behind and results of these experiments can tell us more about developers' language ideologies and the new contradictions that would accompany AI-assisted deliberation.

There are several steps before prompting Claude to write consensus statements. First, Claude is prompted to do topic modeling on the submitted comments. Topic modeling is a common type of statistical modeling in natural language processing, designed to discover semantic structures in a text body. It works based on statistical co-occurrence of words. Topic modeling reproduces Polis's fundamental assumption that meaning can be discovered in decontextualized texts, not understanding that denotational textuality relies on interactional textuality (Silverstein, 1976; Urban & Silverstein, 1996). Topic modeling also differs from older approaches that separate denotation from interaction, like formal semantics (Castelle, 2024). From those topics, developers then prompted Claude to summarize participants' comments, including how much agreement and disagreement there was about the topics Claude had identified. Claude's output was short paragraphs. The developers were concerned that those summaries might be incorrect, so they imagined a way for participants to check them in real time. Instead of adding a step to the Polis process where humans check summaries, as other LLM experiments in deliberative democracy have done (e.g., Tessler et al., 2024), they fed these AI-generated summaries back into Claude, prompting it to use the summaries to “generate novel comments that are likely to get consensus” (Small et al., 2023, 43). Unlike the longer summaries it had produced first, Claude's output here was supposed to look like statements written by humans. Developers imagined that other participants could be asked to agree, disagree, or pass on them in the same way they would for a human-written comment.

Developers thought of these LLM-produced statements to be not only a strategy to check Claude's summaries, but also as a way to make participants “feel understood” (Small

et al., 2023, 15). CompDem diverges further from a Habermasian kind of public sphere, where individuals are supposed to be able to set aside their private, emotional selves to achieve consensus through rational, depersonalized discussion. Instead, CompDem is influenced by psychologist Marshall Rosenberg's Nonviolent Communication, which stresses acknowledging and respecting where each side of a conflict is coming from as you look for points of common ground. Nonviolent Communication is used in some other deliberative democratic initiatives and usually requires a facilitator who works like a therapist to bring participants to points of agreement. Based on these influences, CompDem hypothesizes that reflecting participants' perspectives back at them will allow "points of common ground or consensus [to] surface faster" (Small et al., 2023, 15). Consensus here is again something that already exists and just needs to "surface," instead of being created through rational discussion. LLMs are not being imagined as solely putting computational patterns into words, but also about *persuading* human participants that those words are consensus, giving them a sense of recognition and of having been heard and understood. This is a fundamentally individualizing model of public formation, where each participant is given a sense of recognition for their own, individual opinion within the whole.

These LLM-generated contributions turn collective computational patterns into individualized voices, where general summaries become statements that circulate on the platform in the same way individual opinion statements do. The LLM poses as an individual, but their words of course do not represent an individual mental state. They are instead strings of words that they predict participants will vote agree on. Polis's creators already encourage human moderators to intervene in similar ways, submitting statements to a conversation that will help them to determine that they understand emerging points of agreement, something de la Rive Box preferred to do in more personal, face-to-face conversations. Yet the study's authors were concerned about the ethics of using LLMs to do this, specifically whether human participants should be told that some of these other "individuals" in their conversation aren't humans. While they state that, ethically, human participants must be informed that a conversation includes LLM-generated content, these statements would display just like human-generated ones so that humans would vote on them in the same way that they would vote on other human statements.

Claude's responses to the prompt to generate consensus were the following (Small et al., 2023, 15–16):

1. Our infrastructure and public amenities could use continued improvement and investment. Things like roads, parks, transit, etc. enhance quality of life for all.
2. We should advocate for better urban planning and transportation options. Making the community more livable, walkable, and connected would benefit everyone.
3. While we may disagree on some issues, we all want a safe, supportive, and thriving community. Compromise and understanding can go a long way.
4. Access to healthcare and economic opportunity is important for any community. We should work to enable wellbeing and prosperity for as many as possible.
5. Technology is constantly changing, but we should ensure changes benefit and connect our community, not isolate us. Education and connectivity for all ages matter.
6. There is always room for improvement, but an appreciation for what we have and the diversity in our community is so important too. Our differences make us stronger.

