

Polis: Scaling Deliberation by Mapping High Dimensional Opinion Spaces

Polis: escalar de la deliberación mediante el mapeo de espacios de opinión de alta dimensión

CHRISTOPHER T. SMALL (The Computational Democracy Project, Polis Technologies Inc),
MICHAEL BJORKEGREN (Polis Technologies Inc), TIMO ERKKILÄ (Polis Technologies Inc),
LYNETTE SHAW (The Computational Democracy Project, University of Michigan Complex Systems)
& COLIN MEGILL (The Computational Democracy Project, Polis Technologies Inc)

Artículo recibido: 26 de octubre de 2020
Solicitud de revisión: 12 de mayo de 2021
Artículo aceptado: 29 de abril de 2021

Small, Christopher T., BJORKEGREN, Michael, ERKKILÄ, Timo, SHAW, Lynette and MEGILL, Colin (2021). Polis: Scaling deliberation by mapping high dimensional opinion spaces. *Recerca. Revista de Pensament i Anàlisi*, 26(2), pp. 1-26.
doi: <http://dx.doi.org/10.6035/recerca.5516>

Abstract

Deliberative and participatory approaches to democracy seek to directly include citizens in decision-making and agenda-setting processes. These methods date back to the very foundations of democracy in Athens, where regular citizens shared the burden of governance and deliberated every major issue. However, thinkers at the time rightly believed that these methods could not function beyond the scale of the city-state, or *polis*. Representative democracy as an innovation improved on the scalability of collective decision making, but in doing so, sacrificed the extent to which regular citizens could participate in deliberation. Modern technology, including advances in computational power, machine learning algorithms, and data visualization techniques, presents a unique opportunity to scale out deliberative processes. Here we describe Polis, an open source web application capable of collecting and synthesizing feedback from people in a scalable and distributed fashion. Polis has shown itself capable of building shared understanding, disincentivizing counterproductive behavior (trolling), and cultivating points of consensus. It has done this in the context of journalistic and academic research, and directly as part of decision-making bodies at local and national levels, directly affecting legislation. These results demonstrate that deliberative processes can be scaled up beyond the constraints of in-person gatherings and small groups.

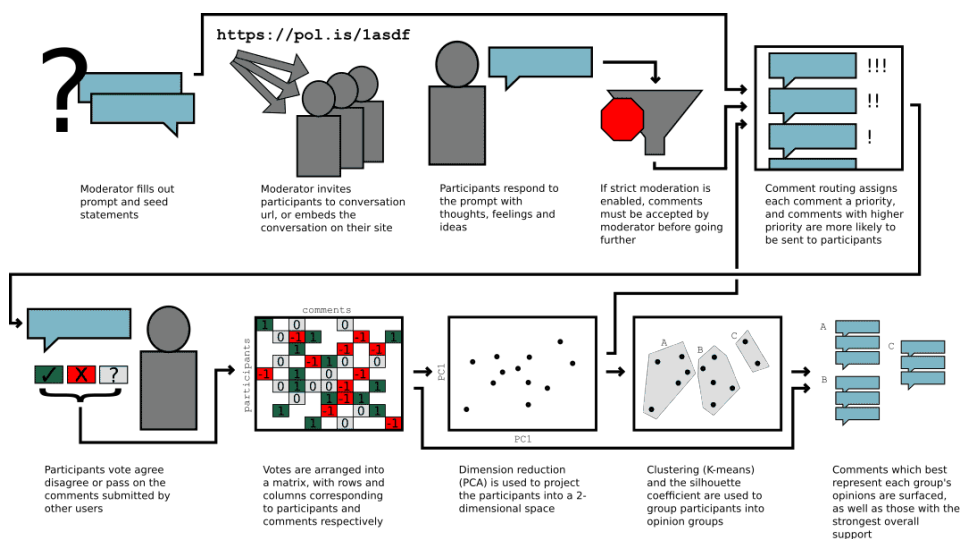
Key Words: deliberation, collective intelligence, unsupervised learning, active learning.

INTRODUCTION

This paper describes the methodology and application of Polis, an open source web application for gathering and synthesizing people’s opinions. Polis combines qualitative and quantitative methods, bridging the divide traditionally spanned by fixed-question surveys and focus groups, respectively. In response to a prompt, participants are able to submit short statements in their own words, ideally expressing a single position or opinion. Participants then have an opportunity to vote (agree, disagree or pass) on statements submitted by other users, presented one at a time, in semi-random order. Using the resulting data, a set of real-time analyses and visualizations are produced which illustrate how the participants break down into opinion groups, what comments distinguish these groups, and where there is rough consensus between groups (see Figure 1 for an overview of the entire process). These methods are closely related to those of the Inglehart and Welzel cultural map (Inglehart, 2005) and DWNOMINATE (Poole, 1985), which attempt to summarize group opinion at scale over time.

Polis was initially developed in 2012, and has since been used by activists and political movements (Ruiz, 2018), academic researchers (American Assembly, 2018), think tanks (Carr, Smith & O’Brien, 2020), journalists (American Assembly, 2018) and governments (Barry, 2016; Hsiao, Lin, Tang, Narayanan & Sarahe, 2018; Horton, 2018; Tang, 2016 and 2019; Miller, 2019a, 2019b and 2019c; OECD-OPSI, 2018; King, 2018 and 2019; Miller, 2020). As a methodology, Polis fits (Bass, 2019) the overarching definition of a *wiki survey* (Salganik & Levy, 2015) by fulfilling the three criteria of being *greedy*, *collaborative* and *adaptive*. Polis is *greedy* in that participants are free to contribute as many votes as they like, sometimes in the hundreds or even thousands of interactions, on an arbitrarily large number of comments submitted by others. Polis is *collaborative* in that participants themselves define the dimensions of the conversation. Polis is *adaptive* in that its model of the opinion landscape informs the order in which comments are shown to participants. Specifically, comments are shown with higher probability if they are more likely to aid in positioning participants in the opinion landscape, build consensus, or are new to the conversation and need a chance to “bubble up”.

Figure 1
Polis Process Overview



A diagram of the participation and analysis process, from framing to conversation and inviting participants, to comment submission, voting and surfacing of opinion groups and representative comments.

