Department: People in Practice
Editors: Melanie Tory, mtory@tableau.com
Daniel F. Keefe, dfk@umn.edu

Designing Technology for Sociotechnical Problems: Challenges and Considerations

Narges Mahyar, Mahmood Jasim, and Ali Sarvghad
University of Massachusetts Amherst

***Abstract*—Designing technology for sociotechnical problems is challenging due to the heterogeneity of stakeholders' needs, the diversity among their values and perspectives, and the disparity in their technical skills. Careful considerations are needed to ensure that data collection is inclusive and representative of the target populations. It is also important to employ data analysis methods that are compatible with users' technical skills and are capable of drawing a representative picture of people's values, priorities, and needs. However, current technical solutions often fail to meet these critical requirements. In this article, we present a set of empirically-driven design considerations for building technological interventions to address sociotechnical issues. We then discuss open challenges and tradeoffs around privacy, ethics, bias, uncertainty, and trust. We conclude with a call to action for researchers to advance the domain knowledge and improve our technological arsenal for addressing sociotechnical problems.**

Digital Object Identifier 10.1109/MCG.2020.3017405
Date of current version 23 October 2020.

Civic decision-making is a complex sociotechnical process that involves engaging and gathering input from various stakeholders, analyzing their input, and making policy decisions that reflect people's perspectives. For example,

Published by the IEEE Computer Society

government officials might require public input to make decisions around public health risks or planning future neighborhoods. Gathering inclusive and representative input from the public and effectively analyzing and reflecting on these inputs is crucial for decision-makers to take actions based on people's needs, issues, and priorities.

The advances of ubiquitous technology and increased access to mobile devices and online platforms have enabled government and city officials to collect public input more efficiently using online engagement platforms (e.g., DemocracyOS, PlaceSpeak, NextDoor), kiosks, and interactive public displays. While these technologies proved successful for broadening data collection, they still fall short of incorporating views of technologically, socially, and economically marginalized populations. Furthermore, as the scale of the collected public data increases, decision-makers struggle to analyze and make sense of this abundance of often qualitative data using traditional methods such as manually examining the data or using generic qualitative data analysis tools.

In the past six years, we have been collaborating closely with both the public and decision-makers to design tools and visualization techniques that facilitate more inclusive data gathering and more effective data analysis. Throughout this process, we learned that decision-makers grapple with issues around collecting meaningful, inclusive, and authentic data from the public, applying appropriate analysis approaches, such as computational algorithms and visualization techniques to analyze public input, evaluating the effectiveness of their approaches in practice, and communicating the decisions back to the public in an understandable and accessible manner. Meanwhile, the public struggles to find opportunities to get involved and contribute to the decision-making process by sharing their input. As a result, they find it difficult to understand the process and trust the outcomes.

In this article, we reflect on our experiences in designing technology to support civic engagement and decision-making. We identify design challenges, provide empirical insights on design considerations based on our prior work, and discuss future directions. Although designing technology for complex sociotechnical* problems is a difficult endeavor with an ever-changing landscape of ethics, values, and ideas in the civic domain, we hope that our empirical insights can better inform researchers and help them navigate uncharted waters in designing impactful tools and visualization techniques for complex real-world problems.

CHALLENGES AND DESIGN CONSIDERATIONS

We describe the challenges involved in the civic decision-making process and design considerations that we explored to address these issues. More specifically, we examine two key areas—collecting public input† and analyzing qualitative public data.

Collecting Public Input

Public input is at the heart of civic decision-making. Collecting meaningful and authentic data from a wide range of stakeholders is a precursor to successful decision-making. From the decision-makers' perspectives, it is crucial to collect meaningful data that represent diverse public viewpoints.[1] To gather such data, they need to engage various stakeholders and provide opportunities for the public to contribute inclusive and authentic input. However, traditional data collection methods often do not allow them to cast a wide net to reach out to a broad population and achieve the desired depth and richness in the collected data.[1] On the flip side, the public faces a different set of challenges. Lack of awareness of workshops, surveys, or meetings prohibits people from being involved in the civic decision-making process. They may also miss the chance to participate due to lack of time or inability to physically attend the venues.[2] Furthermore, people are often coerced to agree or disagree with proposed ideas put forward by decision-makers, rather than openly express their opinions.

In our CommunityCrit[3] project (see Figure 1), we partnered with a local planning group working

*Unlike traditional computer-based systems, *sociotechnical* systems include in their architecture and operation organizational and human actors along with technical systems, and they are affected by broader societal, political, economic, cultural, and organizational context.[22]

†*Public input* is used to refer to comments, new ideas, thoughts, or opinions shared by the public in qualitative textual format.