These statements have many similarities in structure. Each is made up of two sentences, the first of which almost always includes a first-person plural pronoun, while the second sentence is a generic statement about the general good for "everyone," "all," or similar. In (4), the order of the two sentences is reversed, while each sentence in (6) includes both elements. These two elements make these statements not only about interventions in the

community, but also a way to perform a collective group. The bounds of who belongs within that group do not seem to be limited to Bowling Green residents, as much as they are any rational, liberal individual. In fact, the generic statements that usually follow the first-person plural sentences sometimes quite explicitly refer to people beyond Bowling Green, for example “any community” (4). Even when this wider group is less explicit, many include verbs in a gnomic present tense that express general facts. While the non-authorship of the generic statements LLMs create (Gershon, 2023) is complementary to the emergence of a shared “voice of the people,” these statements are potentially *too* generic, no longer distinct to Bowling Green. While these comments might receive many agree votes because they are so general, they also may provide an example of how the generic is not the unmarked (Gershon, 2023; MacLochlainn, 2022) but could come to index an LLM posing as a human participant. MacLochlainn has argued that, while the generic can be a template crucial to how we make meaning in the world, it can also be discarded, like muzak or elevator music (2022, 17–19). Just as muzak is backgrounded and even derided for being clichéd and inane, a human seeing these Claude-generated consensus statements might reject them for lacking specificity. On the other hand, as MacLochlainn points out, Muzak may be designed to be backgrounded, but it was never there just to fill silence in an elevator. It has been linked to psychological and behavioral studies since the 1940s, believed to increase worker efficiency, and faced accusations of “brainwashing.” This more sinister side might also emerge in LLM-assisted deliberation, where there would appear to be a lot of agreement within the group because there is a high proportion of similarly vague comments. It would also reinforce the LLM’s training to write statements that make human participants “feel understood,” as discussed above. Messeri and Crockett (2024) have cautioned that using AI in research risks “illusions of understanding” when people think they understand more than they actually do because they trust the AI. This tendency, coupled with these inoffensive, generic comments, could lead to similar illusions of consensus, where a group thinks it is in agreement, but that agreement lacks any substance that could be turned into action.

Stylistically, Claude’s statements are not only too generic and too inclusive, but it is also important to note that they all describe a collective “we” that takes on an active role in relation to society. Several of the “we”s in the generated statements are followed by the modal *should* and an action: “we should advocate” (2), “we should work” (4), “we should ensure” (5), which puts the responsibility for change onto the actions of the imagined public itself. The public invoked here is making demands on itself as much as the government, declaring that it is responsible for the changes that it wants to see. Meanwhile, when there are suggestions that some other actor is responsible, as in the call for improved infrastructure in (1) and making the community more connected in (2), there is no clear agent who is responsible for these changes, and the modal verb is no longer about obligation (“should”), but possibility (“could,” “would”). The study’s authors were not able to test these statements on Bowling Green residents, so we cannot know if they would have been widely agreed with, or if the emphasis on personal responsibility would have made participants “feel understood.” Strikingly, the authors do not discuss the denotational content of these statements. They never mention that many of these points do not easily translate into concrete actions the local government could take. Instead, they focus on how the statements “capture the typical *tone* and *style* of consensus statements [they]’ve observed in real conversations” (Small et al., 2023, 16, emphasis in original). This similarity would not be surprising, as many have noted LLMs’ skill at identifying and replicating genres and styles (Gershon, 2023), not only textual but also musical improvisation styles (Wilf, 2023). Yet we would imagine that human participants would also be looking for similarities in denotational content, especially who is assumed to have agency and responsibility for making changes.