Using algorithms to order and scale democratic participation spans the history of democratic practice, beginning, arguably, with the Greek *kleroterion* for selection by lot (Cartledge, 2016: 170). Indeed, the very idea of a congress as a deliberative body implies, by definition, an exercise in collective intelligence (Mulgan, 2017: 181). Efforts to explicitly target improvements in collective intelligence outcomes, specifically with regards to scale, include 20th century pre-internet innovations such as the Delphi Method, pioneered to integrate many experts' opinions into a single prediction (Turoff & Linstone, 1975). Polis fits into a contemporary ecosystem of technology aimed at empowering deliberative and participatory democratic processes and enhancing exercises in collective intelligence (Berditchevskaia & Baeck, 2020). Tools like Loomio (<https://loomio.org/>), Consul (<https://consulproject.org/>) and Decidim (<https://decidim.org/>) provide comprehensive packages for debating, proposing and voting on initiatives, giving feedback on legislation, and participating in budget allocation. Polis, meanwhile, focuses squarely on the challenge of getting

descriptive but actionable information out of deliberation at scale, producing a kind of emergent map, generated by the population being mapped, to be fed to subsequent processes.

Polis's innovation in the space of online deliberation platforms lies in its nuanced revealing of the overall opinion landscape in a way that preserves opinion groups and respects minority dissent, as well as assuming no relationships between various comments, other than that they could be compared. Other platforms which focus explicitly on scaling deliberative processes include All Our Ideas (AOI, <https://allourideas.org>) and Make.org. AOI, which was created as a demonstration of the aforementioned notion of a wiki survey, presents participants with a series of pairs of statements and asks them to choose which they prefer. This generates more information about the overall relative favorability of the available options, making it ideal for deciding between ideas which participants generally agree with, but with different priorities. Make.org, similarly, allows participants to respond to a question with proposals, vote on others' proposals, and elevate points of consensus. Make.org appears to be collaborative and greedy, and has been used in France to deliberate issues with hundreds of thousands of participants, and break down responses into groups or "axes".

Public authorities are increasingly seeking and integrating all kinds of deliberative methods; as a recent report from the OECD states simply: "Public authorities from all levels of government increasingly turn to Citizens' Assemblies, Juries, Panels and other representative deliberative processes to tackle complex policy problems" (OECD, 2020). To date, our most consequential application of the technology in direct connection with government decision-making processes was in Taiwan as part of the vTaiwan platform, which was put together by members of the civic tech community in response to the Digital Minister of Taiwan's call for a platform that the entire nation could use to deliberate issues of national importance (Barry, 2016). Polis was used to deliberate the regulation of Uber and AirBnB, and to address issues such as corporal punishment, online alcohol sales, and the nonconsensual sharing of sexual images (Horton, 2018). It is now part of the successor Join platform (Tang, 2019; Horton, 2018), to which half of Taiwan's citizens are subscribed (Bertelsmann Stiftung, 2020).

1. METHODS

To start a Polis conversation, moderators simply provide a brief description of what sort of feedback they would like from participants, and decide on a handful of settings. They also have the opportunity to “seed” the conversation with comments which help set the tenor, and ensure the very first participants have comments to vote on. People can be invited to participate anonymously by logging in through a social media account (e.g., Facebook or Twitter), or in association with external identity and metadata by embedding the participation experience within the moderator’s web property.

As participants join the conversation, they have the opportunity to vote (agree, disagree or pass) on the comments submitted previously. If they want to mention perspectives that have not yet appeared in the comments, they can submit a new comment. This comment is then sent out to others in the conversation. If the conversation has been set up with “strict” moderation enabled, these new comments will have to be explicitly accepted by moderators before being sent out to other participants. With strict moderation turned off, comments may be sent out immediately, but moderators may choose to remove comments from the conversation if they are deemed off-topic or abusive.

The participation interface can optionally be configured to show a visualization summarizing the results of the conversation. The simpler visualization interface shows just a small number of divisive and consensus comments, with donut charts representing the vote distributions. The more advanced visualization shows participants’ positions in an abstract “opinion landscape” (as described below), and outlines around these participants indicating the presence of opinion groups. The visualization is interactive; by clicking on groups, users can see which comments best separated participants in a given group from the rest of the conversation, in addition to a list of the majority opinions. Finally, a more detailed and interactive report is provided as part of the tool, which allows facilitators to examine the results of the conversation more deeply and surface actionable information.

1.1 The matrix

As participants vote, a *vote matrix* is derived, where rows correspond to participants, and columns to comments.

$$D = \begin{bmatrix} v_{12} & v_{11} & \cdots & v_{1C} \\ v_{21} & v_{22} & \cdots & v_{2C} \\ \vdots & \vdots & \ddots & \vdots \\ v_{N1} & v_{N2} & \cdots & v_{NC} \end{bmatrix}$$

Each value $v_{i,j}$ above corresponds to the vote of participant i on comment j . Agree votes are coded as +1, disagrees as -1, and passes as a 0. Missing values, corresponding to comments the participant in question did not see, are imputed by taking column-wise means of the non-missing values associated with the given comment, a common method of dealing with this issue in downstream analyses (Dray & Josse, 2015). Participants who voted on fewer than seven comments are removed from the conversation to avoid the “clumping up” of participants around the center of the conversation. This number is somewhat arbitrary but tuned as a hyperparameter based on experience with the domain. We will leave discussion of better metrics and functions for deciding whether to keep or drop a participant to a future paper dealing with missing data.

1.2 Dimensionality reduction

Dimensionality reduction is performed using principal components analysis (PCA), a common method in exploratory data analysis (Pearson, 1901). This produces a lower dimensional (in our case 2-D) representation of the data, which can be thought of as a 2D “map” of the opinion space, and is suitable for visualization and further analyses. The particular algorithm used to perform PCA is the power iteration method (Roweis, 1998). This method allows us to use the previous principal component eigenvectors as the starting point for further iterations of the method when there are new votes to process, allowing the method to converge very quickly, and with reduced computational load.

Because participants can respond to different numbers of comments, and given the abovementioned method of imputing means from nonempty column entries of the vote matrix, participants who voted on fewer comments will naturally tend to be projected more closely to the center than those who voted on lots of comments. This results from the first step in PCA analysis being the “centering” of the data matrix, by subtracting the mean of each column from every entry in that column, leaving a row corresponding to someone with no votes with mostly 0 entries. To correct for this, a simple procedure is applied that scales the projected positions of participants by the factor $\sqrt{C/C_p}$, where C is the total number of comments, and C_p is the number of comments voted on by participant p .¹ This value should be 1 for a participant who has voted on all of the comments, and greater than 1 otherwise.

The intuition behind the C/C_p factor is that, as the inverse of the fraction of comments which have been voted on, it can be taken as a proxy or estimate for the inverse of the total possible variance that has been explored via the given participant’s votes. Meanwhile, the square root term has the effect of softening the scaling factor, and was introduced based on testing, which suggested that scaling by C/C_p alone was overly aggressive early on, likely because in effect it makes very strong assumptions based on the first few votes about how participants will respond to later votes.