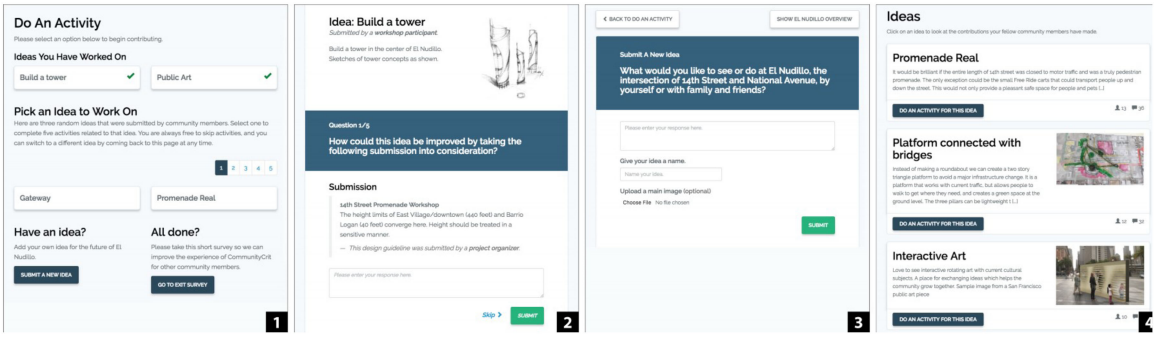

Figure 1. Microactivity-based workflow of CommunityCrit: 1) participants select an idea to share thought and comments on it; 2) a sample microactivity where the participants answer a question asked by another participant (they can do or skip an activity); 3) participants can submit new ideas; and 4) they can view an aggregation of ideas and contributions made by all participants. These steps are nonlinear: people can exit the workflow, pick another idea, submit a new idea, view all/my contributors at any time. After moderators approve new ideas, they will be online and visible to the community.

on redesigning a major intersection in San Diego, CA, USA. The group sought ideas, thoughts, and opinions from the community to help them make informed plans. The planning group was dedicated to engaging the public by organizing public tours, workshops, and surveys, but they could only collect a handful of public input. The amount of gathered data was insufficient, inconclusive, and unrepresentative of the domain problem. We observed their approaches by attending their face-to-face workshops and surveying both the community members and the planning group to understand the potential of online technology for gathering broader input from the affected community. Based on these close observations, we built and deployed a low-barrier online "microactivity" based platform called CommunityCrit.[3] This online platform (see Figure 1), proved to be successful in empowering the public to provide meaningful feedback, thoughts, and new ideas in a short amount of time. The microactivities allowed the community members to provide their opinions and reasoning as free-form text at their own volition without being forced to agree or disagree with preselected ideas and from anywhere without fixed time commitments. This tool also empowered decision-makers to overcome some of their challenges in involving and inviting the public to provide ideas and feedback. CommunityCrit made the data collection process more streamlined compared to the cost and manpower needed for running workshops, inviting speakers, providing refreshments, recording, and transferring ideas into reports. It enabled them to collect around 400 free-form text-based contributions from 63 unique participants over four weeks that provided much richer insights into community members' perspectives. We also interviewed the public and gained qualitative feedback that demonstrated how CommunityCrit's microactivity workflow helped to engage members of the public at their convenience. Participants appreciated the ability to contribute, especially given the difficulty of attending workshops, and expressed interest in using the system again. Our focus group session with the local planning group revealed that the tool helped raise awareness about the planning process and efforts, and gathered useful insights on public opinions (for quotes and detailed feedback, see the article by Mahyar *et al.*[3]).

Another issue is around biases as they can occur even in the most participatory processes during the collection and analysis of public input. Visualizing the gathered data without explanations of caveats and communication of uncertainty might imply that this particular portion of the data is representative of everyone's perspectives, where in reality, it is from a self-selected group and not necessarily representative of the affected populations. The data incompleteness and the negligence towards the incompleteness of gathered data can potentially inject biases in civic decisions.[4–6] One way to alleviate the biases associated with incomplete data is to ensure inclusivity in the data collection process. It is imperative to collect input from all stakeholders within the

