CommunityPulse: Facilitating Community Input Analysis by Surfacing Hidden Insights, Reflections, and Priorities

Mahmood Jasim
mjasim@cs.umass.edu
University of Massachusetts Amherst
USA

Enamul Hoque
enamulh@yorku.ca
York University
Toronto, Ontario, Canada

Ali Sarvghad
asarv@cs.umass.edu
University of Massachusetts Amherst
Amherst, Massachusetts, USA

Narges Mahyar
nmahyar@cs.umass.edu
University of Massachusetts Amherst
Amherst, Massachusetts, USA

ABSTRACT

Increased access to online engagement platforms has created a shift in civic practice, enabling civic leaders to broaden their outreach to collect a larger number of community input, such as comments and ideas. However, sensemaking of such input remains a challenge due to the unstructured nature of text comments and ambiguity of human language. Hence, community input is often left unanalyzed and unutilized in policymaking. To address this problem, we interviewed 14 civic leaders to understand their practices and requirements. We identified challenges around organizing the unstructured community input and surfacing community's reflections beyond binary sentiments. Based on these insights, we built CommunityPulse, an interactive system that combines text analysis and visualization to scaffold different facets of community input. Our evaluation with another 15 experts suggests CommunityPulse's efficacy in surfacing multiple facets such as reflections, priorities, and hidden insights while reducing the required time, effort, and expertise for community input analysis.

CCS CONCEPTS

• **Human-Centered Computing → Human Computer Interaction (HCI)**.

KEYWORDS

Digital civics, visual analytics, text analysis, community input

ACM Reference Format:
Mahmood Jasim, Enamul Hoque, Ali Sarvghad, and Narges Mahyar. 2021. CommunityPulse: Facilitating Community Input Analysis by Surfacing Hidden Insights, Reflections, and Priorities. In *Designing Interactive Systems Conference 2021 (DIS '21), June 28-July 2, 2021, Virtual Event, USA.* ACM, New York, NY, USA, 18 pages. https://doi.org/10.1145/3461778.3462132

DIS '21, June 28-July 2, 2021, Virtual Event, USA

ACM ISBN 978-1-4503-8476-6/21/06…$15.00
https://doi.org/10.1145/3461778.3462132

1 INTRODUCTION

Community engagement is paramount in the practice of participatory democracy [59]. There has been a surge of interest within the researchers and practitioners of HCI and digital civics towards advancing technology to address issues of civic engagement and participatory democracy [17, 41, 57, 66, 73, 76, 80]. The advancements in digital civics have enabled civic leaders to broaden their outreach and collect thousands of comments, new ideas, thoughts, or opinions from the community, which we refer to as *community input*, through online engagement platforms (e.g., CommunityCrit [72], Decidim [7], DemocracyOS[5], and Pol.is [1]). Previous research found that civic leaders have a strong preference towards gathering rich qualitative input by moving past simple measures of public preference (e.g., surveys, voting, polling) to understand people's priorities, issues, and viewpoints [17, 40, 41, 73]. Civic leaders need to analyze, understand, and utilize these rich qualitative inputs to inform policy decisions around various civic proposals. Their analysis and decision-making can have critical impact on government policies [25], health [85], transit [26], urban planning [71], and environment [51] at a local or national scale.

However, as the scale of community engagement and gathered data increases to thousands of comments, civic leaders face obstacles in effectively and efficiently analyzing and making sense of textual community input for decision-making. Exploring, analyzing, and sensemaking[1] of online text data is challenging due to high dimensionality, lack of structure, and ambiguity inherent in natural human languages [14]. The challenges intensify in the civic domain due to informal language, redundancy, and the unclear boundary between positive and negative opinions shared as comments in community engagement platforms. Current practices around analyzing community input continue to rely heavily on manually coding or using qualitative data analysis tools (e.g., [3, 4, 11]) for thematic analysis. Civic leaders often struggle to make sense of community perspectives effectively using manual analysis, which is tedious, time-consuming, and labor-intensive [58, 73]. Existing civic technologies mostly enable opinion sharing (Opinion Space [50]), consensus building (ConsiderIt [67]), and providing high-level overviews of community-generated data (Pol.is [1]). However, these tools mainly provide aggregated high-level statistics and put less emphasis on

[1] Sensemaking is an information-processing task that serves as a critical prerequisite for decision-making [94]

the in-depth investigation of multiple facets of community input. Thus, existing tools fall short of enabling civic leaders to peruse and sublimate large-scale community input into concrete and actionable insights. Due to the lack of effective analysis methods, community input often remain unanalyzed and unutilized in civic decision-making, which is not only a missed opportunity in civic engagement but also a threat to participatory democracy [13, 73].

A first step towards addressing the aforementioned issue is learning about civic leaders' challenges in order to inform the design of technologies that can cater to their specific requirements and help them analyze nuanced, abstract, and often ambiguous textual community input [58, 73, 88]. To that end, we conducted in-depth interviews with 14 civic leaders with extensive experience in analyzing community input to gain a deeper understanding of their data analysis practices and specific technology needs. In these interviews, several critical requirements emerged including 1) organizing unstructured community input, 2) the need for meaningful and impactful methods to explore and make sense of unstructured community input at different levels of granularity, 3) going beyond binary sentiment analysis methods that group comments into positives and negatives to have a deeper understanding of community's emotions towards civic proposals, and 4) surfacing the main discussion points from community input.

Based on these findings, we distilled four design goals. These deign goals drove the development of CommunityPulse, an interactive visual analytics system that scaffolds different facets of community input such as their priorities, emotions, and reflections by combining text analysis methods such as emotion classification, topic modeling, and keyphrase extraction with data visualization. CommunityPulse supports multilevel exploration of community input starting from a high-level overview of the community's emotions along with the main discussion topics and allows interactions to drill down to individual text comments. We conducted an expert evaluation of CommunityPulse with 15 civic leaders (not involved in the initial interviews) who explored and examined the system for two weeks. Then, we conducted semi-structured interviews with these experts to understand their experiences of using and interacting with CommunityPulse's functionalities to analyze community input. The analysis of their qualitative feedback and suggestions during the interviews along with analysis of their usage patterns during the two-week exploration period revealed that they found that CommunityPulse enhanced their community input analysis practices by surfacing the community's priorities and identifying hidden insights that would be impossible to find using their current approaches.

The primary contributions of our work are as follows: 1) identification of key requirements for analyzing community input based on 14 in-depth interviews with civic leaders that highlight their challenges around meaningful exploration of unstructured community input, the need to go beyond binary sentiments for more fine-grained categorization of community's perspectives, and support for flexible role-based data exploration; 2) design and development of CommunityPulse by combining text analysis and interactive visualization while using visualization as a scaffolding to organize text analysis results for multi-level exploration and analysis of multifaceted community input; and 3) discussion of insights gained, and lessons learned from our study that highlight the utility and impact of our approach along with challenges faced and possible future directions to further explore designing interactive systems for facilitating sensemaking of unstructured community input. We also discuss how such systems could be expanded to undertake complex sociotechnical problems in other domains.

2 BACKGROUND

In this section, we discuss current practices and the challenges that civic leaders face around analyzing large-scale community input. We then review existing technologies for community input analysis and tools for online text analysis in various domains.

2.1 Community Input Analysis Practices and Challenges

Researchers in HCI and digital civics have been exploring various technological interventions, including online platforms and forums, to enable broader community engagement [40, 57, 73, 80, 92]. These online methods have enabled civic leaders to collect large-scale community input on various civic proposals, mostly as textual comments [1, 5, 7, 67, 72]. With the increase in community engagement, civic leaders, who collect, analyze, and make decisions based on community input, have to grapple with large-scale community-generated data. Previous research showed that civic leaders struggle to effectively analyze community input due to lack of time, analytical skills, and specially designed technology [58, 73]. To compensate for the lack of time to analyze community input, decision-makers often outsource the data analysis, which results in an iterative back and forth between decision-makers and analysts that further strain valuable time and resources. Furthermore, recent work [73] suggests that the community input analysis practices are limited to either manual data coding or utilizing an assortment of spreadsheet applications and generic data analysis tools, such as Atlas.Ti [3], Dedoose [4], and NVivo [11], to conduct thematic analysis. Although these tools offer some features to facilitate thematic coding of qualitative text data, they demand expertise from the users to be operationally effective. Even with such expertise, large-scale manual data requires copious amounts of time and labor for coding, resulting in an inefficient analysis process. Prior work has called for attention towards the lack of well-thought-out specially designed technologies for analyzing community input catered towards sociotechnical issues in digital civics [58, 73]. In this paper, we answer the call by making attempts to introduce visual analytics techniques that can help civic leaders to effectively analyze community input without sacrificing efficiency.

2.2 Current technologies for Community Input Analysis

Previous research in digital civics has resulted in tools that enabled opinion sharing, consensus building, and providing an overview of community input [1, 21, 50, 67, 91, 102]. For example, ConsiderIt builds a pro-con list to augment personal deliberation and help participants to identify common grounds from diverse opinions [67]. Procid enables consensus-building for distributed design decisions [102]. It allows users to post new ideas and criteria along with options to explore previous interactions. MetaViz visualizes

Table 1: This table shows needfinding interview participants, their roles, role descriptions, and their community input analysis approaches. Based on civic leaders' main responsibilities, we characterize them as Decision-Makers who are in charge of key decisions, Community Envoys who work to foster change on behalf of a community and Analysts who rigorously analyze and interpret community input. The majority of their analysis practices involve using document editor and spreadsheets while some civic leaders use qualitative data analysis tools or outsource the community input analysis.