However, this assumption about C/C_p ends up being invalidated by the introduction of comment routing (the method by which comments are routed to participants for voting, as described later in this section); because comments with high variance (specifically principal component loading) tend to be shown first, the scaling factor described above has the potential to be overly aggressive early on. Future work will look at how taking principal component loadings directly into account might mitigate this issue. Other methods have also been developed for addressing these sorts of issues, which will be evaluated (Dray & Josse, 2015).

¹ Unfortunately, a bug was recently discovered in this method which counts comments that have been explicitly moderated out of the conversation as having been voted on. For conversations with heavy moderation, this has the effect of reducing the scaling factor to close to 1, softening the effect of the correction across the board. This will be fixed in a future release of the software.

1.3 Opinion groups

From this 2-dimensional projection, the system performs a preliminary fine-grained clustering analysis using the K -means clustering (MacQueen, 1967) algorithm (specifically, Lloyd's algorithm; Lloyd, 1982), with K set to 100. These fine-grained clusters serve primarily as a performance optimization to reduce the data payload necessary to update the data visualization when new votes are processed.

These fine-grained clusters then serve as the basis for a more coarse-grained clustering, also using K -means. Here, multiple runs of K -means are performed for values of K between 2 and 5. The K for which the silhouette coefficient (a measure of within-cluster similarity vs. between-cluster dissimilarity) (Rousseeuw, 1987) is optimal is chosen for the opinion groups, which is a common technique for deciding on a number of clusters. However, it is possible for the optimal value of K to switch briefly, especially early on in a conversation when the landscape is shifting more quickly with each vote. In order to avoid frequently changing numbers of opinion groups in the visualization, a smoothing function is applied. A particular K value must be found optimal at least 4 times in a row before that value will be accepted.

1.4 Comment statistics

Once opinion groups have been defined, comments are analyzed for how strongly they represent each opinion group. This *representativeness* metric $R_v(g, c)$ for group g , comment c and vote v estimates how much more likely participants in group g are to vote v on said comment than those outside group g . Letting $N_v(g, c)$ be the number of participants in group g who cast vote v on comment c , and $N(g, c)$ be the total number of votes on comment c within group g , we derive $P_v(g, c) = \frac{1+N_v(g,c)}{2+N(g,c)}$ as an estimate of the probability that a given person in group g votes v on said comment. The 1 and 2 pseudocount terms ensure that the ratio defaults to $\frac{1}{2}$ in the absence of votes. Then the representativeness is defined as the estimated relative odds ratio

$$R_v(g, c) = \frac{P_v(g, c)}{P_v(\underline{g}, c)}$$

Here, \underline{g} is the complement of g , that is, everyone in the conversation *not* in g .

The selection criterion for which comments are to be shown involves looking at the two-property test (in essence, the Fisher exact test; Fisher, 1922). The corresponding Fisher Z -statistic is multiplied by $R_v(g, c)$ to reflect both the estimated effect size and the statistical confidence associated with the effect.

These metrics are computed for both agree and disagree on every comment, and for every group, and a selection procedure is carried out that first attempts to select those comments which are representative for agreement, and if there are none, selects those representative for disagreement.

The groups also inform a *group-aware consensus* metric, which is computed as

$$C_v(c) = \prod_{g \in G} P_{v=a}(g, c)$$

This metric is highest when *all* groups tend to agree with a comment in question, helping to protect from tyranny of the majority and respect minority dissent. Polis's detailed report enables users to order comments by this metric to find points of common ground and rough consensus.

1.5 Comment statistics

In order to take the fullest advantage of users' time, comments are sent to participants for voting in a semi-random order,² probabilistically weighted according to a metric which reflects how likely they are to help place participants in the opinion landscape or build consensus, as well as highlight comments new to the conversation. The specific formulation of this priority metric is as follows:

$$Priority(c) = \left[\frac{P_{v=a}(c) \cdot (1 - P_{v=p}(c)) \cdot (1 + E(c)) \cdot (1 + 2^{3-N(c)/5})}{2} \right]^2$$

Here, $P_v(c) = P_v(G, c)$, as defined above, where G is the set of all participants. Similarly, $N(c)$ is the total number of votes on comment c . $E(c)$ is the

² More precisely, according to a random, but non-uniform distribution.

extremity of comment c , and is defined as the distance from the center of the conversation to a theoretical participant who *only* voted (agree) on c and no other comments.

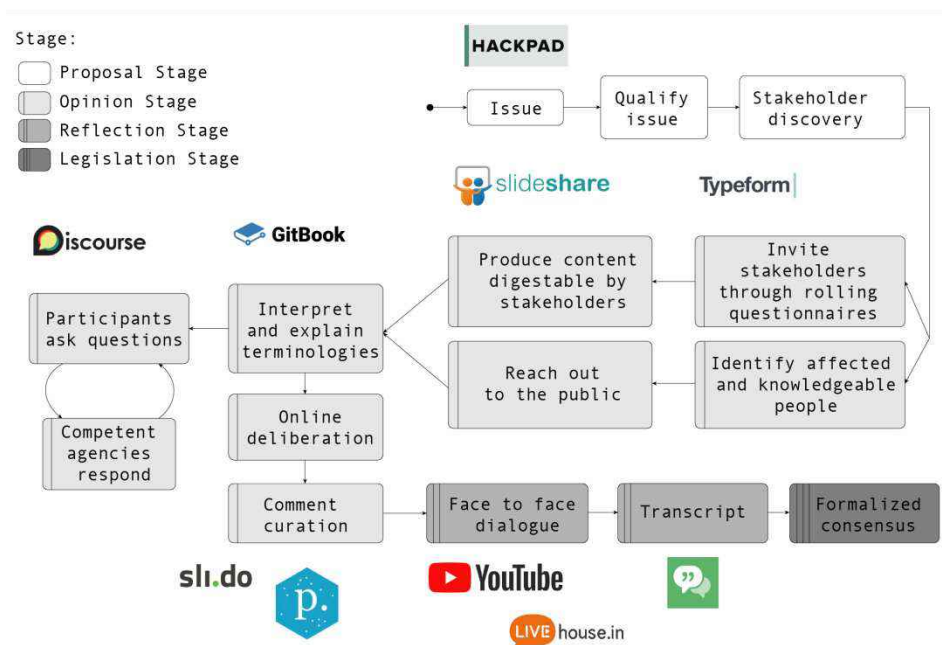
This equation is constructed such that each of the terms in the product are greater than 1 for comments which should be sent to more users, and less than 1, decreasing toward 0, for comments which should not be shown as much. $P_{v=a}(c)$ boosts consensus and shrinks to 0 for comments with no support. The $1 - P_{v=p}(c)$ term goes to 0 for comments which have been mostly passed on. $1 + E(c)$ elevates comments with a high PCA loading, which helps us place participants in the conversation. The final $1 + 2^{3-N(c)/5}$ term starts off at 9 for comments with no votes, and asymptotically approaches 1, the more votes there are for the comment. The outer square term is used to strengthen the effect of the bias toward comments boosted by each of these factors.