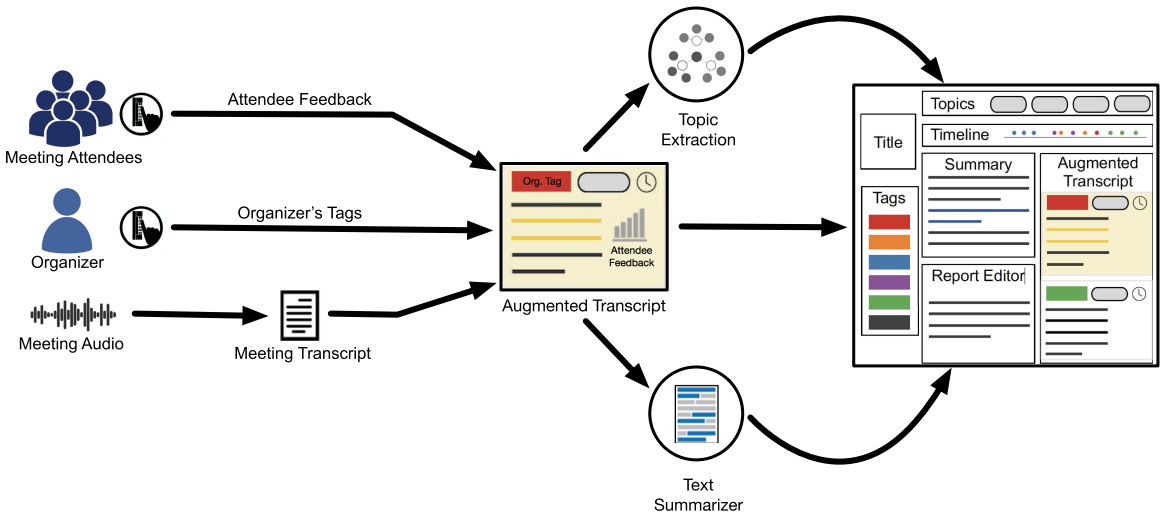

Figure 2. Snapshot of CommunityClick's workflow. During the meeting, attendees and organizers can use iClickers to share feedback and tag the meeting, respectively. The meeting is also audio-recorded for transcription. The audio recordings are transcribed automatically and then augmented with the attendees' feedback. Furthermore, we generated the feedback-weighted discussion summary and extracted the most relevant topics. The interactive interface enables the exploration and utilization of augmented meeting discussions, which is available online for organizers to examine and author meeting reports.

problem domain. Although achieving inclusivity at a holistic level might not be possible, it is worthwhile to strive for it within a smaller scope as a start. For example, in public town hall meetings, people who are privileged with training in analytical and rhetorical reasoning often find themselves at an advantage. Hence, broadening inclusively in these meetings can be a viable starting point. To take the first step towards this goal, we attended several town halls in the town of Amherst and surveyed the public attendees as well the decision-makers who organized these town halls. We found that 20% of attendees (out of 88) felt that they could not share their opinions during town halls despite their desire to get involved and being physically present. Education and financial backgrounds and social dynamics such as shyness and tendency to avoid confrontation with dominant personalities were contributing factors for these people not to speak up. As a result, their opinions remained invisible and left out of the data.

In our CommunityClick[‡7] project, we partnered with the local government officials to design tools that can help them to collect more inclusive and representative public data from town halls. To capture the opinions of silent participants, we used iClickers, an audience response technology popular in the education domain. We modified iClickers to allow reticent participants to provide real-time feedback on ongoing discussions silently by clicking one of the five options on the iClickers. In keeping with the idea of going beyond binary agreement or disagreement, we made the options fully customizable so that appropriate and relevant response options based on the topics discussed in a meeting. Figure 2 shows the workflow of our CommunityClick and how it enabled the silent attendees to provide their real-time feedback on any discussion points using iClickers and how these feedback were integrated with an automated transcripts for more inclusive analysis process.

In-the-wild deployment of CommunityClick in town hall meetings demonstrated the tool's efficacy in providing a more inclusive platform and especially empowered attendees who were not comfortable speaking up to share their thoughts silently during the meetings. We also found that the organizers were more confident in making policy decisions equipped with input from both vocal and silent public attendees. Based on prior work and our experiences, we argue that we should pay careful attention to ensure that the collected public data is as inclusive and representative as possible. We emphasize the importance of going beyond

[‡]https://cc.cs.umass.edu

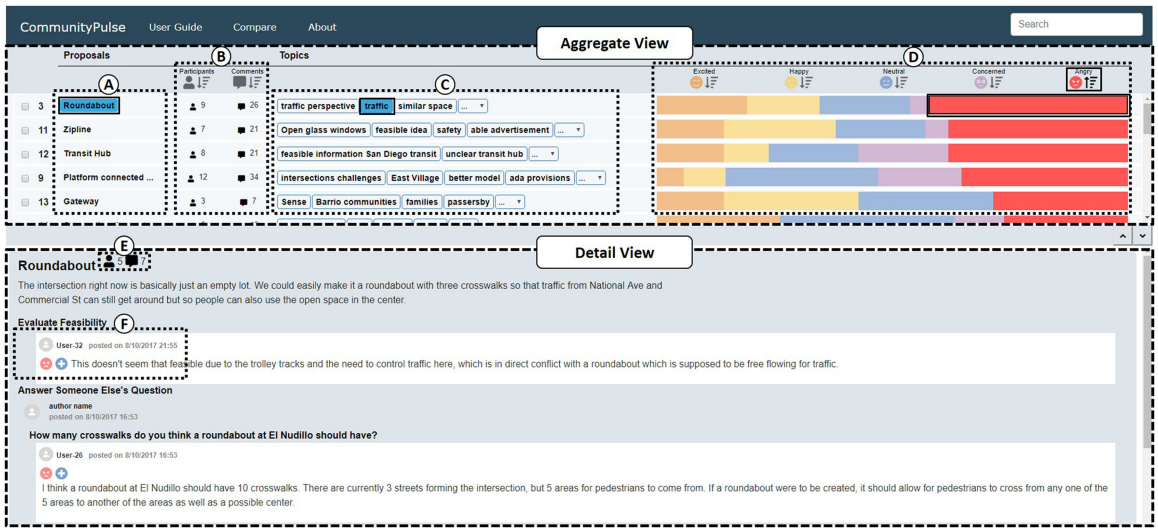

Figure 3. Snapshot of CommunityPulse. There are two views. The aggregate view shows: A) a list of proposals; B) the number of people and comments for each proposal; C) a list of topics for each proposal; and D) emoticons to sort the proposals based on emotions and stacked bar charts to present people's emotion distribution and drill down to actual comments. The detail view is updated based on the filters used in the aggregate view. This example shows E) meta-information based on the user-selected angry comments; and F) user information for each comment, with icons to represent associated emotion and option to save the comment as a note.

traditional data collection methods to reach a broader audience and include otherwise neglected populations to accumulate more comprehensive data that better represent the problem domain.