ID	Role	Role Description	Data Analysis Approach(es)
N1	Decision-Maker	Facilitates engagement initiatives within the community	Spreadsheets
N2	Decision-Maker	Partners with groups within city for development projects	Google Docs and Spreadsheets
N3	Decision-Maker	Leads city wide initiatives in community outreach programs	Outsource
N4	Decision-Maker	Directs organization that promote community engagement	Spreadsheets
N5	Decision-Maker	Makes decisions in community engagement efforts	Pen and paper with no tools
N6	Community Envoy	Created online platform to collect public opinion	Spreadsheets
N7	Community Envoy	Directs a non-profit community group	Outsource
N8	Community Envoy	Researches the impact of technology on society	Spreadsheets
N9	Analyst	Coordinates city-wide community initiative	Google Docs and Spreadsheets
N10	Analyst	Advises government on environmental issues	Pen and paper with no tools
N11	Analyst	Coordinates distributive network for civic technologies	Atlas.Ti and Dedoose
N12	Analyst	Brings assessment methods to government	Atlas.Ti and Dedoose
N13	Analyst	Researches urban transformation in communities	Pen and paper with no tools
N14	Analyst	Creates strategic plans for varying scales	Spreadsheets

computationally identified metaphors from political blogs by focusing on the users' ability to perform critical thinking [21]. MyPosition allows users to vote on public policies and visualizes the results on public displays [91]. Opinion Space supports public opinion understanding by enabling users to browse online opinions [50]. Pol.is utilizes real-time visualization to provide a high-level overview of opinions [1]. While these tools provide an overview of community input and votes, they do not support multifaceted exploration and analysis of large-scale unstructured community input with systematic access to actual text comments in an organic way.

2.3 Technologies to Analyze Online Text Data

In the last decades, researchers proposed various tools and techniques [69] to analyze online text across several genres including customer feedback [33, 34, 54, 63, 96, 98], social media text [62, 74, 95], and online conversations[49, 61]. We reviewed the approaches used in these tools to better inform our design.

For visualizing user reviews for analysis, OpinionBlocks provides aspect-based summary of user reviews [63] using text snippets. Gregory et al. proposed ways to explore affective contents in product reviews [54] and Chen et al. analyzed conflicting opinions from customer review text using multiple connected views [34]. OpinionSeer uses various visualization techniques to present customer feedback on hotels. It applies feature-based opinion mining and an uncertainty model to extract customer opinions from text comments [96]. Review Spotlight provides a user interface for summarizing user-generated restaurant reviews using objective noun pairs and visualizing the summary result [98].

Recently, analysis of social media text has received significant attention [74, 95]. TwitInfo supports the exploration of large collections of tweets and visualizes positive and negative tweets [74] while OpinionFlow visualizes opinion diffusion in social media [95]. Opinion analysis is also examined in blogs, forums, and multi-party conversations [49, 61]. For example, ConVisIt [61] enables multi-faceted exploration of blog conversations based on topics, authors, and sentiments. ConToVi [49] visualizes speaker dynamics and behaviors such as sentiment, politeness, and eloquence, on different conversation topics using various animations.

Many of these tools were designed for experts who are comfortable working with complex visualization with multiple views. However, previous work highlights that community input analysts often lack extensive training in data and visual analytics [73]. The existing tools are not directly applicable in community input analysis due to incompatible and often complex design choices and the expertise required to utilize them effectively. Furthermore, often the text analysis techniques used in these methods are inaccurate [56, 63] beyond a reasonable margin to be utilized in the civic domain where critical policy decisions are made based on the analysis [13]. To address these shortcomings, a deeper understanding of civic leaders' requirements to design technologies that can cater to their needs in essential.

3 NEEDFINDING INTERVIEWS

In this section, we describe our needfinding interviews that revealed civic leaders' community input analysis practices and requirements, followed by four design goals that we derived from our findings from the interviews.

3.1 Participants

We conducted two rounds of interviews with 14 civic leaders over six months in 2019. We initiated contact with multiple civic leaders by connecting with local officials, community leaders, and individuals from large scale online community engagement initiatives across the united states. We recruited more interviewees willing

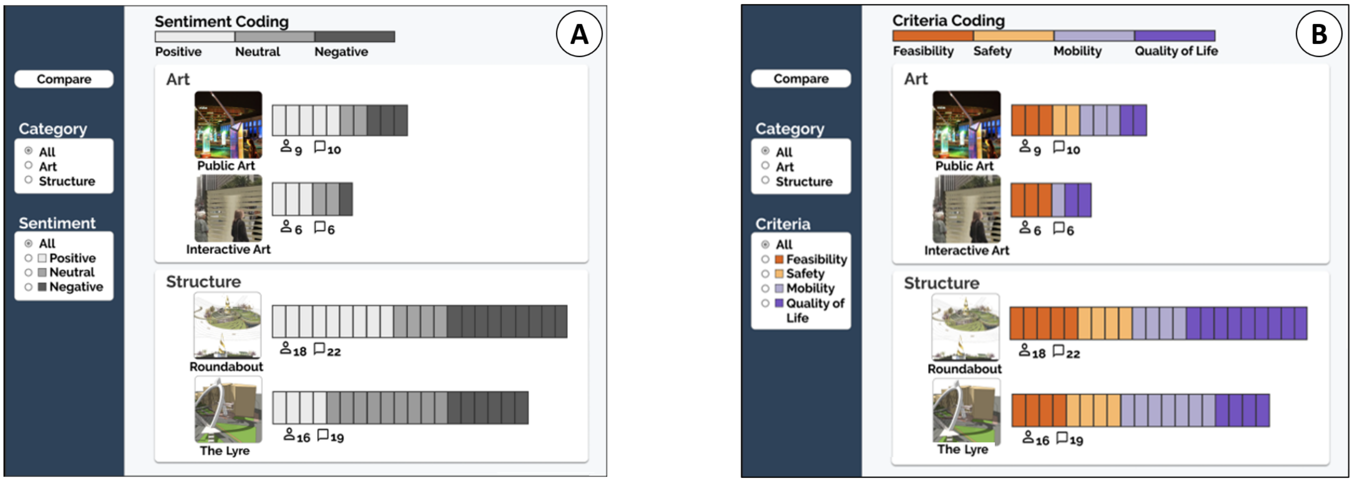

Figure 1: Examples of rapid prototypes that we presented to the interviewees to discuss about their technology requirements. Both of the prototypes were focused on providing high level overviews. (A) In this prototype, we focused on visualizing the community's sentiments towards proposals that are extracted using sentiment analysis from the community input. (B) In the second prototype, we focused on categorizing the comments into fixed criteria (e.g. safety, mobility) to highlight community's interests in those criteria.

to participate by emailing the civic leaders who were suggested by our initial contacts using the snowball method [23]. To capture diverse objectives and approaches towards community input analysis, we engaged with several local and national organizations who grapple with the challenges of analyzing and utilizing large-scale community-generated data. (e.g. Imagine Boston [8], Every Voice Engaged [6], and Participatory Budgeting Project [9]). In total, we reached out to 32 civic leaders among whom 14 agreed to participate in our study. These participants are operating across different regions of the United States working in projects with a common goal towards community engagement and utilizing community input for inclusive civic decision-making. 11 of the participants are male and 3 of them are female. They have an average of 12.61 years of experience working with their respective communities. Table 1 presents civic leaders we interviewed, their primary roles, and their community input analysis approaches.

3.2 Procedure

During the initial interviews, we asked the civic leaders open-ended questions regarding their needs and challenges of analyzing community input. We began by asking *What do you want to learn from the community input and why?* We also asked, *What method or technology do you use to analyze the data?* Furthermore, we asked questions around how to address the challenges of analyzing these inputs, such as, *What could help you in analyzing community-generated data?* Finally, we touched upon the desired attributes of potential analysis tools that would help them better analyze and interpret community input, asking *What features would you want in civic data analysis tools?* Based on these interviews and prior work on community engagement [50, 67], we performed rapid prototyping using illustrator tools to further engage in discussions around their technology requirements in another round of follow-up interviews with each participant where we received their feedback on the prototypes (Fig. 1). We designed the prototypes to be informative yet accessible for civic leaders with different levels of visual analytic expertise. Initially, we designed a sentiment-based prototype following sentiment analysis approaches used in digital civics [67] (Fig.1(A)). We also designed a criteria-based interface based on prior work in decision support systems [33](Fig.1(B)). We sent digital copies of these prototypes to participants through email a week prior to the follow-up interviews.

We conducted semi-structured interviews over video conferencing that lasted between 30-45 minutes. All participation was voluntary. The interviewees were not reimbursed for their participation as mentioned in the consent form provided prior to the interviews. We conducted each interview with two members of our research team to divide the roles of interviewer and note-taker. We audio-recorded the interviews and took extensive notes. Later, we transcribed over 900 minutes of collected interview audio for analysis. We used a spreadsheet application to thematically analyze the interview transcripts and notes taken during the interviews using the open coding method [31]. The main themes from our analysis are presented in the following section.