There were no moderator-produced consensus statements in the Bowling Green conversation that we could compare Claude’s statements to. In the data available on github

from other conversations, there is variation in moderator-produced consensus statements, from collective obligation, ("We should recognise and value our indigenous ecosystems for their contribution to global biodiversity"), to general declarative statements ("The biodiversity system needs more collaborative decision-making that involves all affected parties"), to prompts that were more about gathering demographic meta-data than achieving consensus ("I am 55 years or older").⁸ However, we can compare these Claude-produced statements to the human-produced ones that received many agree votes from Bowling Green participants, and they look quite different.⁹ The top 10 statements with the highest proportion of agree to disagree votes were the following:

1. The area of downtown and the riverfront should continue to be developed for entertainment and recreation for both day and night activities. (144 agree, 3 disagree)
2. More citizens need to be involved in the future development of BG. Too many developers who benefit financially are calling the shots. (143 agree, 4 disagree)
3. Impact on traffic flow should be considered with any proposed new developments. (656 agree, 20 disagree)
4. Competition for cable companies and an end to heir [sic] price gouging. (446 agree, 14 disagree)
5. We need a fiber option for internet. (214 agree, 7 disagree)
6. Bowling Green needs more competitive internet rates. (481 agree, 17 disagree)
7. More choices when it comes to the internet. BGMU has been offering service to businesses for a while; they should expand to offer to residents. (606 agree, 25 disagree)
8. Bowling Green needs more competitive cable rates. (550 agree, 20 disagree)
9. City/County officials with connections to a project or proposal should abstain from voting on the project or proposal. (509 agree, 23 disagree)
10. Traffic flow needs to be improved throughout Bowling Green, especially on Scottsville Road. (612 agree, 29 disagree)

If we think of different kinds of publics as discursive achievements (Yeh, 2018), we can see Claude's statements as creating a neoliberal public of a citizenry who looks out for its own interests, while these human-produced statements suggest a public that is more concerned about the obligations of local government. This is visible in the content and style of the popular statements during the Bowling Green event. Statements like (2) that "too many developers who benefit financially are calling the shots," (4)'s accusation of cable companies' price-gouging, and (9)'s concern about potential corruption among City/County officials suggest that the participants in this Polis conversation had serious concerns about how decisions in Bowling Green get made, and they were pushing for a more transparent and democratic process. These themes, as well as comments about specific traffic issues and complaints about cable and internet rates, were highlighted in the official report on the Polis conversation, and in media coverage about it (e.g., Karaganis & Ford, 2018). However, these conflicts are completely smoothed over in Claude's consensus statements.

Stylistically, modals are even more present in these human-written statements, but they have different agents and always highlight obligation or necessity; only one statement (4) does not include "should" or "need." There is also only one widely agreed on statement that includes a first-person plural pronoun (5). Instead of sets of actions we, the public, should take, the responsibility for the change is more or less explicitly placed on agencies within the municipal government. Many of these statements are passive: "the area of downtown and the riverfront should continue *to be developed*" (1), "more citizens need *to be involved*" (2), traffic flow "should *be considered*" (3), and "needs *to be improved*" (10). In the active sentences, the agent is never a collective we, as in Claude's statements, but instead BGMU (Bowling Green Municipal Utilities) or City and Council Officials. Less popular statements

did include things that a general “we” should do, but many of these focused on specific tasks, like planting more trees, prioritizing kinds of downtown development, or encouraging people not to drive drunk, not broad statements about shared values.

These differences can be explained by the fact that Claude was prompted to “generate novel comments that are likely to get consensus” based on the summary Claude itself had already produced. This summary was already quite generalized; it stated that there was agreement about “traffic, public amenities, media/technology, urban planning, and transportation,” but it did not say what that agreement had been or include specific street names or internet companies. The wording of Claude's new comments then came from Anthropic's general training data, not only the specific comments from Bowling Green residents. It is beyond the scope of this paper to know whether these examples are representative of a larger trend where Claude produces statements that prioritize citizen obligations over governmental responsibility. However, these statements do raise the question of whether human participants really would see AI-generated statements to reflect their own opinions back at them, as the study's authors anticipate, let alone whether they would accept these statements as “consensus.”