Analyzing Large-Scale Public Input

As the scale of collected public data increases, it becomes difficult for decision-makers to analyze and understand the data manually. Furthermore, public input often includes short free-form, unstructured text with ambiguous boundaries between agreement/disagreement or positive/negative. The data is also full of redundancies due to semantically similar input shared by multiple people. These issues make effective exploration and analysis of unstructured textual public data an open challenge that requires attention.[8] To this end, we conducted interview studies with 21 civic decision-makers across a broad range of organizations and responsibilities to gain deeper insights into their data analysis practices and challenges. We learned that they were constrained by time, manpower, lack of training and expertise to work with complex off-the-shelf data analysis tools, and scarcity of specialized data analysis tools to analyze large-scale public input.[1]

Computational methods are used to address some data analysis challenges faced by decision-makers. For instance, NLP techniques can be used to reduce redundancy in textual data and extract the most relevant information. However, it is equally important to consider what type of information needs to be extracted using computational approaches for informed decision-making. Our interviews with decision-makers revealed that they did not prefer popular text analysis methods such as sentiment analysis that aggregate all opinions into positive or negative clusters. In contrast, they wanted to know the nuances present in public input. For instance, they believed emotions played an important role in shaping public opinion, and they wanted to learn how/why people felt about civic issues. They also wanted to understand the main themes of public discussions.

We engaged in a long-term collaboration with these decision-makers and took a human-centered approach to accumulate their domain knowledge into designing and developing an interactive visual analytics tool called CommunityPulse.[§] This tool combines text analysis and

[§]https://communitypulse.cs.umass.edu

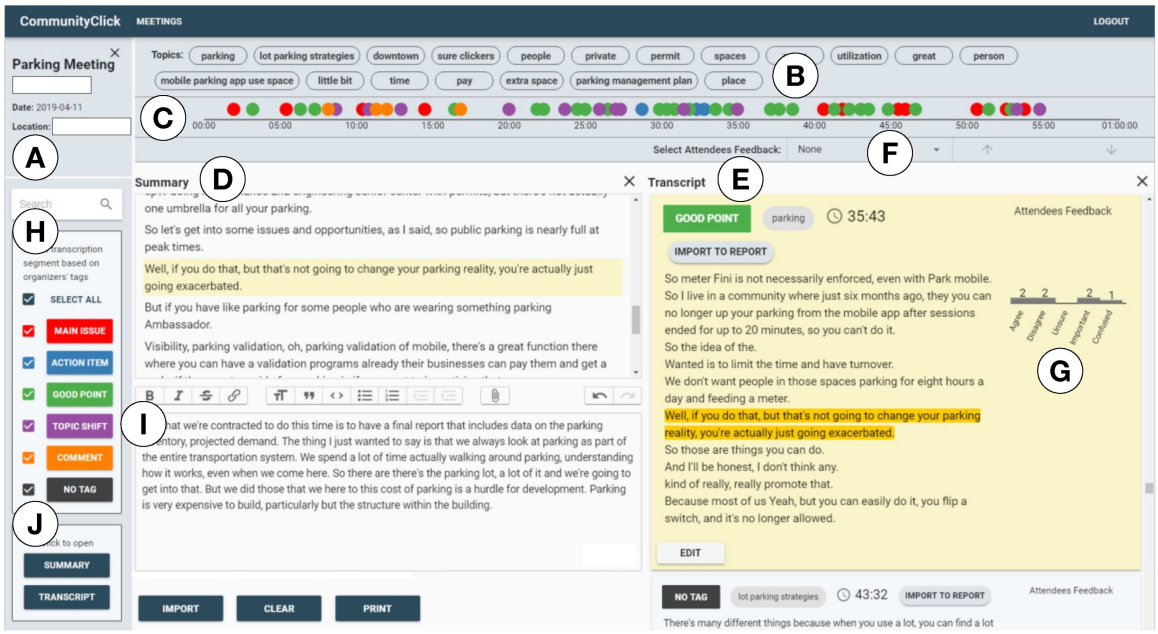

Figure 4. Snapshot of CommunityClick's interface. A) Title provides metadata about town hall meetings such as date and location. B) Main topics extracted from the meeting transcript. C) Timeline visualizes organizers' tags in chronological order. Each circle represents a tag. Clicking on a circle brings organizers to the corresponding transcript segment. D) Interactive feedback-weighted summary. E) Transcript view displays the transcript text alongside organizers' assigned tag, the main topic, and aggregated attendees' feedback in that time interval for each segment. F) Filters for attendee's feedback (based on what was provided on iClicker). G) Bar chart displays attendees' feedback. H) Filters for organizer's tags. I) Rich text editor for organizers to author the report. J) Options to view or collapse the summary or the transcript view.

visualization techniques to scaffold different facets of public input and enables multilevel exploration of public data. In particular, we applied affect analysis and topic modeling to go beyond sentiment analysis and extract main discussion topics (see Figure 3). We also structured public input by using interactive visualization as a scaffolding to organize unstructured text data around it. Our evaluation with another 12 experts demonstrated that CommunityPulse enabled them to understand public reflections on critical civic issues and identify hidden insights buried in the data while minimizing the required time, effort, and expertise.