3.3 Findings from Needfinding Interviews

3.3.1 Organizing different facets of community input through visual summary. All interviewees mentioned their struggle to keep up with the increasing scale of unstructured community input text gathered from online platforms. The collected textual data is often loaded into spreadsheet applications as a flat list of comments where the analysts and decision-makers have to explore and analyze the data without helper functionalities to provide them with additional support and insights for exploration and sensemaking in an effective and efficient way. The number of comments vary from

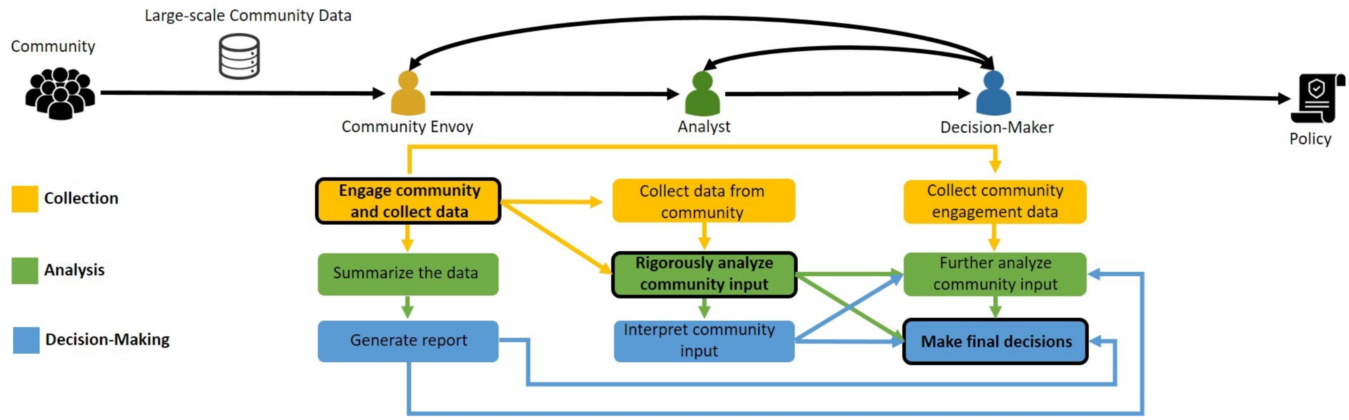

Figure 2: A simplified overview of the complex community input analysis process involving various actors. In our study, we refer to these actors as *Civic Leaders.* Among them, *Community Envoys* work as buffers between community members and decision-makers and engage community members to collect input, *Analysts* examine and interpret community input, and *Decision-Makers* make planning and policy decisions. The rectangles with black borders indicate the primary activities of these actors, but they often assume multiple roles. Depending on their assumed role, civic leaders often have different agendas and practices towards community input analysis which contribute to the complexity of the analysis process.

hundreds to thousands of comments, making analyzing the data to inform civic decision-making extremely tedious, time-intensive, and challenging. All of them wanted better organization of community input to gain a high-level understanding. In the follow-up interviews, we probed further into what kind of data organization might help them efficiently analyze community input. They wanted quick ways to gain high-level visual summaries of different facets of community input including what emotions they have expressed towards each proposal, what topics participants have discussed, and how many participants contributed to such discussions to examine and compare between proposals. However, some noted that visual summaries could be insufficient in contextualizing the community input resulting in "*uncertainties towards the data*" (N3). For example, one interviewee (N14) said of the prototypes (Fig. 1), "*Visually, the bar charts focus too much on the number of comments. This might be misleading for users to think the proposals that did not get many comments were not important.*" They mentioned the importance of preserving the actual text from the community input to combat decontextualization. Reading the original text comments also matched their "*traditional methodology*" (N2) for community input analysis. They also emphasized on a "*more interpretable interface*" (N6) that would be easier to work with.

3.3.2 Supporting role-based flexible exploration of community input. From our interviews with civic leaders, we gained a deeper understanding of the current community input analysis process, various actors involved in the process, and their needs and requirements (Fig 2). The responsibilities of these actors involve collection, analysis, and decision-making based on community input. The actors include community envoys who establish dialogue between decision-makers and community members and collect raw data, analysts who analyze and interpret community input, and decision-makers, government-affiliated authorities, who make planning and policy decisions. In practice, the transition between these responsibilities can be fluid as one individual often assumes multiple roles during the community input analysis process. For example, community envoys often summarize the community input and forward the summary results directly to decision-makers, instead of the usual raw data. Analysts often make decisions about how to interpret community input and what to add in the analysis results. Decision-makers sometimes collect community input directly from the community members, bypassing community envoys. Furthermore, depending on their assumed role, civic leaders often have different agendas and goals towards community input analysis. We found that the interviewees who are primarily involved with decision-making wanted high-level summaries of community input to help them concentrate on critical issues efficiently without interpreting every input individually. In contrast, the others, who are mainly involved with data analysis, perceived that without considering individual input, summaries alone might oversimplify the underlying narrative, leading to an incomplete interpretation of community input. One of them mentioned, "*Summaries can suppress minority viewpoints*" (N13), which may lead to marginalization. They wanted ways to drill down to individual community input and examine each input in detail. Some analysts noted, "you have to understand their specific needs and what works for them" (N10) which is important "*especially with regards to lower-income neighborhoods*" (N13). While both groups thought summaries to be useful for high-level exploration of community input, they were divided on to what extent summaries should be utilized. They wanted to negate marginalization by including all community input while maintaining reasonable time and labor requirements for the analysis process.

3.3.3 Understanding community's opinions beyond binary sentiments. All interviewees needed to understand the community's reactions towards a proposal. They emphasized the importance of understanding the "*community's emotions towards the circumstances around them and the prospects*" (N6). Another interviewee mentioned,

> *The impetus for involving the community in decision-making has changed in the last decades. Now, we are more focused on understanding how the community feels about the proposals and changes. We need to ensure that they have ownership and control over the planning and we are not imposing these ideas without their consent.*

(N9). During our discussion on what would be the ideal categorization of the community's opinions, some interviewees mentioned their familiarity with sentiment analysis in extracting opinions from the text. However, they were concerned that the binary categorization of viewpoints into positive and negative was ineffective in surfacing the community's emotions in online civic discussions (Fig. 1). One of interviewees mentioned, "*We aren't usually thinking positive, neutral or negative, we are usually thinking: what are the different categories in the way people are thinking about something*" (N11). We further asked them about the community's emotions and mentioned some popular emotion categories, such as Anger, Disgust, Surprise, Happiness, etc. However, they mentioned that some of these categories seldom appear in civic discussions, saying, "*Disgust and Surprise are not relevant to civic discussions [...]. I don't see how these emotions factor in*" (N2). In contrast, the majority of the civic leaders (11/14) specifically wanted to learn whether the community is excited or angry towards a civic issue, or if the community is concerned, happy, or if they do not care at all. Based on these insights, we crystalized five categories of emotions suitable for community input analysis - **Excitement**, **Happiness**, **Neutral**, **Concerned**, and **Angry**. **Excitement** represents people's *"enthusiasm towards civic proposals"* (N3), **Happiness** portrays their *"sense of approval"* (N11) regarding proposal contents, **Neutral** demonstrates people's lack of empathy or interest for a particular proposal, **Concerned** highlights *"apprehension towards unforeseen consequences"* (N8) that might be associated with the proposal, and finally, **Angry** portrays people's repulsion or displeasure towards the proposal. Some examples of such categorization on real-world data is presented in Table 4 where we discuss the impact of fine-grained text categorization algorithms.

3.3.4 Identifying the main discussion topics. There was a general consensus among civic leaders around extracting the main topics that are the key points of discussions on the proposal. However, they objected against the fixed categorization of topics (Fig. 1). Several interviewees (7/14) drew parallels with product review analysis, where fixed aspects of products are extracted (e.g., price, weight, etc). They emphasized on the open-ended nature of civic discussions and how pre-modeled categorization can restrict the insights that can be gained from them. One participant mentioned, "*sometimes it's easy to have predetermined categories. But often, we found, some of the most surprising things came through kind of open-coding and that often attracted most interest*" (N2). They wanted means to automatically extract underlying topics from the community in a way that reflect the unique discussion agendas without a forced characterization. This need is aligned with their current practices of using qualitative thematic analysis to analyze community input.

3.4 Design Goals

Our interviews with civic leaders surfaced several key insights and requirements. We utilized the key findings from these interviews to distill four design goals to drive the design and development of our prototype. First, we found that the unstructured nature of community input is one of the primary barriers for civic leaders to effectively explore and analyze these input. Furthermore, they wanted visual summary of various facets present in the data. Thus, we need to provide ways to organize community input and visualize it's different facets (G1). Second, depending on their roles and practices, civic leaders required to access and explore community input as both high-level summaries and low-level drill downs. To support exploration at multiple granularity, our second design goal is to support role-based multilevel exploration of community input (G2). Third, The interviewees unanimously wanted to understand community's reactions towards civic proposals. However, they were weary about gross categorization of nuanced opinions into binary positive or negative dispositions. Hence, our third design goal is to surface opinions that go beyond sentiments (G3). Finally, in keeping with the current community input analysis practices involving thematic analysis, civic leaders needed to identify the main discussion topics. However, the topics need to be automatically extracted, dynamic, and relevant to the corresponding proposal as opposed to being clustered into predefined fixed categories (G4).

4 COMMUNITYPULSE

Grounded by the insights gained from the multiple rounds of interviews with civic leaders, we designed and developed CommunityPulse, a visual analytics system to help civic leaders with community input analysis. CommunityPulse has 3 views. The Aggregate View provides a visual summary and scaffolds various facets of community input. The Detail View expands from this scaffolding and allows the inspection of individual comments (Fig 3). Finally, Compare View enhances the community input analysis by supporting comparison within or between proposals (Fig 4). All three views focus on presenting textual information while preserving the context of community input.