CONCLUSION

Polis's attempt to integrate LLMs is not just a failed experiment. It is also a shift in how participatory democratic projects understand consensus, from a discursive achievement to something an LLM can “surface” and put into words that humans recognize and are persuaded to accept. Political conflict like the Taiwanese government's “backroom dealings” and material concerns about corruption in Bowling Green become “noise” that the LLM can filter out. It is no surprise then that de la Rive Box's organization does not recognize what Polis produces. That organization understands consensus as a discursive achievement requiring face-to-face deliberation, and material concerns are not “noise,” but the basis on which collective recommendations to the government can be made.

While CompDem's founders think consensus is a pattern an algorithm can surface, what remains a black box is what consensus is for Claude.¹⁰ This would depend on the statistical associations it developed from training data, which likely included Habermas, but also more probabilistic understandings of consensus. As far as I have been able to find, no one at CompDem questions this, taking consensus to be a self-evident concept that can be used in prompting Claude without further explanation. Claude was never prompted to explain *why* the statements it generated were “likely to get consensus,” but its output was taken at face value. A closer look at this output suggests that consensus for Claude is overly generic statements that have the “right tone” and incorporate a wide variety of perspectives, but, as with all LLM output, there is nothing referential in these statements.

The black box of what consensus is for Claude is simultaneously what gives Claude's output authority for technologists. AI is frequently presumed to be able to unveil truths about the world that humans could not see on their own, from plants and animals (Handman, 2026) to human biosignals (Semel, 2026). “Consensus” becomes an object with the same qualities; it is presumed to exist out in the world and an LLM can help humans to understand it by transducing it into human language (Keane & Nakassis, 2026; see also Gal, 2023; Silverstein, 2003). While this is a departure from a Habermasian deliberative democracy, it is still complementary to liberal ideals of a view from nowhere, belonging to no one in particular. When consensus is a truth surfaced by an LLM instead of emerging through discussions between humans, it is even further abstracted away from embodied individuals. As the inclusion of LLMs in democratic deliberation moves from experiments to actual use, it will be crucial to attend to ways that these presumptions around consensus produce

new ideals around what democracy is and how it is practiced, just as machine learning is already changing ideas of what a perfect language looks like (Handman, 2025). CompDem and Anthropic's experiments suggest that LLMs, under the aegis of "improving democracy," could create conditions in which it is even more difficult for people to make claims on their political leaders and discuss material concerns. Linguistic anthropologists, with their long focus on the discursive practices through which publics are created and maintained, are well-positioned to understand the changes LLMs may bring to democratic publics and problematize the language ideologies through which they are constructed and naturalized.

ACKNOWLEDGMENTS

This paper would never have been written without Rogier de la Rive Box's initial rhetorical question that inspired the project. Thanks also to him for his generous comments on an earlier draft. This paper has further benefitted from substantial and productive feedback from an anonymous reviewer, as well as helpful comments from James Costa, Ilana Gershon, Anna Weichselbraun, Kathryn Woolard, and members of the Language Machines Network. Thank you to Hannah De Mulder for help with the Dutch abstract.

CONFLICT OF INTEREST STATEMENT

The author has no known conflicts of interest.

ORCID

Janet E. Connor  <https://orcid.org/0000-0002-4877-9266>

Endnotes

¹ For a diagram of the Polis process, see Small et al. (2021, 3).

² <https://compdemocracy.org/Polis/>.

³ Het is echter zeer de vraag of het woord consensus hier terecht wordt gebruikt. In onze proef waren er geen inhoudelijke verbanden tussen de stellingen. Toch bleken er consensusrelaties te worden benoemd. Wanneer er geen onderling verband tussen de stellingen bestaat, lijkt ons het woord consensus niet aan de orde.