Although computational approaches can help automate parts of the analysis process, one size does not fit all. It is essential to carefully consider the appropriate computational approaches that best match users' requirements and data characteristics. For instance, during the early days of our CommunityPulse project, we realized that off-the-shelf text analysis methods did not closely meet our collaborators' requirements and the characteristics of public input. Accuracy of supervised methods are dependent on trained domain data and do not generalize well in other domains. In contrast, the unsupervised methods are designed for general problems, but they ignore unique characteristics within the data. In our work, we iteratively refined the text analysis process by experimenting with several approaches and identified the most suitable one based on fine-tuning bidirectional encoder representations from transformers (BERT) to better identify contextual information for our affect analysis task. We also used a method for topic representation designed for online conversations instead of common standard libraries.

It is important to note that true inclusivity is only achieved when attempts to collect inclusive data is matched by attempts to achieve inclusive analysis. In CommunityClick (see Figure 4), we collected inclusive attendees' feedback during several town hall meetings by modifying iClickers as a real-time feedback system so that attendees could

provide feedback during the ongoing discussion silently using one of the five iClicker options (e.g., agree, disagree, important, unsure, and confused). CommunityClick's workflow presented in Figure 2 shows how the data gathered from in-person meetings is transformed by applying inclusive computational approaches to help decision-makers to better reflect on the data. The attendees' vocal and silent feedback gathered from town halls are fed into an automatically-generated meeting transcript that is augmented with the attendees' feedback. In order to accelerate the analysis process, we summarized the augmented transcript and incorporated it into an interactive authoring tool for organizers to write reports based on their analysis. We experimented with different text summarization methods but found that popular text summarization methods usually consider text at a lower level as a combination of words and sentences that only focus on the syntactic and semantic features. Hence, it was impossible to incorporate silent attendees' feedback into the aggregated text summary using these methods. To circumvent this issue, we proposed a novel text summarization method. This method took the meeting discussion transcript and attendees' feedback as the input and generated a *feedback-weighted summary* that prioritized the meeting discussion segments that the attendees engaged with the most. This way, the summarization method incorporated both verbal and silent input into the summary. We found from our evaluation of the tool by decision-makers that this augmentation allowed them to identify the attendee's interests on key discussion points.

Exploration and analysis of large-scale public input are challenging due to their unstructured nature, inherent ambiguity, and redundancy. Applying computational methods and visualization techniques to enable surfacing the public's reflections, priorities, and hidden insights from an inclusive population can address this gap. However, we echo previous work that demonstrates the absence of a universally generalized model for text analysis and how off-the-shelf methods lead to ineffective systems.[9]

Design Considerations in Civics

The rise of the internet and ubiquitous access to mobile devices and online platforms has enabled decision-makers to collect information about people's needs and issues more efficiently than ever. Public opinion polls, surveys, forums, and mobile apps are examples of such data collection tools. Generally, these tools attempt to lower access and reporting barriers for the public to help them report issues and share ideas while enabling city officials to hear from a broader, more representative range of community members. These tools have increased access for a broader spectrum of the public to share their ideas, issues, and comments on civic issues. However, technology for gathering data from the public must carefully consider engagement, provide adequate information and resources, and create opportunities for flexible involvement based on an individual's interest, available time, and expertise. To provide manageable context and upfront resources, extracting key excerpts from lengthy and complex civic design guidelines can enable the public to understand the context and get started on tasks and activities in a shorter amount of time. Breaking down complex tasks to sets of microactivities proved to be successful in our work.[3] Finally, providing social interactions in online spaces allows people to explore the context, design alternatives, and input provided by others, which ultimately helps the public to explore and understand the context and provide more meaningful feedback.

When visualizing extracted information from large-scale public input, we found two design considerations to be critical—*simplicity* and *familiarity*. In our interviews with decision-makers, we found that their expertise in analyzing public input does not necessarily translate to the expertise of working with visual analytics tools. To account for decision-makers' limited expertise in data analysis tools, simple visualization encoding such as bar charts can reduce the learning curve and enable decision-makers with various skill levels to utilize the tool without extensive training. In our CommunityPulse project, we observed how the decision-makers with no significant prior skills with data visualization tools quickly acclimatized to the simple visual encodings.

It is essential to maintain familiarity when introducing new technology into the existing data analysis pipeline.[10] It is especially critical in

the civic domain due to the wide variety in civic decision-makers' expertise levels when it comes to working with data analysis tools. In civic decision-making, the traditional method of analyzing public data is examining text data manually or using generic qualitative data analysis tools. Hence, visualization tools and techniques for decision-makers should incorporate more textual elements from the data that resemble and align with their current data analysis practices, rather than providing aggregated statistics. Furthermore, discussion data can be represented in a forum style, invoking a sense of familiarity when examining public input.

DISCUSSION

We discuss the major barriers to wide-scale design and integration of tools and techniques for decision-making on complex sociotechnical issues. We also highlight future opportunities and approaches to overcome these barriers.