4.1 Scaffolding different facets of community input (G1)

To provide a high-level summary of community input, we scaffolded different facets such as proposals, discussion topics, and emotions using visualization. Prior works suggest that visualization provides a starting point for interactive exploration of patterns [93] to detect the expected and discover the unexpected from a large amount of information [90]. We took an overview + detail approach following the visual information seeking mantra [86] which is effective for text comprehension using visualization [37]. The Aggregate View provides an overview of the community input in a tabular format. It consists of a list of proposals (Fig 3(A)) and meta-information including the number of comments and number of people (Fig 3(B)). The automatically extracted topics follow that present the main discussion points for each proposal (Fig 3(C)). Finally, community emotions are shown using a stacked bar chart that provides a distribution of the community's feelings towards each proposal (Fig 3(D)).

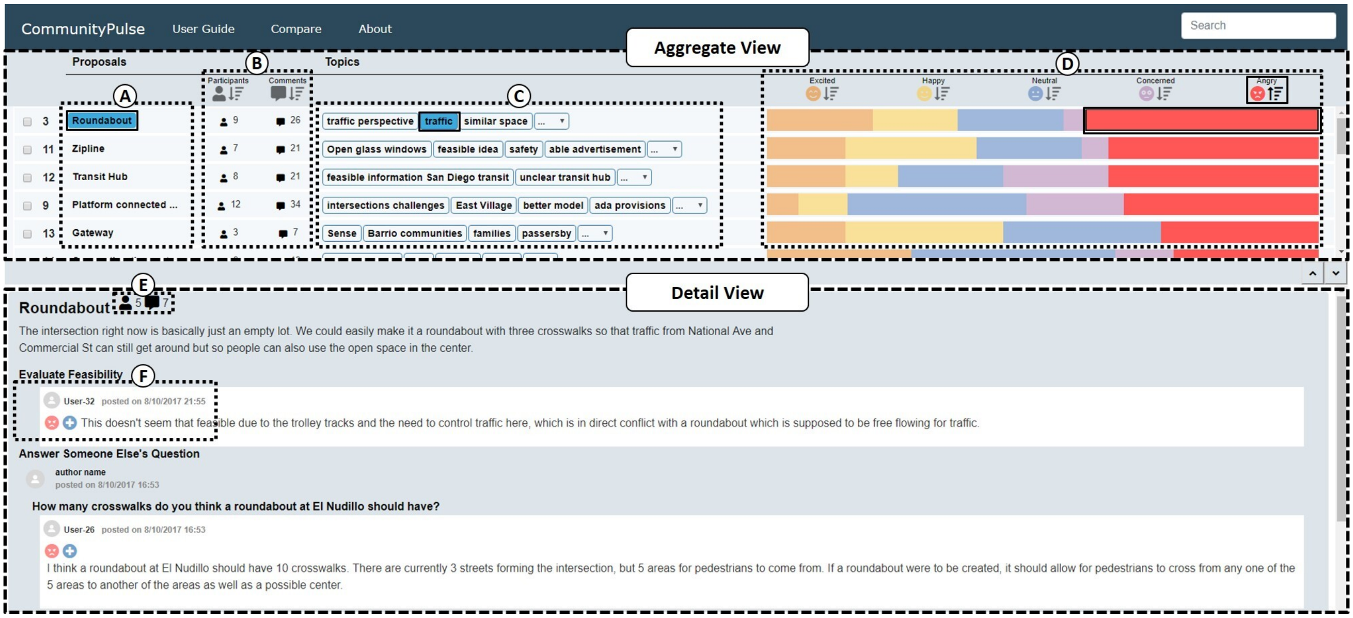

Figure 3: A snapshot of CommunityPulse. The Aggregate View shows: (A) a list of Proposals (Roundabout is selected), (B) the number of people and comments for each proposal, (C) a list of topics for each proposal (Traffic is selected), and (D) emoticons to sort the proposals based on emotions and stacked bar charts to present people's emotion distribution and drill down to actual comments (In this view, the proposals are sorted by *Angry* emotions and angry comments from Roundabout are selected). The Detail View is rendered and 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.

At the top of the bars, we used emoticons to represent each emotion and provide interactions.

The proposals are arranged row-wise to enable rapid comparison between them. The rows can be sorted in ascending or descending order based on the meta-information or emotions by clicking the interactive icons at the top. Each row contains information regarding a particular proposal. We used fixed-width columns for proposal titles, list of topics, and emotion bars to establish a tabular formation to separate these components. We also truncated the proposal titles to reduce text clutter and provided a dropdown to access the complete list of topics. Hovering on the emotion bars highlights them and displays the number of comments categorized as the highlighted emotion and the total number of comments. We used fixed-width for the emotion bars to provide the same weight to all proposals. The fixed bar-width communicates that every proposal is important regardless of the number of comments it garnered to avoid marginalizing proposals with fewer comments. To find a suitable color palette, we explored ColorBrewer and Tableau [12] to identify a divergent set of colors to reduce eye fatigue. We also reviewed Lyn Bartram's work on using colors to represent emotions in visualization to color-code the emotion bars [18]. All the components in the Aggregate View afford organization of the underlying community input. The proposal titles and meta-information provide useful information on proposals, the topics group the community input on proposals into several key discussion points, and the bar charts categorize the community's emotions towards the proposals. All of these components can be used as interactive filters. Hovering on any filter highlights and provides tooltips to show important information or signify the interaction results. Clicking on any filter displays the Detail View containing the text comments pertaining to the selected filters. Interacting with filters in the Aggregate View helps users to reach specific information clusters in the Detail View. These clusters are manageable pockets of information extracted from the community input that is compartmentalized and yet robust enough to yield actionable insights that can help make informed decisions.

The Aggregate View is geared towards civic leaders who are mainly involved with decision-making and want quick summaries of community input. It might allow civic leaders to answer analytical questions such as, *Which proposals are most discussed? Which proposals attracted the most attention from community members? Which proposals are generating Angry emotions? What topics are they interested in for a particular proposal? Which of these topics are more significant?*

4.2 Supporting role-based multi-level exploration (G2)

To support the civic leaders who are primarily involved with analysis or decision-making, CommunityPulse supports community input exploration at multiple levels of granularity. While the Aggregate View provides a high-level overview for community input exploration, the Detail View allows a deeper exploration of individual comments by working in tandem with the Aggregate View.

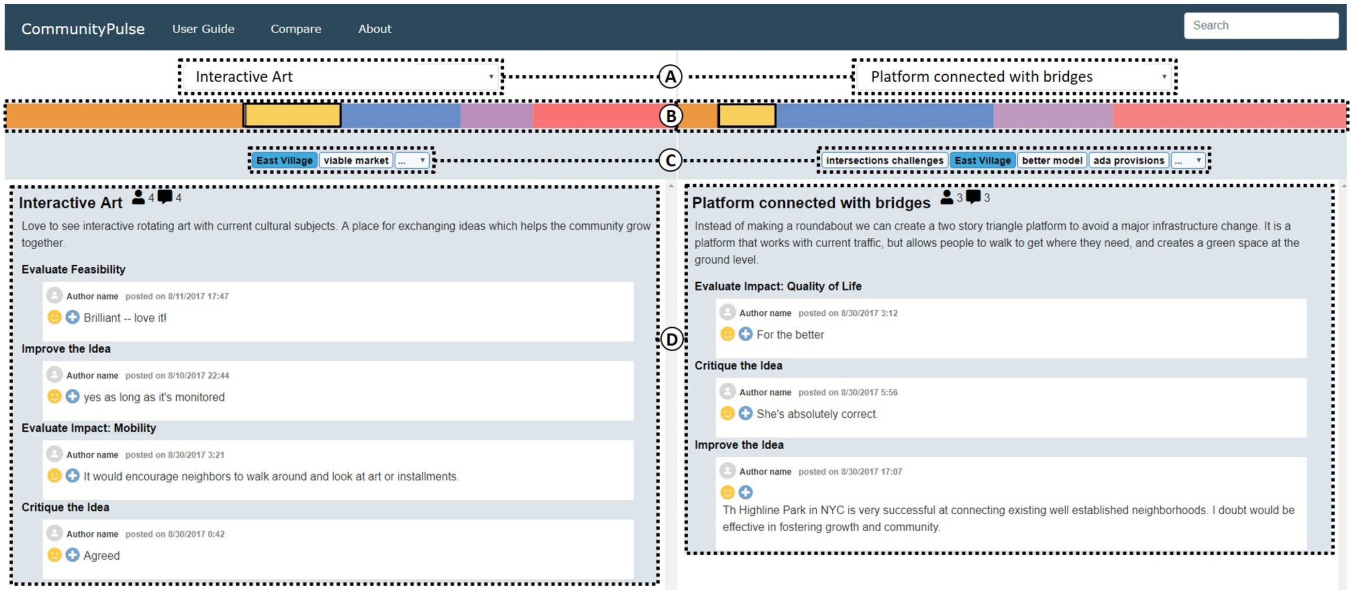

Figure 4: The Compare View contains two sections with identical components: (A) proposal dropdowns to select proposals (Interactive Art and Platform connected with bridges are selected), and (B) stacked bar charts representing emotions of the selected proposal (*Happy* is selected for both proposals), (C) a list of topics for each proposal (East Village is selected from both proposals) and (D) comments based on the selected filters.