When a participant loads the page, or finishes voting on one comment and is ready for another, the probability that the algorithm will send that participant comment c is $Priority(c)$, normalized by the sum of all such values for other comments the participant has voted on.

2. METHODS

Our first opportunity to field test Polis on a large scale in a national policy-making context came in 2015 in Taiwan. Early in the life cycle of product development, the Polis team came into contact with a group of open source civic technologists and activists in Taiwan called g0v (pronounced “gov zero”). Initial experimentation with the platform led by these technologists progressed to facilitate rulemaking in a participatory regulation process called vTaiwan (Figure 2), set up in response to the Sunflower Movement. The facilitators of vTaiwan were interested in ways of understanding large groups of people while reducing the need for moderating the complex, nested discussion threads typical of forums.

Figure 2
vTaiwan Process Overview



A diagram of the vTaiwan policy-making process, from identifying issues and explaining them to the public, to online deliberation facilitated by Polis, face-to-face dialogue and eventual law.

We note that the entire platform and all its methods worked out of the box with Traditional Chinese, since there is no use whatsoever of natural language processing. Polis's algorithms are agnostic to language and even content: the content voted on could be pictorial, audio or video.³ Facilitators used Polis's output to examine the latent space of opinion across hundreds of issues and to instruct policy deliberations.

Here we focus on data from a conversation run as part of the vTaiwan processes examining the legality of Uber operating in the Taiwanese transportation marketplace. Facilitators used Polis to invite Uber drivers, taxi drivers, transportation users and the general public into a shared conversation, allowing

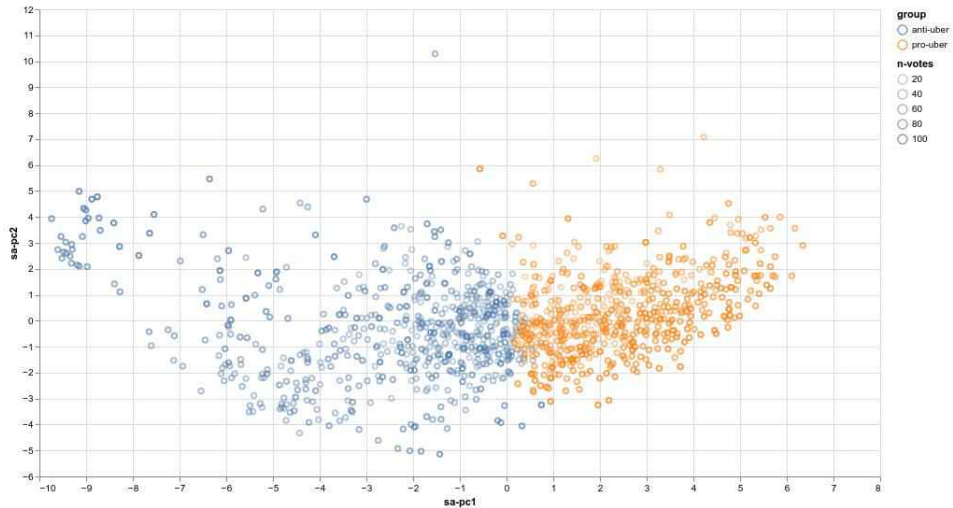
³ To be clear, this is not currently supported by the Polis software as written, but would be trivial to implement given that the algorithm itself is agnostic to content.

them to interact with each other's ideas in the interest of understanding a general topology of the opinion landscape as it existed among those communities.

In the vTaiwan conversation, close to 2000 people participated and 100 people submitted nearly 200 comments. These participants were drawn in from invitations to stakeholder groups, including taxi drivers and Uber drivers, as well as citizen outreach via paid and organic social media posts. Advertised posting on Facebook was specifically carried out in order to recruit a diverse set of participants, in terms of geographic location and gender. After applying comment moderation and removing participants with fewer than seven votes, 98 of these comments and 1238 participants remained. Facilitators stopped recruiting new participants when levels were deemed sufficient for gauging public sentiment, on a par with or exceeding participation levels often associated with traditional public opinion polling. Comments that were moderated out were either unclear, irrelevant, or expressed a sentiment already reflected in the conversation. At the time of this conversation, we did not yet have comment routing in place, and so we generally advised that moderators limit the number of comments more aggressively than we do now, in order to better make use of participants' time.

These participants broke down into two distinct groups: those in favor of Uber and ridesharing apps more broadly, and those opposed to them. This division is well illustrated by the first component of the PCA (see Figure 3), which captured 22% of the variance in the conversation. By contrast, the second principal component captured only 6.2% of the variance, with further components representing decreasing variance. The second component seemed to correlate most strongly with attitudes on regulation, and was largely independent of attitudes toward Uber. The sharp decline in variance explained between the first two principal components is reflective of the fact that in general, people were supportive of regulatory measures, lending less variance overall along this axis. Interestingly, visual inspection of the PCA projection suggests that the participants with the strongest views against regulation tended to be more staunchly either pro-Uber or anti-Uber (and pro-taxi).

Figure 3
PCA Projection of Participants in Uber conversation

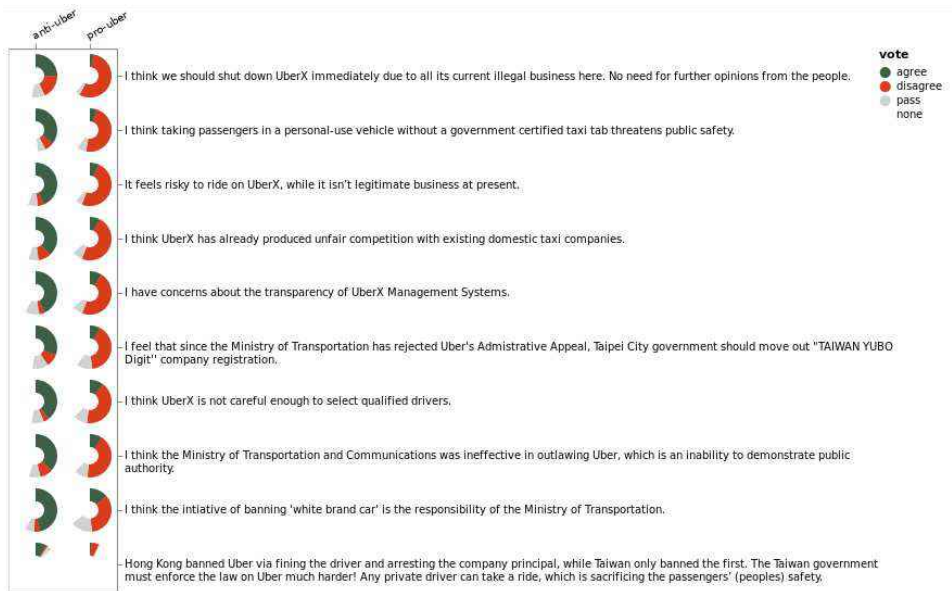


Participants are plotted according to the sparsity-aware corrected PCA projection, colored by K-means assigned opinion group. Participants with fewer votes are less opaque.