Broader Societal and Technical Issues

Designing tools and techniques to tackle complex sociotechnical problems is a challenging endeavor. This is due to the intricate interplay among socioeconomic, cultural, and computational components, and their long-term impact on society. Particularly, understanding a wide range of potential users, including the public, analysts, and decision-makers' perspectives, and adapting the design based on their requirements and technical abilities is a nontrivial task. Design requirements can vary based on different issues and selected populations. Therefore, close collaboration and partnership with affected populations are necessary to understand their needs, challenges, skill sets, and accessibility to technology, among other concerns.

In our projects, we took a human-centered design approach where we worked with domain users very early in the project and involved them in the iterative design process by frequently receiving feedback, discussions on possible designs, and persistent communication. Despite the rapport we created with our partner groups during our long-term collaboration, their eagerness, and the apparent need to adopt new tools, we experienced many challenges regarding adoption. In general, adopting new technologies in civics is a great challenge due to the sensitivity of the community input (personal information embedded into the comments, such as personal, experience, livelihood, neighborhood, etc.), and bureaucratic issues of convincing the upper echelon (e.g., government, management), to provide infrastructural support for adoption.

Privacy concerns are another critical issue related to the collection and analysis of public data. There is a tradeoff between gathering rich insights from the public and preserving privacy. Depending on the sociopolitical system's turbulence, the public often prefers to maintain privacy when sharing their input and do not feel comfortable providing demographic information. However, the decision-makers consider demographic information crucial for understanding the reasoning behind public input and whether the information is credible enough to act upon. While obfuscating personalized information is often desirable,[4] it also limits the richness of the data critical to making informed decisions. Although there are visualization design approaches around differential privacy and data encryption to preserve user privacy,[11] these methods aggregate the data to obscure specific information. As a result, they could discard the richness of qualitative data around the reasoning and disposition behind people's input. There is a lack of definitive guidelines for this paradoxical issue of designing systems that enable exploration of personalized information such as demographics while preserving user privacy. Future researchers might explore this open problem to find pragmatic solutions suitable for the civic domain.

Ethical Considerations and Biases in Visualization

Historically visualizations have been considered to be ethically neutral. However, previous researchers questioned this neutrality by emphasizing the subjective nature of design decisions made throughout the data visualization pipeline by the designers.[4,12] The presence of choice in the method of collecting data by observing real-world events, data interpretations and quantification, use of computational approaches to analyze data, and selection of outliers (both data and users) for omission from the visualization demands ethical considerations. This is especially true for sociotechnical problems. Gitelman

argued that data should not be considered as neutral and as objective facts that fully represent the real-world.[13] Furthermore, as quantity does not ensure quality, collecting more data does not improve the neutrality of data. As Boyd and Crawford suggested, "bigger data are not always better data."[5] Instead, we should put more focus on collecting inclusive data in ways that do not neglect silent and unpopular individuals from a marginalized group. We should also take great care while omitting data points as outliers as such omissions might suppress marginalized opinions. In recent years, there have been much debate and studies on biases in machine learning models, natural language models, text analysis methods, and similar statistical approaches that try to generate models and make inferences on the real-world data.[14] Although the debate circulates around how the bias gets injected into the models, evidence shows that it can happen at any stage during the model generation process. Various visual analytics tools heavily use such models, especially when the underlying data is too large to analyze manually. It is important to consider and account for the injection of biases that can translate from these computational methods to the visualization process.

Furthermore, there is an emerging interdisciplinary field of studying cognitive biases in visualization that mostly focuses on how the decision-makers who use visual analytics systems to make decisions, often make inherently biased decisions as opposed to rational ones.[15] However, little attention has yet been paid to the process of how designers' biases in creating visualization systems propagate into the visual design. In this regard, Cairo suggested that visualization researchers should feel obligated to build visualizations that enable comprehension, but such designs should be based on a thorough analysis of information to bring attention to what matters the most.[6] Visualizations designed with such considerations have the potential not only to increase awareness but also to enhance understanding and knowledge. One way to incorporate such considerations into the visualization design process is to look at the process from a critical perspective at different stages of design.[16] Such perspectives can help visualization designers to consider the people involved in the process, emphasize their values, promote equality and emancipation throughout the design, and address possible consequences and impact of the design on civic decision-making and knowledge acquisition.[17]

Lack of Trust in Visualization in Civic Domain

Coupled with logistics and privacy concerns, there is often apprehension towards replacing existing human-driven and often manual public input analysis pipeline with semiautomatic visualization tools. One of the key reasons behind this concern is the questionable accuracy of visual analytics tools. The root of such inaccuracies can be traced back to a lack of domain-specific public data to train language models.[9,18] Major public engagement platforms such as PlaceSpeak, DemocracyOS, or NextDoor do not allow the use of public opinion data for research purposes due to privacy concerns. Furthermore, local authorities often collect public opinion data using physical pen and paper methods that are often not digitized or archived before or after use. As such, training models to predict often ambiguous public opinions becomes challenging and results in inaccurate systems. Such inaccuracies result in a lack of trust in automated systems.