Clicking on the proposal titles, topics, and emotion bars contracts the Aggregate View and expand the Detail View to render the filtered comments. This transition between the views is performed using a smooth animation to avoid sudden drastic changes that would otherwise break the flow of exploration and induce unnecessary cognitive load [61]. The Detail View presents the title and description of the selected proposal, followed by the numbers of comments and participants, and finally, the list of filtered comments (Fig 3(E)). Each comment contains the user avatar, user name, time of posting, and comment text. It contains an emoticon to represent the emotion associated with each comment. There is also an option to save the comment as a note (Fig 3(F)) for later analysis.

The filters can be cascaded to focus on particular proposals, topics, and emotions in conjunction. For example, a civic leader might be interested in exploring a particular proposal among dozens. Clicking the proposal title will display all comments that can number in hundreds or thousands. If the civic leader wants to focus on a specific topic, clicking it will filter the Detail View and display the comments associated with the selected topic only. Moreover, the civic leader can select any of the emotions (e.g. Angry) from the emotion bars to further filter the comments. The Detail View might allow the civic leaders to find answers to analytical questions such as, *What are the community members saying in support/oppose of this proposal? Why does the community consider this topic to be important? Why do they think this might not be good for the community? Is someone trying to influence a discussion to gain a favorable outcome?*

The Compare View can be used to compare the comments between two proposals or within the same proposal side by side. It provides filtering options to have a closer look at the community input to find important differences or similarities among proposals. Unlike the other views, Compare View is rendered in a separate window. It can be accessed by using checkboxes with each proposal and clicking the Compare button in the navigation bar (Fig. 3). There are two separate sections with proposal dropdowns, associated topics, and emotion bar charts similar to Aggregate View. The topics and emotion bars can be used as filters similar to the exploration in the Detail View. This view might allow the civic leaders to answer the following analytical questions, *Why do people like proposal X but not proposal Y? Why is the community divided on a particular proposal? Why does the community view a topic positively in one proposal but negatively in another proposal?*

4.3 Surfacing opinions that go beyond sentiments (G3)

From our needfinding interviews, we identified five emotion categories that civic leaders want to extract from the community input (Excitement, Happiness, Neutral, Concerned, Angry). We utilized text classification to identify these categories from each comment. We experimented with various techniques including support vector machines (SVM) [100], random forests (RF) [97], convolutional neural networks (CNN) [24], long-short term memory (LSTM) [65], and bidirectional encoder representations from transformers (BERT) [46]. We also experimented with commercially available text analysis software including IBM Watson Tone Analyzer and Microsoft Azure Text Analytics. We used the CrowdFlower dataset [2] to train emotion prediction models to utilize these techniques. We found that fine-tuning the pre-trained BERT model for our emotion prediction task provided the best result, yielding an F1-score of 0.912 on the CrowdFlower dataset while SVM, RF, and CNN+LSTM achieved 0.61, 0.62, and 0.72 respectively.

Compared to these text classification methods, the commercially available tools yielded poor performance with F1-scores of 0.52 and 0.55 respectively. This suggests that by effectively capturing the context of words, BERT leads to more accurate emotion classification for our community input analysis task.

To visualize the predicted emotions in the Aggregate view, we used interactive stacked bar charts to show the normalized distribution of emotion categories for corresponding proposals. We used stacked bar chart as it provides a simple yet comprehensive way to visually compare the emotion distribution between two proposals [33]. In the Detail view, we used emoticons to represent emotions at the comment level due to their ubiquity in online conversations [45].

4.4 Supporting the exploration of the discussion topics (G4)

To automatically extract discussion topics from each proposal, we applied Latent Dirichlet Allocation (LDA) [53] on relevant comments. We also applied keyphrase extraction on each topic cluster to find the keyphrase that best represents that topic. We adopted Joty and Carenini's method which is designed to deal with online conversations [64]. This method applies a syntactic filter to select only nouns and adjectives as candidates and then ranks these candidate words based on an unsupervised graph-based ranking model. Finally, it combines the adjacent words to form phrases and applies the Maximum Marginal Relevance criterion to select the relevant but not redundant topics. The combination of topic modeling and keyphrase extraction helped to identify the most important topics.

The Aggregate View presents the list of topics alongside the meta-information (proposal title, number of people, and comments) in a tabular format. The meta-information and the number of topics indicate the diversity in a proposal discussion and the amount of attention it generated. For example, a proposal with fewer topics and more participants indicate the discussion is around similar concepts. Alternatively, a large number of topics indicate a diversity of opinions.

4.5 Implementation Details

CommunityPulse is developed as a responsive and mobile-friendly web application. It uses PHP to perform back-end calculations and store user interaction logs. The front-end is designed using a combination of CSS and JavaScript. CommunityPulse is designed to be capable of handling community input datasets containing several thousand comments. The system was stress-tested in real-time with sample datasets containing over 10 thousand text comments from Reddit posts distributed in 10 proposals. It takes CommunityPulse an average of 0.614 seconds, measured over 100 attempts, to render comments from a proposal with a thousand data points. The tests were performed on a laptop computer with Intel Core i5 7th generation processor (7300HQ) and 8 gigabytes of RAM, running on local host.

5 EVALUATION WITH DOMAIN EXPERTS

In this section, we describe our evaluation process which involved 15 civic leaders who are different from the needfinding interviewees. We did not involve the interviewees from our needfinding studies for evaluating CommunityPulse because we wanted to study how CommunityPulse can appeal to potential users who were not involved in the design and development process and whether it can satisfy users' needs in a broader spectrum. We also describe the results and key findings from these interviews.

5.1 Participants and Dataset

To evaluate CommunityPulse, we engaged civic leaders with various roles and expertise in community input analysis. We approached civic leaders involved in planning, development, analysis, and decision-making. We reached out to a total of 36 civic leaders across the US who are experienced with large-scale community engagement and input analysis. We followed the Snowball method to connect with the participants using emails [23]. 15 civic leaders (7 females) responded and agreed to participate in evaluating CommunityPulse. Table 2 presents our interviewees and their goals for community input analysis. Over half of them (8/15) utilize manual coding to analyze community input and the rest use various software (7/15). Together, the interviewees have an average of 13.16 years of experience working with communities. Their breadth of expertise in various civic engagement practices at a local and national scale allowed us to observe their unique approaches to using CommunityPulse. While the table only presents their primary roles in the community input analysis process, several participants (P3, P6, and P12,) are also involved in community outreach and frequently partake in the responsibility of acting as community envoys.

CommunityPulse is designed to be data agnostic. It can be populated with any dataset with a list of discussion agendas and the associated text comments. However, popular community engagement platforms for public consultation such as Nextdoor [10] does not provide APIs to collect civic discussion data for academic purposes. Furthermore, after the consultation, the actual comments are often removed from public access and replaced with a summary of the analysis result or a set of decisions. Although CommunityPulse can operate on Twitter and Reddit discussions, the discussion dynamics in such social media platforms are different than online community engagement platforms and hence are not suitable for proper evaluation. Eventually, the data we used to evaluate CommunityPulse was collected from a real-world urban planning project that involved redesigning efforts in a major US city using an online community engagement platform called CommunityCrit [72]. The data was collected through a partnership with a local planning group in San Diego who were leading a redesign effort for a major downtown street called the"14th Street Promenade". To collect community input on the proposals to redesign the street, the planning group used CommunityCrit [72] and collected community-generated data in the form of short free-form text comments. The community members shared new ideas, reviewed and commented on existing ideas, and provided suggestions for improvements. During the deployment phase of 4 weeks, the planning group was able to collect over 352 qualitative free-form textual community input from 39 community members on 19 civic proposals. We used this data collected from a real-world scenario to evaluate CommunityPulse. We used the same data for all interviewees to reduce confounding factors due to variation in the dataset. We reflect on the challenges of accessing community input for evaluation in Section 6.

Table 2: This table shows the civic leaders who evaluated CommunityPulse, their primary roles in community input analysis, their analysis goals while using CommunityPulse, and their interactions with CommunityPulse's various features (Filter comments by proposals, topics, emotion bar charts, sorting proposals by clicking on emoticons, and making notes from the detail view) captured in their respective interaction logs. The interactions are presented in percentage distribution rounded to the nearest integer using a heatmap where the color saturation corresponds to the percentage of interaction.

P(#)	Primary Roles	Analysis Goals	Proposals (%)	Topics (%)	Emotion Bars (%)	Emoticons (%)	Notes (%)
P1	Decision-Maker	Understand reflections	13	17	51	18	2
P2	Decision-Maker	Understand reflections	23	13	46	14	4
P3	Decision-Maker	Understand reflections	30	9	42	18	2
P4	Analyst	Understand reflections	28	21	27	22	3
P5	Analyst	Understand reflections	41	27	19	11	2
P6	Decision-Maker	Identify insights	11	19	41	22	7
P7	Analyst	Identify insights	25	33	19	12	11
P8	Decision-Maker	Identify insights	18	43	23	9	10
P9	Analyst	Identify insights	12	32	33	14	10
P10	Analyst	Identify insights	11	21	32	25	11
P11	Analyst	Identify insights	13	27	42	16	2
P12	Decision-Maker	Identify main topics	17	36	32	11	4
P13	Analyst	Identify main topics	35	36	14	9	7
P14	Analyst	Identify main topics	16	42	13	11	18
P15	Decision-Maker	Identify main topics	14	45	17	13	11

5.2 Procedure

We deployed CommunityPulse in a public domain. We sent the link with a written introduction to CommunityPulse's functionalities via emails to the civic leaders who agreed to participate in evaluating CommunityPulse. In the emails, we asked them to explore the system at their own times and conveniences. We also asked for interviews to discuss their experiences of working with CommunityPulse after two weeks, potential issues that emerged, and how it can help them with community input analysis. We scheduled the interviews at least two weeks after we reached out to the civic leaders who agreed to participate to provide them sufficient time to explore CommunityPulse's various features. The average difference between the time we reached out to them and the time of interview was 18.4 days apart. We collected time-stamped usage logs of their interactions with CommuntyPulse to monitor and reconstruct their usage patterns. The usage logs show that each participant had 3 or more different exploration sessions where they spent an average of 20.42 minutes per session. On average, each participant spent a total of 52.61 minutes exploring the data using CommunityPulse. Table 2 presents a distribution of their interactions with various features. We conducted the interviews in a semi-structured manner using video conferencing. Each lasted around 45-60 minutes and the interviewees shared their screens with us. We recorded and transcribed the audio and shared screens for analysis. In the interviews, we asked them open-ended questions regarding their experience of using CommunityPulse. For example, we asked, *What are your goals when analyzing community input? How did you use CommunityPulse to achieve these goals? What features did you find useful and which ones you did not? How does CommunityPulse compare with your current analysis method? How can CommunityPulse be improved?* We also encouraged them to provide specific examples of how they used different features to get useful information from community input.