Comments representative of the anti-Uber group revolved around concerns with public safety, lax regulations, unfair competition with traditional taxi companies, and a lack of transparency (see Figure 4).

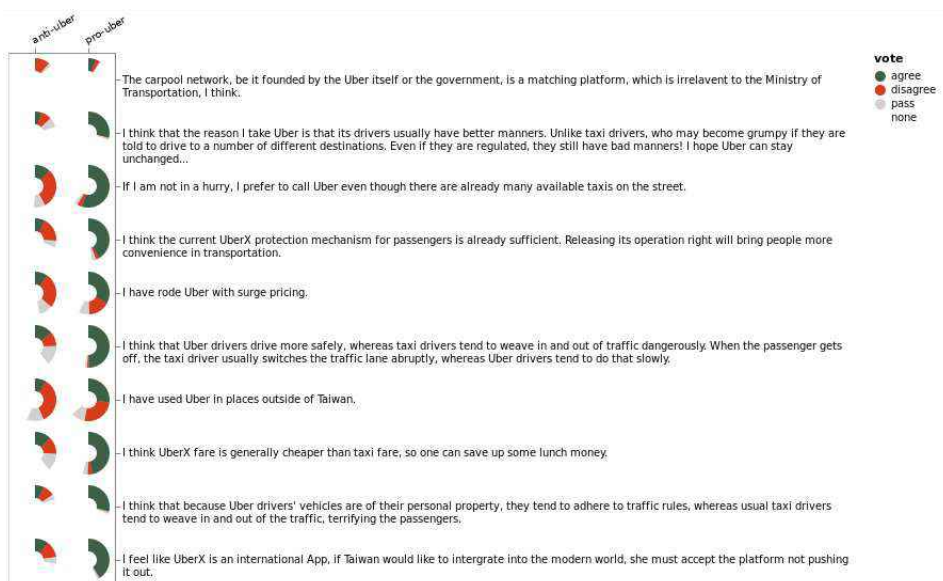
Meanwhile, comments representative of the pro-Uber group revolved around Uber's superior service, its drivers' safer driving practices, and it being sufficiently regulated (see Figure 5). They also reflected an overall higher preference for riding with Uber, and a sense that ridesharing should not be considered in the same category as traditional taxis.

Figure 4
Voting patterns for comments representing the anti-Uber group



Donut charts illustrate the percentage of people in each group who voted agree, disagree or pass. Missing segments of the donut correspond to participants who did not vote on the comment in question. Comments are sorted according to how well they represented the anti-Uber group, from most to least representative.

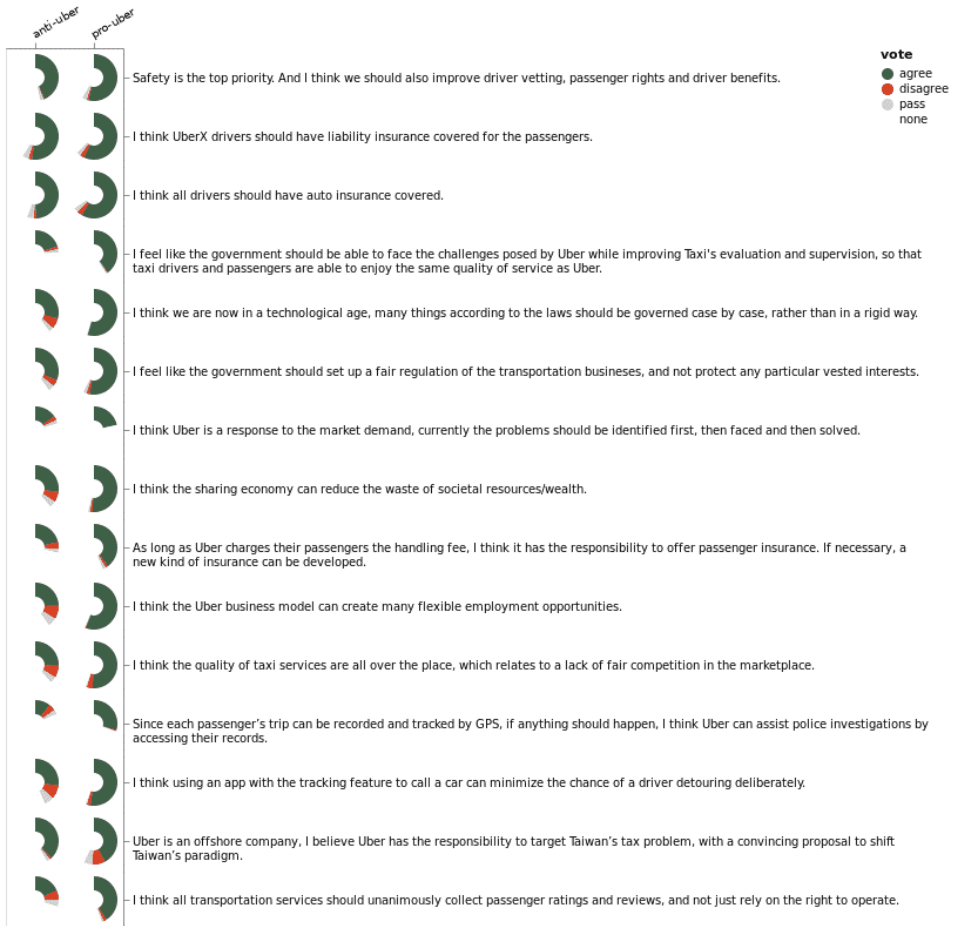
Figure 5
Voting patterns for comments representing the pro-Uber group



Donut charts illustrate the percentage of people in each group who voted agree, disagree or pass. Missing segments of the donut correspond to participants who did not vote on the comment in question. Comments are sorted according to how well they represented the pro-Uber group, from most to least representative.