Lack of trust in visualization also stems from the absence of techniques to communicate data quality issues such as incompleteness.[1] Designers and researchers need to consider how to visualize such uncertainties in the visual analytics systems, not as an afterthought or additional component, but as a core design goal.[19] Recently, visualization researchers have focused on visualizing uncertainties.[20] However, the dynamics and relationships between visualizing uncertainty and its impact on the trust in visualization remain debatable. While visualizing uncertainties can increase a system's credibility, it can also lead to lack of trust if the users consider the system to be incapable of functioning in the presence of incomplete data.[21] Without further empirical studies, it is difficult to identify the interplay between uncertainty and the trustworthiness of visual analytic systems in the civic domain. Currently, visual analytics systems are rarely institutionalized and made a part of the public opinion analysis pipeline. As a result, the efficacy and

impact of such systems in addressing complex sociotechnical problems in the wild remains uncertain and needs further assessment.

Opportunities and Possible Future Directions

Despite the challenges involved, from our experience working with the stakeholders in the civic domain, we observed interest and enthusiasm from both the public and decision-makers on the potential of using online tools. They highlighted the value of such tools for gathering and analyzing public input and establishing a tighter communication channel between the public and decision-makers. This abundance of interest shows promise and potential for designing visualization tools to address complex real-world problems in the civic domain.

To circumvent the existing challenges of designing tools for sociotechnical problems, as designers of analytics systems, we can start by rethinking the design and development approach of such tools. First, we can acknowledge, consider, and embrace the presence of uncertainty in the design process. Second, we can make fewer assumptions about people's needs when analyzing data and pay careful attention to surface what matters for the intended audience. Finally, we can welcome the dynamic characteristic of a community-oriented design process where the target users are incorporated into the design process to better adjust to the ever-changing social and political landscape of the civic domain.

Evaluating data visualization systems have been a challenging issue for decades. It is difficult to objectively assess the performance of such systems because their impact and value are intertwined with the reception by their target audience. With the aforementioned design considerations in place, we can strive to gain users' trust in the visualization systems. We can start by gradually integrating the new tools and techniques with the existing ecosystem, not as a replacement of the current norm, but as a complement to support the analysis process.[22] We posit that demonstrating and establishing the efficacy and efficiency of such systems gradually over a long period will build users' trust by achieving tangible outcomes. This might also smoothen the integration process of such systems into the civic domain to address sociotechnical issues.

Finally, we call for more attention towards designing visualization for social good. While designing such tools, we can carefully consider presenting the data to diverse audiences with variable visual literacy and analysis skills. We can incorporate uncertainties inherent with the data and analysis process and examine the ethical impacts and biases associated with the complete design and development process, including the data collection, analysis, design, and implementation. To achieve this, we call for tighter interdisciplinary collaboration between HCI and visualization researchers, designers, practitioners, government officials, and data analysts.

CONCLUSION

In this article, we emphasize the major barriers to designing visualization tools and techniques to address complex sociotechnical issues. We highlight how uncertainty, noninclusive data collection, and biased computational methods can result in ineffective visualization tools. We draw from our experience to suggest considerations around ethics, bias, uncertainty, and trust during the visualization design pipeline, including data collection, analysis, design, development, and evaluation. We conclude with a call to action for researchers to make appropriate design considerations when approaching complex sociotechnical problems.

ACKNOWLEDGMENTS

The authors would like to thank S. P. Dow, E. Hoque, A. X. Zhang, P. Khaloo, S. Wadhwa, M. R. James, M. M. Ng, and R. A. Wu for their valuable contributions to these projects. They would also like to thank all interviewees and participants who closely worked with them during the last six years to inform the design, development, and evaluation of these systems.

REFERENCES

1. N. Mahyar, D. V. Nguyen, M. Chan, J. Zheng, and S. P. Dow, "The civic data deluge: Understanding the challenges of analyzing large-scale community input," in *Proc. Des. Interactive Syst. Conf. ACM*, 2019, pp. 1171–1181.