⁴ De la Rive Box's decision was also due to concerns around data protection, given Polis is an American company.

⁵ Thank you to my anonymous reviewer for flagging that the term "mini-public" sits uneasily with linguistic anthropological understandings of publics as interactional accomplishments. There is also a rich debate about these so-called mini-publics in political theory. See, for example, Fishkin (2018), Lafont (2019), and Landemore (2020).

⁶ All caps used to indicate emphasis.

⁷ Available at: <https://commons.wikimedia.org/wiki/File:Constitution-of-the-Athenians-in-the-4th-century-BC.png>.

⁸ These comments are all from a conversation a New Zealand national newspaper ran to help a governmental agency develop policies around biodiversity. Available at <https://github.com/compdemocracy/openData/blob/master/scoop-hivemind.biodiversity/comments.csv>.

⁹ Polis's full report on this conversation is available at <https://pol.is/report/r2xcn2cdbmrzjmmuuytdk>. All 896 of the comments are available at <https://github.com/compdemocracy/openData/blob/master/american-assembly.bowling-green/comments.csv>.

¹⁰ Thanks to Kathryn Woolard for first pointing this out to me.

REFERENCES

- Agence France-Presse. 2023. "Romania PM Unveils AI 'Adviser' to Tell him What People Think in Real Time." *The Guardian*, March 2, 2023. <https://www.theguardian.com/world/2023/mar/02/romania-ion-ai-government-honorary-adviser-artificial-intelligence-pm-nicolae-ciuca>
- Anderson, Benedict. 2006. *Imagined Communities: Reflections on the Origin and Spread of Nationalism*, Revised ed. London: Verso.