In spite of these differences, numerous comments found broad support across the conversation (Figure 6). These comments were characterized by thoughtful nuance, and expressed the overall importance of safety, requirements for drivers to have liability insurance, fair regulation, the opportunity for ridesharing to reduce waste, the value of flexible employment opportunities afforded by ridesharing, and inconsistent quality of traditional taxi services. One particularly outstanding comment was translated as “I feel like the government should be able to face the challenges posed by Uber while improving Taxi’s evaluation and supervision, so that taxi drivers and passengers are able to enjoy the same quality of service as Uber”. Nearly everyone, even the taxi drivers, agreed that Uber’s entry into the market was an opportunity to improve the quality of transportation services across the board.

Figure 6
Voting patterns for comments with the most agreement, across the conversation



Donut charts illustrate the percentage of people in each group who voted agree, disagree or pass. Missing segments of the donut correspond to participants who did not vote on the comment in question. Comments are sorted according to how well they represented the group, from most to least representative.

It is worth noting that several of the topics which drew majority support across the conversation were also found to be divisive when framed differently. For example, members of the anti-Uber group tended to disagree that taxi drivers “have bad manners”, but they didn’t dispute that the taxi service could be

improved, and saw Uber as a forcing function for these improvements. Meanwhile, while pro- and anti-Uber groups differed in whether they thought Uber was sufficiently regulated, there was consensus that fair regulations and insurance in particular are important. Moreover, while the groups differed in their responses to whether Uber drivers were safer drivers and the role of regulations in relation to public safety, both groups ultimately saw safety as a top priority. These subtleties highlight the value of the open-ended wiki survey. A traditional poll with fixed questions could have easily missed much of the subtlety that emerged organically from this deliberation, and almost certainly would not have covered an equivalent scope.

3. DISCUSSION

Today, opinion polling and focus groups serve as levers through which journalistic institutions help shape the public's understanding of public opinion, giving citizens' voices an indirect influence in governance. As technology has improved, so has the ability for people to weigh in on issues in real time of their own accord. This has been observed even in platforms not explicitly built for deliberation or democratic engagement, such as the use of Reddit by Podemos (Blitzer, 2014). However, the coherent aggregation that results from intentional meaning-making and platform design has lagged, and left emergent, qualitative dimensions undervalued as an input to policy-making processes.

Moreover, existing structures of representative democracy have proven inadequate in expressing the collective will of the people (Lee, Zhang & Yang, 2017), and have been shown to contribute to political polarization (Yang, Abrams, Kernell & Motter, 2020). Political parties themselves can be seen as a sort of clustering or reduced dimensionality across all possible issues of interest, frequently forcing voters to sacrifice their positions on issues of importance to them, and often not having any candidate at all representing their positions on other issues. While ballot initiatives give voters the right to vote on individual issues, these measures themselves are then presented as binary options, failing to capture the nuanced perspectives constituents might have in relation to the matter at hand. More open, participatory and deliberative processes afford citizens the ability to weigh in on individual issues before measures have been drafted, leading to legislation which better represents the collective will.

Polis is intentionally designed without threads, which reduces the ability of individuals to take a conversation off topic, and with moderation, which allows conversation owners to better utilize people's attention over time. Surfacing consensus emergently — for example, the way both parties emerged as heavily in favor of requiring liability insurance in the Uber case — also nudges the deliberative exercise toward productive outcomes. Polis, as a platform, aims to shift agenda-setting power away from those running a conversation or survey toward those who are participating in it. Its approach rests on a fundamental shift away from a focus on responses to pre-established statements and closed form questions, to the creation of interpretable representations of the opinion spaces that are revealed organically through groups' own deliberation processes.

Part of the inspiration for Polis visualizing the resulting groups back to the participants themselves lies in the theory and practice of nonviolent communication (Rosenberg, 2003). We set out with the hypothesis that visualizing common ground in real time, in the context of mirroring identity back to those participating via group identities, would be a powerful way to induce progress on deadlocked policy issues.

Polis has shown that the same class of methods which OKCupid and Netflix (Madrigal, 2014) use to match users with like-minded people and content can instead be used to help us understand each other and build consensus in a policy-making setting. Indeed, algorithms such as collaborative filtering intentionally and consciously leverage collective intelligence (Segaran, 2007) but, perhaps predictably, have been exploited most powerfully by industry rather than the public sector. As Carr, Smith & O'Brien (2020) note with regard to the deployment of the tool: "The key difficulty though, as with all democratic innovations, will be making the case for opinions generated here to become policy and legislation".

3.1 Sociological perspective

While Polis was initially developed with the objective of creating a more effective digital platform for the deliberation processes of large and diverse groups, it may also be used in lieu of, or in addition to, more traditional survey and focus group approaches. In particular, by spanning the divide between qualitative and quantitative methods, Polis offers researchers a powerful tool to gather quantitative open-ended opinion data.

All measurements inherently reduce a high dimensional reality into a handful of observations, reflecting a lower dimensional set of features. Selecting

which of these features to use in a measurement would, ideally, be motivated by unbiased theoretical understanding of the underlying reality being measured (Lakatos, 1978). In practice, however, the universe of features that might provide a basis for a measurement is significantly bounded by the technological limits of what can be observed or analyzed using the available instruments.

In attempting to measure social or political phenomena, the risk of doing even further “violence to reality” (Weber, 1969) than is already committed by the pragmatic limits of instrumentation are multiplied by unobserved failures in researchers’ own reflexivity (Bourdieu & Wacquant, 1992) and the uniquely performative dependencies within social systems by which measurements of social life can, over time, lead to an over reification of the relative importance of the dimensions they assess (Berger & Pullberg, 1965).

The difficulties inherent in reducing the complexity of social phenomena into interpretable and analyzable observations are especially pronounced in the study of political ideologies and opinions. Historically, they have led social science researchers into a deeply unsatisfying choice. One option is to rely on qualitative methods that use observations which better preserve the dimensionality of a particular opinion space but do so at the cost of limited interpretability and generalizability. The other standard pathway is to employ quantitative methods that produce more straightforward and generalizable interpretations but which require the aggressive reduction of the dimensionality of opinion spaces in ways that can be highly biased by researchers’ survey design or misleadingly shaped by their choice of formal analysis techniques.

Through the platform’s leveraging of contemporary advances in user interface design, data visualization, and machine learning, *Polis* offers a new middle way for researchers interested in preserving the complexity of people’s opinion landscapes, while still enabling the collection of data that are amenable to computational or quantitative analyses. Because of the ways in which sociological research, and survey methods more broadly, contribute to shared understanding of public opinion, this application empowers the public to shape a more nuanced understanding of itself.