The interviews were conducted with two members of our research team and extensive notes were taken. We thematically analyzed the collected notes and over 500 minutes of interview audio transcripts using the open coding method [31]. Two members of our research team independently coded the data at the sentence level using a spreadsheet application. The inter-coder reliability was 0.91, measured using Krippendorff's alpha [68]. We had several iterations of discussions among the research team to condense the codes into the themes presented in Table 3. We also utilized the captured usage logs to complement the information they divulged in the interviews.

5.3 Findings from Interviews with Experts

Analysis of participants' usage logs and interviews revealed insights on their analysis goals, strategies, how CommunityPulse helped them to identify critical information and potential issued they faced.

5.3.1 CommunityPulse can provide a structure for community input analysis. One of the main design goals of CommunityPulse was to organize textual data to support multifaceted and systematic exploration of community input. To this end, we used a combination of visualization and text analysis techniques to structure data around visualization, topics, and emotions. To better understand the usefulness of our approach, we analyzed participants' interactive behavior using a Markov Model [52]. Fig. 5 depicts their interactions with proposal, topics, emoticons, and emoticon bars. Fig. 5 suggests two recurring high-level strategies employed by participants for investigating community input. In the first strategy

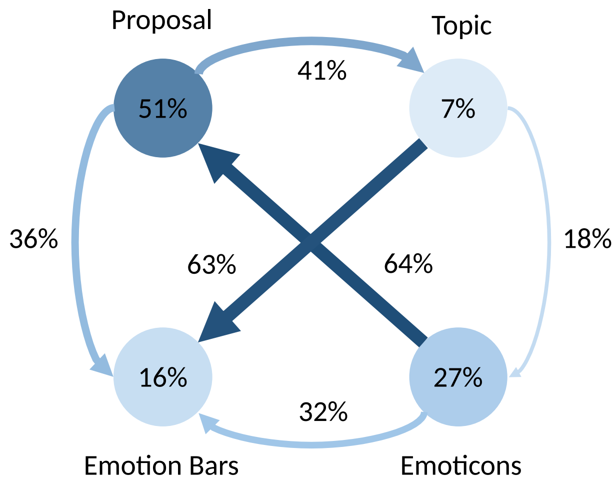

Figure 5: Analysis of interview participants' interactive behavior with proposals, topics, emotion bars, and emoticons, depicted by circles. The arrows depict the percentage of interactions between the two states. The circles and arrows are color-coded where the color saturation corresponds to the percentage of interactions. Two high-level strategies (STs) emerge from these patterns: (ST1): Proposal →Topic →Emotion Bars and (ST2): Emoticons →Proposals →Topic. These patterns demonstrate the differences in participants' approaches to community input analysis.

(ST1), participants started with proposals, then moved to topics, and emotion bars. In the second strategy (ST2), participants started by sorting the proposals using emoticons, then moved to sorted proposals, topics, and eventually emotion bars. By taking a closer look at their usage patterns (Table 2) and their interview transcripts, we found that ST1 was mostly employed when participants' goal was to closely investigate community input in detail and ST2 was utilized for a high-level overview before drilling down to individual comments. Regardless of the differences between the two high-level strategies and their goals, we observed that to carry them out, participants utilized the interactive components from the Aggregate view (proposals, topics, emoticons, and emotion bars). These recurring analysis strategies to identify critical information suggest that our approach towards structuring the community input using visualization and text analysis is a promising direction. P11, who is undertaking the planning of development of an open community space in one of the major cities in the United States mentioned,

> *I can use this system to quickly measure the temperature of the community as we progress into the planning process. It can also help us make a stronger case for the changes by having a record of community's interests and emotions that can help you articulate and justify your decision.*

We also investigated the relationship between the participants' primary roles and their interactive behaviors. However, we did not find any distinctly identifiable patterns, which might suggest that civic leaders' analysis goals are the driving force behind community input analysis. This is aligned with the findings of our needfinding studies in section 3.3.2 that highlights how civic leaders roles can be fluid as they assume different roles in the community analysis process.

5.3.2 Surfacing emotions can help to understand community's reflections. The majority of the participants (13/15) thought that emotions can be useful indicators of the community's perspectives towards civic proposals and for measuring the pulse of the community. One participant mentioned, "*This is a great way of gathering all the feelings of the people, their concerns, and attitudes*" (P6). P1 brought a specific example from his exploration of CommunityPulse, where he sorted the proposals based on the emotion *Angry* and identified that the proposal **Roundabout** have almost half of the comments (11 out of 26) categorized as *Angry*. He mentioned, "*The first three topics [of the proposal] was about traffic, it was clear that they are not happy about the traffic situation [...]*". It enabled him to identify community's negativity towards the proposal. The emotion bars can also provide accounting for the community's support or opposition towards a proposal as P2 mentioned, "*I see that by saying Happy, the people in the community are endorsing the idea and it makes a lot of sense.*". P11 compared his experience of working with CommunityPulse with their current analysis process using spreadsheets full of comments on different proposals. He mentioned,

> *When analyzing community input manually, different analysts can have different subjective interpretations of comments. The emotion graphs (bars) give a neutral interpretation based on the data and it has a great value in the analysis process. This system can help decision-makers to understand people's reactions early on which is very important for decision-making. The crux of it is to be able to associate emotions as objectively as possible with people's comments and make them readily observable.*

5.3.3 CommunityPulse can help to identify where community's priorities lie. Civic leaders wanted to examine the community's preferences towards civic proposals while avoiding redundancies in comment text. The majority of interviewees (12/15) mentioned CommunityPulse helped them to identify the community's priorities by utilizing the extracted topics and meta-information. One participant mentioned, "*It is a very useful system that is capable of delivering quick snapshots of major trends in public values and emotions in a compact way*" (P15). P11 mentioned,

> *These aspects (topics) can be used as the starting point of your analysis to get a sense of the most talked about aspects and identify where the community's concerns or supports lie as the comments come in. Once you have a grasp of their priorities, you can look at it from different angles to really make sense of what the community wants.*

P9 presented an example scenario when he explored the proposal **Interactive Art** and observed from the meta-information that there are 10 users and 28 comments. He filtered the comments by the topic *East Village* and immediately identified that all 10 users were

Table 3: This table shows themes that emerged from analyzing the interviews with the civic leaders to evaluate CommunityPulse. The codes associated with the themes and their description is also presented in the table.

Themes	Codes	Descriptions
Community's reflections	Emotions towards proposals, understanding points of views, support or oppose an idea	Reflecting on community's perspectives
Community's priorities	Common trends, preferences, opinions of interest, important topics	Identifying what matters most for the community
Hidden insights	using combinations of filters, identify easy to miss details, hidden patterns	Identifying hidden insights present in the data
Speeding up analysis	Quick snapshot of public's opinions, automating analysis steps, managing large-scale data	Accelerating the community input analysis process
Visual analytics	Access to text comments, gathering insights, simplicity, ease of use	Interactions with visual analytics to interpret community input
Concerns and caveats	Lack of user input, misinterpretation, learning curve	Concerns regarding CommunityPulse

commenting about East Village which constituted 17 out of 28 comments. It demonstrated the community's interest around that area which helped him identify the community's priorities. He mentioned, "*It's really helpful to have the data categorized and most common trends [identified] for you automatically. Usually, I have to do it myself.*".

5.3.4 *Multilevel exploration can surface hidden insights buried in the data.*

Some interviewees mentioned that combining various features and filters, they were able to extract hidden insights from the data that would have been difficult, if not impossible to extract using their current community input analysis practices. Here, we provide two examples of such resourceful exploratory analysis.

When exploring the proposal **El Parquecito**, P6 stumbled upon an anomaly when he filtered the comments by the topic *Own Space*. He observed that two *Angry* comments came from one user, user-44. He dug further by filtering only the angry comments from the whole proposal using emotion bars. He discovered from the meta-information at the Detail View that all three angry comments on this proposal came from the same user, user-44. P6 read the comments to find that user-44 was arguing for strict restrictions on private motor traffic. He also found revealing patterns in the comments that made him to suspect that user-44 might not be a resident of that area. He conjectured that user-44 was biasing the discussion, saying, "*While every opinion matters, some people might direct the discussion in a particular way by being vocal and suppressing others. It's important to identify that you don't have a biased sample*". He was intrigued to find such critical information quickly which would usually require him to closely read all comments on spreadsheets. He also mentioned that he might have completely missed the information without the *Angry* emotion categorization.