- Applebaum, Anne, and Peter Pomerantsev. 2021. "How to Put out Democracy's Dumpster Fire." *The Atlantic*. April 2021 Issue. https://www.theatlantic.com/magazine/archive/2021/04/the-internet-doesnt-have-to-be-awful/618079/?utm_source=twitter&utm_campaign=the-atlantic&utm_medium=social&utm_content=edit-promo&utm_term=2021-03-09T22%3A46%3A23
- Barber, Benjamin R. 1984. *Strong Democracy: Participatory Politics for a New Age*. Berkeley and Los Angeles: University of California Press.
- Bauman, Richard, and Charles L. Briggs. 2003. *Voices of Modernity: Language Ideologies and the Politics of Inequality*. Cambridge: Cambridge University Press.
- Bourdieu, Pierre. 1979. "Public Opinion Does Not Exist." In *Communication and Class Struggle*, edited by A. Mattelart and S. Siegelau, Translated by Mary C. Axtmann. New York, NY: International General.
- Calhoun, Craig. 1992. "Introduction: Habermas and the Public Sphere." In *Habermas and the Public Sphere*, edited by Craig Calhoun. Cambridge, MA: MIT Press.
- Castelle, Michael. 2024. "Contextualizing High-Dimensional Communication: The Relevance of Linguistic Anthropology for Theorizing Large Language Models." Language Machines Reading Group Presentation, September 23
- Cody, Francis. 2009. "Daily Wires and Daily Blossoms: Cultivating Regimes of Circulation in Tamil India's Newspaper Revolution." *Journal of Linguistic Anthropology* 19(2): 286–309.
- Cody, Francis. 2011a. "Echoes of the Teashop in a Tamil Newspaper." *Language and Communication* 31(3): 243–54.
- Cody, Francis. 2011b. "Publics and Politics." *Annual Review of Anthropology* 40: 37–52.
- Cody, Francis. 2015. "Populist Publics: Print Capitalism and Crowd Violence beyond Liberal Frameworks." *Comparative Studies of South Asia, Africa and the Middle East* 35(1): 50–65.
- Congress of Local and Regional Authorities. 2024. "Participatory Democracy Is Antidote to Populism and Key to Resilient Democracies." Council of Europe. October 18, 2024. <https://www.coe.int/en/web/portal/-/participatory-democracy-is-antidote-to-populism-and-key-to-resilient-democracies>
- Coy, Peter. 2023. "Can A.I. And Democracy Fix each Other?" *The New York Times*, April 5. <https://www.nytimes.com/2023/04/05/opinion/artificial-intelligence-democracy-chatgpt.html>
- Dewey, John. 1970. *The Public and its Problems*. Chicago, IL: The Swallow Press.
- Farrell, Henry, and Hahrie Han. 2025. *AI and Democratic Publics*. New York: Knight First Amendment Institute at Columbia University. <https://knightcolumbia.org/content/ai-and-democratic-publics>.
- Fishkin, James S. 2018. *Democracy when the People Are Thinking: Revitalizing our Politics through Public Deliberation*. Oxford: Oxford University Press.
- Fraser, Nancy. 1992. "Rethinking the Public Sphere: A Contribution to the Critique of Actually Existing Democracy." In *Habermas and the Public Sphere*, edited by Craig Calhoun. Cambridge, MA: MIT Press.
- Gal, Susan, and Kathryn Woolard. 2001. *Languages and Publics: The Making of Authority*. Manchester: St. Jerome Publishing.
- Gal, Susan. 2006. "Contradictions of Standard Language in Europe: Implications for the Study of Practices and Publics." *Social Anthropology* 14(2): 163–81.
- Gal, Susan. 2023. "Rethinking Translation and Transduction." In *A New Companion to Linguistic Anthropology*, edited by Alessandro Duranti, Rachel George, and Robin Conley Riner, 178–93. Hoboken, NJ: Wiley.
- Gershon, Ilana. 2023. "Bullshit Genres: What to Watch for when Studying the New Actant ChatGPT and its Siblings." *Suomen Antropologi* 47(3): 115–31.
- Golumbia, David. 2009. *The Cultural Logic of Computation*. Cambridge, MA: Harvard University Press.
- Graan, Andrew. 2022a. "Marketing Logics and the Politics of Public Spheres: On Discursive Engineering and Enclosure." *Journal of Linguistic Anthropology* 32(2): 301–25.
- Graan, Andrew. 2022b. "Publics and the Public Sphere." In *Oxford Research Encyclopedia of Anthropology*. Oxford: Oxford University Press.
- Habermas, Jürgen. 1984. *The Theory of Communicative Action, Volume 1: Reason and the Rationalization of Society*. Boston, MA: Beacon Press.
- Habermas, Jürgen. 1989. *The Structural Transformation of the Public Sphere: An Inquiry into a Category of Bourgeois Society*. Translated by Thomas Burger and Frederick Lawrence. Cambridge, MA: MIT Press.
- Handman, Courtney. 2025. "The Chatbot's Real Self: On the Archaeology of Artificial Personas." *Journal of Linguistic Anthropology* 36(1): e70040. <https://doi.org/10.1111/jola.70040>.
- Handman, Courtney. 2026. "Animal Translations: AI and the Intelligibility of Non-human Worlds." *The Journal of the Royal Anthropological Institute*. <https://doi.org/10.1111/1467-9655.70132>.
- Hirschkind, Charles. 2006. *The Ethical Soundscape: Cassette Sermons and Islamic Counterpublics*. New York, NY: Columbia University Press.
- Horton, Chris. 2018. "The Simple but Ingenious System Taiwan Uses to Crowdsource its Laws." MIT Technology Review, August 21. <https://www.technologyreview.com/2018/08/21/240284/the-simple-but-ingenious-system-taiwan-uses-to-crowdsource-its-laws/>