3.2 Challenges and limitations

The method of gathering opinions as implemented in the *Polis* platform has a number of limitations. As a natural consequence of the platform being built for scale, the method does not handle small numbers of participants or comments very well. The method has a lower bound of dozens of people and

tens of comments. In general, institutional applications of Polis have been able to recruit more than enough participants to generate useful analyses of public opinion. However, in the context of smaller communities or organizations using the tool for decision making or internal feedback, participation levels can be important to producing meaningful results, and failure to achieve these thresholds has been observed at the margins. Nonetheless, even a small but committed group can successfully make use of Polis, as evidenced by reports we received of a small Dungeons & Dragons gaming group which used the tool to collaboratively decide how to run their campaigns.

Another limitation is that overly-narrow prompts can lead to conversations lacking in dimensionality. Specific binary questions like “Should we change the time for this meeting?” are inappropriate for the method, as are leading questions like “Should we ban all guns?”. A good prompt can be as general as “What problems are you facing?”, given a participant group that understands the context of the question, such as a group of employees.

Similarly, Polis tends not to perform as well with questions which attempt to rank or prioritize a set of approaches to something, such as “What color should we paint the bike shed?”. Delivering a definitive global ordering of ideas is one of the strengths of All Our Ideas relative to Polis. Future work will look at adding *importance* to the responses, so that participants may indicate that “This comment is important to me”, in addition to merely agreeing or disagreeing.

Despite the comment routing system prioritizing new comments, Polis conversations remain sensitive to the time ordering of participants’ involvement. Participants who arrive early will not see comments submitted by those who arrive later unless they return, and those who arrive later will have more opportunities to vote because of the greater number of previously submitted comments. Conversely, comments submitted by earlier participants are more likely to be seen and interacted with, and thus will also be over-represented in terms of vote density. Work to further improve upon this issue will include adding daily notifications which let participants opt in to receive an email if there are new comments.

The system does better with light moderation to maximize participants’ time on task (removing nonsensical or obviously off-topic statements), but this can introduce bias. Future work will explicitly investigate how crowd moderation might be implemented while mitigating the risk of participants using moderation tools to suppress others’ voices. At present, however, the platform relies

on the presence of benevolent moderators who proactively seek to avoid introducing biases by removing statements on certain topics. In this and several other regards, facilitation training for discussion moderators and facilitators remains a vital part of the Polis process.

In general, Polis produces better results, the more people participate. However, as the number of comments in a conversation grows, users become more likely to submit what are effectively duplicate comments. Ideally, as a person starts typing out their comment, the interface would show comments semantically similar to the idea they wanted to express (similar to StackOverflow), so they have the opportunity to vote on those comments instead of submitting a new one. This would help Polis conversations scale even further without proliferating the comment space unnecessarily, and ease the job of moderation. However, there are challenges to consider in relation to how this ability affects voting and participatory behavior, and more study is warranted.

As mentioned in the Methods section, Polis has different visualization settings: off, simple (top majority comments), and full (opinion space visualization, and comments broken down by group). It has been found that while some (typically younger, more educated, and more technologically savvy) audiences are very engaged by the full visualization, it is off-putting for other audiences. For these situations, the simpler visualization is a better fit. On the other hand, in some cases, facilitators may want participants to be able to respond without the potential bias of a visualization influencing their behavior. However, no study has yet examined the precise effect of each of these participation modes on behavior in the conversation.

The current approach to dealing with missing data by adjusting the projection of participants is somewhat ad hoc, and based on assumptions violated by comment routing. Moreover, this method does not account for scaling the effect of comments with few votes. There are many approaches to dimensionality reduction that deal with these limitations in more principled ways, which we hope to explore and eventually employ with Polis (Dray & Josse, 2015).

It is also the case that the current approach to comment routing is somewhat ad hoc and heuristic. Currently the “decay rate” of the factor which highlights new comments, giving them a chance to “bubble up”, does not scale with the total number of comments in the conversation, although it probably should do so. Building a better understanding of missing data (as described above) may dovetail well with this. Many challenging questions might be asked about how each of the factors should be weighted relative to each other, and these considerations will additionally have to account for the presence of “importance” data

from participants, once this is added. Building a systematic framework for thinking about these decisions and their implications is an active area of research for us.

4. CONCLUSION

By striking a balance between the quantitative approaches historically associated with survey research and the type of content usually extracted from qualitative methods such as focus groups, Polis has sought to offer a middle road between the former's high generalizability and the latter's richness of insight that directly supports the public's ability to build shared understanding and surface points of consensus. These features have helped decision-making bodies break through political gridlock, resulting in successful legislation at the national level. In a world of growing political polarization, we believe that these and other participatory and deliberative methods will help the public better work together toward making decisions in the interest of the common good. We maintain that the application of machine intelligence to public deliberation holds great promise in reimagining decision making in public institutions.

5. ACKNOWLEDGEMENTS

We would like to thank Audrey Tang, Elizabeth Barry and Carl Miller for documenting the vTaiwan process for audiences outside Taiwan. We are also grateful to Olga Cheng, Cassandra Lim, Tzu-Chi Yen, Bonnie Chan and Arne Brasseur for their assistance in translating comments into English, C sar Hidalgo for his early feedback in the preparation of this manuscript, and Julien Cornebise and Deep Ganguli for their feedback during our final revisions.

BIBLIOGRAPHY

- American Assembly at Columbia University (2018). The Bowling Green Civic Assembly. 30 June 2018.
- Barry, Liz (2016). vTaiwan: Public participation methods on the cyberpunk frontier of democracy. *Civicist*, 11 August 2016.