Another example came from P14. She mentioned exploring topics from different proposals when she came across two different proposals, **Interactive Art** and **Platform Connected with Bridges** and both of them have a topic named *East Village*. She wanted to know how this particular topic related to two different proposals. She used the Compare View to load two proposals side by side. Then she filtered the comments based on the topic in question. She discerned that while the comments discussed two different ideas, an art installation and a platform with bridges, they both focused on areas around East Village. After reading the comments, she realized that these two proposals complemented each other and could be merged into a single project. She mentioned in the interview, "*I use separate files in [Microsoft] Excel [for different proposals] and flip through them, [...] but this [Compare View], with the topics solves that problem*".

5.3.5 *CommunityPulse can accelerate community input analysis process.*

The current community input analysis method is tedious and time-intensive that often depends on manually coding large amounts of text. After exploring CommunityPulse to explore and analyze text comments, our interviewees (11/15) emphasized how the computational approach can potentially accelerate the process without loss of rigor. One participant mentioned, "*I am impressed that machine learning can be used to review all the comments, which make it much easier to tease out the most relevant information without having to go through a tedious annotation process*" (P4). The advantage becomes more prevalent as the scale of community input increases which requires categorization of data for focused exploration, as P6 mentioned,

> *This system has great potential and will be particularly useful where the community is large and the civic leaders need a quick way to get a glimpse of public opinions. Clearly, there is a time-saving advantage of using this tool.*

5.3.6 *Visual analytics can help to effectively interpret community input.*

Similar to the needfinding interview participants, we found that some civic leaders who evaluated CommunityPulse also preferred actual text over aggregated statistics, as P7 mentioned, "*It's good that CommunityPulse does not generate anything that looks like a statistic and shows more comments, [that] allows me to make specific inferences*". Others mentioned the inefficiencies inherent in their current community input analysis practices that complicates gathering important insights, especially during unforeseen events such as the recent pandemic. One such participants (P10) mentioned,

> *Due to the pandemic [COVID-19], online engagement has increased exponentially, which is the only viable source for gathering comments from the community.*

> *And the only means to analyzing these data we have is to dump everything together and try to make sense of individual data pieces. This system [CommunityPulse] lets me analyze data in an effective way and I can interpret the data based on these guidance [topics and emotion].*

The participants also found the visualizations incorporated intuitive and easy to work with. Regarding the visual interface of CommunityPulse, P11 mentioned, "*It [CommunityPulse] is very intuitive and not complicated at all. Overall, its very well-designed and easy on the eyes.*" Others emphasized how simplicity is critical for such tools in practice, especially for the civic leaders who are not trained extensively in complex data visualization. One such participant (P14) mentioned,

> *This tool presents a lot of data in a compact way. Although it is a little crowded, it is an acceptable tradeoff. There exists is a danger of making the system too complex for the average users that might prevent them from interacting with it.*

5.4 Concerns and Caveats

Despite extolling CommunityPulse for enabling more effective community input analysis by supporting and augmenting their current analysis process, some civic leaders we interviewed highlighted some limitations of our system. P3 and P5 added two caveats to depending solely on the visual summary for interpreting the underlying data. P3 mentioned,

> *It [emotion bars] is good for experienced community engagement analysts and professionals, but there is a danger for a novice user to make incorrect inferences that are not actually supported by the data [comments].*

Similarly, P5 highlighted, "*"There is a learning curve involved and it takes some time to start understanding the bigger picture*". In addition, some participants (3/15) who wanted to further analyze community input, wanted more interaction functionalities to add their own topics, on top of the automatically extracted ones. One participant (P14) mentioned, "*I can take notes [using the Notes function] from individual comments, but sometimes when we dig deeper, we find some patterns that we want to add as topic categories*".

6 DISCUSSIONS

Our interviews with civic leaders and their usage logs highlighted how civic leaders can utilize CommunityPulse for analyzing community input. The majority of the experts we interviewed for validation found CommunityPulse to be effective and efficient in extracting community's emotions, main discussion topics, and actionable insights in a fraction of time with negligible effort. In this section, we discuss lessons learned from our study and possible future directions.

6.1 Combining interactive visualization and text analysis for community input analysis

Previous researchers have conducted studies in the civic domain that mostly focused on how the public engages with the early stages of civic design [17, 38–41]. However, such engagement models do not necessarily translate to the later stages that includes community input analysis [58, 73]. Responding to the call from prior researchers [58, 73] to design specialized technologies to enable civic leaders to analyze community input, we extended prior work by conducting multiple rounds of interviews with civic leaders and performing rapid prototyping which resulted in deeper insights into the analysis process involving various actors and specific requirements for such interventions. These requirements include finer categorization beyond sentiments, multi-level exploration to investigate community input, and surfacing latent themes.

Exploration, analysis, categorization, and sensemaking of large-scale informal language textual comments are well-known for being particularly challenging tasks [61, 62]. In the civic domain, community input analysis adds to this challenge due to the open-ended nature of community discussions. In our work, we experimented with various text analysis techniques including SVM, RF, CNN, LSTM, BERT, etc. and iteratively refined them to achieve reasonable accuracy for community input analysis. We echo previous work that demonstrated the absence of a universally generalized model for text analysis and how off-the-shelf methods lead to ineffective systems [36]. For real-world text analysis tasks, especially when accuracy is critical, future work should use an iterative approach to identify suitable, domain-specific, and reliable solutions capable of addressing the problem in question.

Furthermore, similar to previous work, we argue that text analysis alone might not be sufficient for the civic domain due to the unstructured and ambiguous nature of the community input [20, 73]. However, visualization enables the extraction of critical information from a heap of unstructured data by facilitating exploration of patterns [90, 93]. In our work, we utilized visualization to scaffold multiple facets of community input that were extracted using text analysis to allow multi-level exploration of large-scale textual community input. Our evaluation with civic leaders suggested the effectiveness of utilizing easy-to-interpret visual encoding such as emoticons and bar charts alongside actual text information that resembled their current practices for community input analysis. This is in keeping with previous work in other domains where the users preferred actual text information that matched their existing practices [89]. We also observed two exploration strategies emerge from the participants' usage logs that further corroborate the efficacy of combining visualization and text analysis to facilitate a structured, multifaceted, and systematic exploration of community input.

Our evaluation suggests different ways in which civic leaders can utilize the combination of text analysis and visualization to identify hidden patterns and actionable insights from often ambiguous and multi-faceted public input in civic discussions. We believe that such a combination could be helpful in other domains where people engage in online discussions and share opinions freely, such as opinion polling [1], discussions in massive open online courses to receive feedback [81], request for comments (RFC) discussions in Wikipedia between experts [99], and customer reviews in online marketplace [79].

Table 4: This table shows a comparison between classifying community input from the dataset we used to evaluate CommunityPulse by the experts. Compared to binary sentiment categories (Positive or Negative) Our fine-grained categorizations can help experts to identify the underlying feelings behind an individual's attitudes towards different proposals.

Community Input Text	Sentiment Categories	CommunityPulse's Emotion Categories
This is fantastic! Art is a wonderful way to connect people within a community.	Positive	Excitement
I like the idea. It seems feasible mobility-wise.	Positive	Happy
Meh, no comments.	Neutral	Neutral
I wonder how it would impact mobility.	Negative	Concerned
This idea is nonsense and a waste of time.	Negative	Angry
Any ideas on how to stop the homeless people from bathing in the fountain?	Negative	Concerned
You cant! We should scrap this idea in its entirety.	Negative	Angry

6.2 Surfacing community's opinions beyond sentiments

Our needfinding study showed that the civic leaders found current sentiment analysis to be ineffective in surfacing nuanced opinions from community input and wanted to go beyond positive, negative, or neutral categories to know how the community feels about different proposals. Previous research in cognitive science emphasized how emotions can be important and influential drivers for decision-making [70]. Our emotion categories helped civic leaders to gain a deeper understanding of community's feelings towards different proposals. Our participants specifically mentioned that *Angry* and *Concerned* helped them to identify issues and main objections, while *Excitement*, and *Happiness* enabled them to recognize community's support towards the proposals. CommunityPulse empowered them to not only understand *what* the community felt about the proposals, but also identify *why* they felt that way. Table 4 provides several examples of why our approach towards providing fine-grained categorization of community input might be useful for civic leaders to tease apart sentiments and the underlying feelings that shape an individual's attitudes towards proposals. In the table we see the same community input text classified using binary sentiment analysis classifications (Positive and Negative) and using our categorization. We can see our fine-grained categorization delineates emotions from these comments as opposed to grouping them in to generic classes. For example, consider the inputs **"I wonder how it would impact mobility."** and **"This idea is nonsense and a waste of time."** made on the same civic proposal. While sentiment analysis categorizes both comments as *Negative*, our proposed system categorizes them as *Concerned* and *Angry*, separating a legitimate concern from venting frustrations, enabling civic leaders to assign appropriate value to these inputs when analyzing and interpreting the community input datasets.