- Karaganis, Joe, and Sam Ford. 2018. "Polarization doesn't Have to Be a Self-Fulfilling Prophecy." *Slate*, December 18. <https://slate.com/news-and-politics/2018/12/political-polarization-bowling-green-study-kentucky-immigration.html>
- Keane, Webb, and Constantine V. Nakassis. 2026. "Introduction: Towards a Linguistic Anthropology of AI." *The Journal of the Royal Anthropological Institute*. <https://doi.org/10.1111/1467-9655.70131>.
- Keane, Webb. 2024. *Animals, Robots, Gods: Adventures in the Moral Imagination*. London: Allen Lane.
- Kelty, Christopher M. 2019. *The Participant: A Century of Participation in Four Stories*. Chicago, IL: University of Chicago Press.
- Kockelman, Paul. 2013. "Information Is the Enclosure of Meaning: Cybernetics, Semiotics, and Alternative Theories of Information." *Language and Communication* 33(2): 115–27.
- Kramer, Elise. 2013. "The Rhetoric of Victimhood: Censorship, 'Silencing,' and Identity Politics in the U.S." PhD diss., University of Chicago
- Lafont, Cristina. 2019. *Democracy without Shortcuts: A Participatory Conception of Deliberative Democracy*. Oxford: Oxford University Press.
- Lamoureaux, Siri, Michael Castelle, and Anna Weichselbraun. 2025. "Language Machines: Towards a Linguistic Anthropology of Large Language Models." *Journal of Linguistic Anthropology* 36(1): e70033. <https://doi.org/10.1111/jola.70033>.
- Landemore, Hélène. 2012. *Democratic Reason: Politics, Collective Intelligence, and the Rule of the Many*. Princeton, NJ: Princeton University Press.
- Landemore, Hélène. 2020. *Open Democracy: Reinventing Popular Rule for the Twenty-First Century*. Princeton, NJ: Princeton University Press.
- Landemore, Hélène. 2024. "Can Artificial Intelligence Bring Deliberation to the Masses?" In *Conversations in Philosophy, Law, and Politics*, edited by Ruth Chang and Amia Srinivasan. Oxford: Oxford University Press.
- Latour, Bruno. 1992. "Where Are the Missing Masses? The Sociology of a Few Mundane Artifacts." In *Shaping Technology/Building Society: Studies in Sociotechnical Change*, edited by Wiebe E. Bijker and John Law. Cambridge, MA: MIT Press.
- Le Bon, Gustave. 1897. *The Crowd: A Study of the Popular Mind*. New York, NY: Macmillan.
- MacLochlainn, Scott. 2022. *The Copy Generic: How the Nonspecific Makes our Social Worlds*. Chicago, IL: University of Chicago Press.
- Mansbridge, Jane. 1983. *Beyond Adversary Democracy*. Chicago, IL: University of Chicago Press.
- Megill, Colin. 2016. "Pol.is in Taiwan." *Pol.is Blog*. <https://blog.pol.is/pol-is-in-taiwan-da7570d372b5>
- Megill, Colin. 2022. "Scaling Deliberation: Polis and the Computational Democracy Project." *Rage & Reason: Democracy Under the Tyranny of Social Media*. Hannah Arendt Center for Politics and Humanities at Bard College. https://www.youtube.com/watch?v=b4j_2JDh_Rc
- Messeri, Lisa, and M. J. Crockett. 2024. "Artificial Intelligence and Illusions of Understanding in Scientific Research." *Nature* 627(8002): 49–58.
- Miller, Carl. 2019. "Crossing Divides: How a Social Network Could Save Democracy from Deadlock." *BBC*, October 26. <https://www.bbc.com/news/technology-50127713>
- Mulgan, Geoff. 2017. *Big Mind: How Collective Intelligence Can Change our World*. Princeton, NJ: Princeton University Press.
- Ovadya, Aviv. 2023. "Reimagining Democracy for AI." *Journal of Democracy* 34(4): 162–70.
- Paz, Alejandro I. 2016. "Speaking like a Citizen: Biopolitics and Public Opinion in Recognizing Non-Citizen Children in Israel." *Language and Communication* 48: 18–27.
- Polletta, Francesca. 2014. "Participatory Democracy's Moment." *Journal of International Affairs* 68(1): 79–92.
- Rousseau, Jean-Jacques. 1762 [2003]. *On the Social Contract*. Translated by G.D.H. Cole. Mineola, NY: Dover.
- Semel, Beth M. 2026. "'Everything Is a Signal': Speaking Circuits and Noisy Signs in the Making of Language-Oriented AI." *The Journal of the Royal Anthropological Institute*. <https://doi.org/10.1111/1467-9655.70135>.
- Silverstein, Michael. 1976. "Shifters, Linguistic Categories and Cultural Description." In *Meaning in Anthropology*, edited by Keith Basso and Henry A. Selby, 11–55. Albuquerque, NM: University of New Mexico Press.
- Silverstein, Michael. 2003. "Translation, Transduction, Transformation: Skating 'Glossando' on Thin Semiotic Ice." In *Translating Cultures: Perspectives on Translation and Anthropology*, edited by Paula G. Rubel and Abraham Rosman, 75–105. Oxford: Berg.
- Slotta, James. 2023. *Anarchy and the Art of Listening: The Politics and Pragmatics of Reception in Papua New Guinea*. Ithaca, NY: Cornell University Press.
- Small, Christopher T. 2018. "Closure on the Cyberpunk Frontier of Democracy." *Closure/Conj* 2018. Durham, NC. <https://www.youtube.com/watch?v=2tBVMAm0-00>
- Small, Christopher T., Michael Bjorkgren, Timo Erkkilä, Lynette Shaw, and Colin Megill. 2021. "Polis: Scaling Deliberation by Mapping High Dimensional Opinion Spaces." *Recerca* 26(2): 1–26.
- Small, Christopher T., Ivan Vendrov, Esin Durmus, Hadjar Homaei, Elizabeth Barry, Julien Cornebise, Ted Suzman, Deep Ganguli, and Colin Megill. 2023. "Opportunities and Risks of LLMs for Scalable Deliberation with Polis." <http://arxiv.org/abs/2306.11932>