- Retrieved from: <https://civichall.org/civicist/vtaiwan-democracy-frontier/>.
- Bass, Theo (2019). Crowdsourcing for democracy using Wikisurveys. *Nesta Blog*, 28 January 2019. Retrieved from: <https://www.nesta.org.uk/blog/crowdsourcing-democracy-using-wikisurveys/>.
- Berditchevskaia, Aleks and Baeck, Peter (2020). The Future of Minds and Machines: How artificial intelligence can enhance collective intelligence. *Nesta*, 10 February 2020. Retrieved from: <https://www.nesta.org.uk/report/future-minds-and-machines/>.
- Berger, Peter and Pullberg, Stanley (1965). Reification and the Sociological Critique of Consciousness. *History and Theory (Wiley)*, 4(2), 196-211.
- Bertelsmann Stiftung (2020). Digital Democracy – What Europe can learn from Taiwan. Web discussion with Audrey Tang. Retrieved from: <https://www.youtube.com/watch?v=I5uMMKaSbSw>.
- Blitzer, Jonathan (2014). In Spain, Politics via Reddit. *The New Yorker*, 7, October 2014. Retrieved from: <https://www.newyorker.com/tech/annals-of-technology/spain-politics-via-reddit>.
- Bourdieu, Pierre and Wacquant, Loïc J.D. (1992). *An Invitation to Reflexive Sociology*. Chicago: University of Chicago Press.
- Carr, Harry, Smith, Josh and O'Brien, Toby (2020). Polis and the Political Process. *Demos*, 3 August 2020. Retrieved from: <https://demos.co.uk/project/polis-and-the-political-process/>.
- Cartledge, Paul (2016). *Democracy: A Life*. Oxford: Oxford University Press.
- Dray, Stéphane and Josse, Julie (2015). Principal component analysis with missing values: a comparative survey of methods. *Plant Ecology*, 216(5), 657-667.
- Fisher, Ronald A. (1922). On the interpretation of χ^2 from contingency tables, and the calculation of P. *Journal of the Royal Statistical Society*, 85(1), 87-94.

- Hsiao, Yu-Tang, Lin, Shu-Yang, Tang, Audrey, Narayanan, Darshana and Sarahe, Claudina (2018). *vTaiwan: An Empirical Study of Open Consultation Process in Taiwan*. SocArXiv.
- Horton, Chris (2018). The simple but ingenious system Taiwan uses to crowdsource its laws. *MIT Technology Review*, 21 August 2018.
- Inglehart, Ronald (2005). *Modernization, Cultural Change, and Democracy: The Human Development Sequence*. Cambridge: Cambridge University Press.
- King, Martin (2018). Deliberation and Decision Making Online: Evaluating Platform Design. *WestminsterResearch*. Retrieved from: <https://westminsterresearch.westminster.ac.uk/download/b445a61f16b8dc5093d5b3f22272a73of17cob40of4bba67dcood6844546e466/4372517>.
- King, Martin (2019). Building Consensus and Compromise on Uber in Taiwan. *Public Impact Observatory* (Center for Public Impact), 18 September 2019. Retrieved from: <https://www.centreforpublicimpact.org/case-study/building-consensus-compromise-uber-taiwan/>.
- Lakatos, Imre (1978). *The methodology of scientific research programmes*. Cambridge: Cambridge University Press.
- Lee, Louisa, Zhang, Siyu and Yang, Vicky Chuqiao (2017). Do two parties represent the US? Clustering analysis of US public ideology survey. *SIAM Undergraduate Research Online*, 12.
- Lloyd, Stuart P. (1982). Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2), 129-137.
- MacQueen, J. B. (1967). Some Methods for classification and Analysis of Multivariate Observations. *Proceedings of 5th Berkeley Symposium on Mathematical Statistics and Probability* (University of California Press), 1, 281-297.
- Madrigal, Alexis C. (2014). How Netflix Reverse-Engineered Hollywood. *The Atlantic*, 2 January 2014. Retrieved from: <https://www.theatlantic.com/technology/archive/2014/01/how-netflix-reverse-engineered-hollywood/282679/>.

- Miller, Carl (2019a). *The Death of the Gods: The New Global Power Grab*. London: Windmill, Penguin Books.
- Miller, Carl (2019b). Can Taiwan Reboot Democracy? *BBC*, 29 October 2019. Retrieved from: <https://www.bbc.com/news/technology-50127713>.
- Miller, Carl (2019c). Taiwan is making democracy work again. It's time we paid attention. *Wired*, 26 November 2019. Retrieved from: <https://www.wired.co.uk/article/taiwan-democracy-social-media>.
- Miller, Carl (2020). How Taiwan's 'civic hackers' helped find a new way to run the country. *The Guardian*, 27 September 2020. Retrieved from: <https://www.theguardian.com/world/2020/sep/27/taiwan-civic-hackers-polis-consensus-social-media-platform>.
- Mulgan, Geoff (2017). *Big Mind: How Collective Intelligence Can Change Our World*. Princeton: Princeton University Press.
- OECD (2020). *Innovative Citizen Participation and New Democratic Institutions: Catching the Deliberative Wave*. Paris: OECD. Retrieved from: <https://doi.org/10.1787/339306da-en>.
- OECD-OPSI (2018). *Pol.is, Official Languages and a Shift Towards 'People-First' Policy Development*. *Observatory of Public Sector Innovation*. Retrieved from: <https://oecd-opsi.org/innovations/pol-is-official-languages-and-a-shift-towards-people-first-policy-development/>.
- Pearson, Karl (1901). On lines and planes of closest fit to systems of points in space. *Philosophical Magazine*, 2(11), 559-572.
- Poole, Keith T. and Rosenthal, Howard (1985). A Spatial Model for Legislative Roll Call Analysis. *American Journal of Political Science*, 29(2), 357-384.
- Rosenberg, Marshall B. (2003). *Nonviolent Communication: A Language of Life*. Encinitas: PuddleDancer Press.
- Rousseeuw, Peter J. (1987). Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20, 53-65.

- Roweis, Sam (1998). EM Algorithms for PCA and SPCA. *Advances in Neural Information Processing Systems*, 626-632.
- Ruiz, Rebecca (2018). Nasty women (and men) to crowdsource 'marching orders' ahead of midterm elections. *Mashable*, 19 January 2018. Retrieved from: <https://mashable.com/2018/01/19/womens-march-march-on-crowdsource-poll-activism/>.
- Salganik, Matthew and Levy, Karen E. (2015). Wiki Surveys: Open and Quantifiable Social Data Collection. *PLOS One*, 10(5).
- Segaran, Toby (2007). *Programming Collective Intelligence*. Newton: O'Reilly.
- Tang, Audrey (2016). Uber responds to vTaiwan's coherent blended volition. *Medium*, May 2016. Retrieved from: <https://blog.pol.is/uber-responds-to-vtaiwans-coherent-blended-volition-3e9b75102b9b>.
- Tang, Audrey (2019). A Strong Democracy Is a Digital Democracy. *The New York Times*, 15 October 2019.
- Turoff, Murray and Linstone, Harold A. (1975). *The Delphi Method: Techniques and Applications*. Boston: Addison-Wesley Educational Publishers Inc.
- Weber, Max (1969). "Objectivity" in social science and social policy. *The Methodology of the Social Sciences*. New York: The Free Press, 49-112.
- Yang, Vicky Chuqiao, Abrams, Daniel M., Kernell, Georgia and Motter, Adilson E. (2020). Why Are U.S. Parties So Polarized? A "Satisficing" Dynamical Model. *SIAM Review*, 62(3), 646-657.