2. J. Mansbridge, J. Hartz-Karp, M. Amengual, and J. Gastil, "Norms of deliberation: An inductive study," *J. Public Deliberation*, vol. 2, no. 1, 2006, Art. no. 7.
3. N. Mahyar, M. R. James, M. M. Ng, R. A. Wu, and S. P. Dow, "Communitycrit: Inviting the public to improve and evaluate urban design ideas through micro-activities," in *Proc. Conf. Human Factors Comput. Syst.*, 2018, Art. no. 195.
4. M. Correll, "Ethical dimensions of visualization research," in *Proc. CHI Conf. Human Factors Comput. Syst.*, 2019, pp. 1–13.
5. D. Boyd and K. Crawford, "Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon," *Inf., Commun. Soc.*, vol. 15, no. 5, pp. 662–679, 2012.
6. A. Cairo, *How Charts Lie: Getting Smarter About Visual Information*. New York, NY, USA: Norton, 2019.
7. M. Jasim, P. Khaloo, S. Wadhwa, A. X. Zhang, A. Sarvghad, and N. Mahyar, "CommunityClick: Capturing and reporting community feedback from town halls to improve inclusivity," in *Proc. 23 rd ACM Conf. Comput.-Supported Cooperative Work Social Comput.*, 2020, Art. no. 29.
8. A. Shrivastava and J. Heer, "ISEQL: Interactive sequence learning," in *Proc. 25th Int. Conf. Intell. User Interfaces*, 2020, pp. 43–54.
9. J. Chuang, D. Ramage, C. Manning, and J. Heer, "Interpretation and trust: Designing model-driven visualizations for text analysis," in *Proc. SIGCHI Conf. Human Factors Comput. Syst. ACM*, 2012, pp. 443–452.
10. S. Faridani, E. Bitton, K. Ryokai, and K. Goldberg, "Opinion space: A scalable tool for browsing online comments," in *Proc. SIGCHI Conf. Human Factors Comput. Syst.*, 2010, pp. 1175–1184.
11. R. Hongde, W. Shuo, and L. Hui, "Differential privacy data aggregation optimizing method and application to data visualization," in *Proc. IEEE Workshop Electron., Comput. Appl.*, 2014, pp. 54–58.
12. D. Sacha, H. Senaratne, B. C. Kwon, G. Ellis, and D. A. Keim, "The role of uncertainty, awareness, and trust in visual analytics," *IEEE Trans. Visualization Comput. Graph.*, vol. 22, no. 1, pp. 240–249, Jan. 2016.
13. L. Gitelman, *Raw Data is an Oxymoron*. Cambridge, MA, USA: MIT Press, 2013.
14. N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan, "A survey on bias and fairness in machine learning," 2019, *arXiv: 1908.09635*.
15. E. Wall, L. M. Blaha, L. Franklin, and A. Endert, "Warning, bias may occur: A proposed approach to detecting cognitive bias in interactive visual analytics," in *Proc. IEEE Conf. Vis. Analytics Sci. Technol.*, 2017, pp. 104–115.
16. M Dörk P. Feng, C. Collins, and S. Carpendale, "Critical INFOVIS: exploring the politics of visualization," in *Proc. CHI'13 Extended Abstracts Human Factors Comput. Syst.*, 2013, pp. 2189–2198.
17. B. C. Stahl and C. Brooke, "The contribution of critical is research," *Commun. ACM*, vol. 51, no. 3, pp. 51–55, 2008.
18. Y. Jin, "The challenges and countermeasures of realizing the modernization of government network public opinion management in the age of big data," in *Proc. Int. Conf. Econ., Bus., Manage. Corporate Social Responsibility*, 2018.
19. J. Hullman, "Why authors don't visualize uncertainty," *IEEE Trans. Visualization Comput. Graph.*, vol. 26, no. 1, pp. 130–139, Jan. 2020.
20. J. S. Mason, D. Retchless, and A. Klippel, "Domains of uncertainty visualization research: A visual summary approach," *Cartography Geographic Inf. Sci.*, vol. 44, no. 4, pp. 296–309, 2017.
21. E. Mayr, N. Hynek, S. Salisu, and F. Windhager, "Trust in information visualization," in *Proc. EuroVis Workshop Trustworthy Visualization*, 2019, Art. no. 20191187.
22. D. A. Norman and P. J. Stappers, "Designx: Complex sociotechnical systems," *J. Des., Econ. Innov.*, vol. 1, no. 2, pp. 83–106, 2016.

Narges Mahyar is currently an Assistant Professor with the College of Information and Computer Sciences, University of Massachusetts Amherst, Amherst, MA, USA. She co-directs the HCI-VIS lab, which focuses on developing novel social computing and visualization tools to facilitate data analysis, communication, and exploration of real-world sociotechnical problems in domains such as civics and healthcare. She received the B.S. degree in electrical engineering from Tehran Azad University, Tehran, Iran, an M.S. degree in information technology from the University of Malaya, Kuala Lumpur, Malaysia, and the Ph.D. degree in computer science from the University of Victoria, Victoria, BC, Canada. Her Ph.D. was funded by SAP Business Objects for four years, where she worked closely with SAP experts and research team to develop novel technologies for collaborative visual analytics tasks. Contact her at nmahyar@cs.umass.edu

Mahmood Jasim is currently working toward a Ph.D. degree with the College of Information and Computer Sciences, University of Massachusetts Amherst, Amherst, MA, USA. He is a member of the HCI-VIS lab. His research focuses on bridging the gap between civic leaders and the community members by building novel technologies that combine machine learning and data visualization for large-scale public-generated data analysis. He received the B.S. and M.S. degrees in computer science and engineering from the University of Dhaka, Dhaka, Bangladesh. Contact him at mjasim@cs.umass.edu

Ali Sarvghad is currently a Research Assistant Professor with the College of Information and Computer Sciences, University of Massachusetts Amherst, Amherst, MA, USA. His research focuses on advancing visual data analysis and data-driven decision making by designing and building novel visualization and interaction techniques. He is also co-director of HCI-VIS Lab. He received the Ph.D. degree in computer science from the University of Victoria, Victoria, BC, Canada, and the M.S. degree in software engineering from the University of Malaya, Kuala Lumpur, Malaysia. Contact him asarv@cs.umass.edu

Contact department editors Melanie Tory at mtory@tableau.com and Daniel F. Keefe at dfk@umn.edu.