- Spitulnik, Debra. 1996. "The Social Circulation of Media Discourse and the Mediation of Communities." *Journal of Linguistic Anthropology* 6(2): 161–87.
- Tang, Audrey. 2019. "A Strong Democracy Is a Digital Democracy." *The New York Times*, October 15. <https://www.nytimes.com/2019/10/15/opinion/taiwan-digital-democracy.html>
- Tessler, Michael Henry, Michiel A. Bakker, Daniel Jarrett, Hannah Sheahan, Martin J. Chadwick, Raphael Koster, Georgina Evans, et al. 2024. "AI Can Help Humans Find Common Ground in Democratic Deliberation." *Science* 386(6719): eadq2852.
- Urban, Greg, and Michael Silverstein. 1996. "Introduction." In *Natural Histories of Discourse*, edited by Michael Silverstein and Greg Urban, 1–17. Chicago, IL: University of Chicago Press.
- Warner, Michael. 1992. "The Mass Public and the Mass Subject." In *Habermas and the Public Sphere*, edited by Craig Calhoun. Cambridge, MA: MIT Press.
- Warner, Michael. 2002. *Publics and Counterpublics*. New York, NY: Zone Books.
- Wilf, Eitan. 2023. *The Inspiration Machine: Computational Creativity in Poetry and Jazz*. Chicago, IL: University of Chicago Press.
- Yeh, Rihan. 2018. *Passing: Two Publics in a Mexican Border City*. Chicago, IL: University of Chicago Press.

How to cite this article: Connor, Janet E. 2026. "Computing Consensus: Language Ideological Work in LLM-assisted Deliberative Democracy." *Journal of Linguistic Anthropology* 36(2): e70071. <https://doi.org/10.1111/jola.70071>.