However, surfacing more nuanced emotions from community input raises a caveat around their impact on the civic leaders while making rational decisions. While emotions can be beneficial for rational decision-making [75], research shows the framing of emotions can significantly impact the way humans make decisions [44]. Long-term deployment and examination of visual analytics systems such as CommunityPulse is needed to study the impact of surfacing emotions on civic decision-making. Furthermore, the emotion categorization should be domain-specific and deeper investigation of users' necessities is essential to identify suitable emotion categories. For example, experts from other domains such as customer reviews, massive open online courses, or e-governance website conversations might not find our emotion categorization to be appropriate indicators of actionable feedback from their users [47, 84]. Our interviews with civic leaders led us to identify emotion categories that enabled our participants to more effectively analyze community's reflections. Researchers in future could further investigate extension of sentiment analysis and identifying appropriate categories for more fine-grained analysis of opinionated text that matches the requirements of target users in different domains.

Another monition to consider is the possibility of inadvertently incurring biased analysis when designing algorithmic interventions to surface opinions from text information. Our careful selection of text analysis methods to extract topics and emotions resulted in few objections from the participants due to high accuracy. However, as P3 cautioned during evaluating CommunityPulse, novice community input analysts might be susceptible to making incorrect assumptions unsupported by the data. While seasoned analysts are more suited no navigate the intricacies of underlying community input by utilizing both algorithm output and close scrutiny of the text data, the risk of influence on data analysis following an algorithmic pipeline is ever-present [32]. Researchers who are adamant to maintaining algorithmic fairness might adhere to methods that increase the explainability of results produced by machine learning algorithms [22]. Visualization can play a key role here by enhancing the interpretability of complex machine learning models by surfacing and communicating uncertainly and confidence levels present in the results to greatly reduce misinterpretation [35, 82]. Another alternative could be to increase user agency and provide them with the means to object or annotate their disagreement on the categorization of topics or emotions produced by the algorithms. This might prove useful to combating misinterpretation by notifying potential inaccuracies produced by the algorithm to other analysts who are using such systems and instill trust in the analysis process [32, 39].

6.3 Undertaking complex sociotechnical problems

Designing innovative solutions for complex sociotechnical problems is challenging due to the intricate interplay among broader societal, political, economic, cultural, organizational, and computational components, and their long-term impact on society [29, 77, 78]. Civic decision-making is a complex sociotechnical problem that has long been considered as a wicked problem [30]. Particularly, understanding and adapting to experts’ perspectives, computational capabilities, and sociopolitical constraints is a difficult task. In this study, we took a human-centered approach to design a visual analytics system that enables community input analysis to aid civic decision-making. While design studies in other domains, such as journalism [27] and data monitoring [28] have advocated for adoption; adopting new technology is extremely challenging in the context of complex sociotechnical problems [78, 87]. Despite the rapport we built with our participants during the study, their necessity to adopt new technology, and their eagerness to adopt our system, we experienced challenges regarding adoption. This was due to the sensitivity of community input (personal information embedded into the comments, such as financial, livelihood, neighborhood, etc.) and bureaucratic issues of convincing the upper echelon (e.g. government, management) to provide infrastructural support for adoption. Future researchers could explore the design of visual analytics platforms that employ data encryption and differential privacy [48, 60] to instill transparency and users’ trust in the system [39]. However, extensive study is required to investigate the dynamics between preserving privacy and allowing drill-down to sensitive information for inferencing and sensemaking in complex sociotechnical problem domains.

7 LIMITATIONS AND FUTURE WORK

CommunityPulse was effective in surfacing the community’s reflections, priorities, and hidden insights in our study. However, our work explored only a few datasets and the experts evaluated the system for a few weeks. Additional work is needed to further study the long-term impact of such systems in practice and to extend it for more general use. To begin with, CommunityPulse does not allow to add civic leaders’ own topics on top of the automatically extracted ones. Furthermore, subtle categorizations such as criteria or issues remain elusive. Criteria are quantifiable aspects such as cost, safety, equity, etc., and issues are notions that neither support or oppose the proposals but are necessary to be taken into considerations while making civic decisions. These nuanced concepts are interwoven within community input which is challenging for text analysis methods to surface [38] due to the lack of labeled data and appropriate prediction models. In the future, CommunityPulse can be augmented with more interaction functionalities to enable civic leaders to add their own data labels. Additionally, these labels can be used to train models that augment statistical approaches with the human in the loop paradigm [15, 43] to identify novel information categories, such as criteria and issues. However, civic leaders might have different interpretations of the community input depending on their analytical skills and their perspectives and their labeled data could lead to biased models. This could be mitigated by establishing moderation functionalities [63, 99] to help civic leaders reach consensus before making decisions that would impact the whole community. To investigate the generalizability of our approach, in the future, we will expose CommunityPulse to more diverse sets of data with different dynamics, such as forum-like discussion or nested social media conversation where users’ can use their own data where users’ can use their own data to evaluate the system. This might help us to identify values and issues with our approach in a more general setting.

In the future, CommunityPulse could be augmented with demographic information, which is an important social indicator [16] that can enable civic leaders to further understand the social dynamics and motives behind the comments made by community members. Future research can focus on investigating ways to integrate demographic information while maintaining the anonymity of users. This line of research can also investigate to find answers to questions around data uncertainty [55] and incompleteness [42]. There is a longstanding apprehension towards data aggregation in civic decision-making [73] as well as information visualization in general [32] where the collected data can be inappropriately interpreted as a representative of the complete domain, which is inaccurate in the presence of uncertainty and incompleteness in data [19]. Including demographic information and integrating uncertainty and incompleteness indicators with the system could enable civic leaders to account for and address the incompleteness in their collected data while inferencing. This can further help them to combat marginalization by identifying indicative patterns in the data. For example, if the community input only accounts for a dominant race, sex, or nationality in a heterogeneous community, it might be an indicator of potential marginalization in the data.

Another avenue of improvement would be to enable visualization of temporal changes of emotions. It might help civic leaders in further understanding nuanced community opinions by identifying changes in community members’ feelings towards a proposal, and perhaps more importantly, gaining insight into *why* the change happened. Furthermore, it can also help civic leaders keep track of how the community-generated ideas evolved from initial concepts to specific artifacts to identify community’s greater involvement with improving an idea as well as the viability of materializing these ideas. To that end, CommunityPulse can be augmented with a provenance module [83, 95, 101] capable of visualizing the chronological changes in community’s emotions. It will also enable researchers to study how changes in emotions might impact community input analysis and civic decision-making.

Furthermore, this work is our first step towards addressing the the community input analysis problem. In this study, we focused on civic leaders’ and learned about their requirements to inform the design of technologies that can cater to their specific requirements around community input analysis. In future, we will draw on the lessons learned from this study and expand the system into a citizensourcing model that involves multiple stakeholders and include community members who are the other side of the participatory democracy equation [13, 17, 41, 57, 59, 73, 76, 80]. To that end, CommunityPulse could be integrated with online community engagement platforms (e.g., DemocracyOS [5], Nextdoor [10], and CommunityCrit [72]) and be extended to allow community members to explore and validate the analysis results and take part in the decision-making process. Such extensions could improve communication between them and enable community members to understand why or how their opinions impacted civic decisions. It could further provide them a way to have a say on how these decisions are shaped. This can mitigate the lack of trust among civic leaders and community members, which is a wicked problem in civics [30, 39, 58]. This iterative exchange and information sharing between community members and civic leaders could not only remove the communication barriers and instill trust between them but also help to refine and inform civic decision-making.

8 CONCLUSION

We conducted an in-depth needfinding study with 14 civic leaders to understand their goals, challenges, and needs to meaningfully explore and analyze community input on important civic proposals. Based on these needfinding results, we crystalized four design goals. We used these design goals to drive the development of an interactive visualization system called CommunityPulse. Our system combines text analysis with visualization to scaffold multiple facets and support the multi-level exploration of community input. Another group of 15 civic leaders with diverse backgrounds and expertise in community input analysis evaluated the usefulness of CommunityPulse. They found the system to be effective in eliciting the community's reflections and priorities. They also found emotion categories helpful in gaining more in-depth insights into the community's dispositions. Our findings indicated that CommunityPulse enhanced civic leaders' current methods of data analysis by enabling them to extract hidden insights that can facilitate critical decision-making. We discussed the challenges faced, lessons learned, and future research directions. The insights gained during our needfinding interviews, crystalizing design goals, and evaluating CommunityPulse in the civic domain could be beneficial for building decision support systems involving large-scale text input from broad participation in online communities. To enable broader access to CommunityPulse, we will make it publicly available for public data analysis. It will allow us to further study the potential impact and utility of our system across a diverse range of domains.

ACKNOWLEDGMENTS

Thanks to Steven Dow, Maggie Chan, Jiayi Zheng, Diana V. Nguyen, Yabo (Olivia) Shi, and Andres Baez for their contributions during the early conception of this project. We also thank all of the participants in the needfinding and evaluation interviews for their dedicated time and effort to the success of this project. Thanks also to the members of the HCI-VIS lab for their valuable feedback, with a special thanks to Rolando J. Franqui Nadal for his assistant with qualitative data analysis.

REFERENCES

[1] 2019. Pol.is. https://pol.is/gov. Accessed: December 2019.
[2] 2019. Sentiment Analysis Emotion in Text. https://www.crowdflower.com/data/sentiment-analysis-emotion-text/. Accessed: April 2019.
[3] 2020. ATLAS.ti. https://atlasti.com/. Accessed: June 2020.
[4] 2020. Dedoose. https://www.dedoose.com/. Accessed: September 2020.
[5] 2020. DemocracyOS. http://democracyos.org. Accessed: September 2020.
[6] 2020. Every Voice Engaged. https://everyvoiceengaged.org. Accessed: March 2020.
[7] 2020. Free open-source participatory democracy for cities and organizations. https://decidim.org. Accessed: January 2020.
[8] 2020. Imagine Boston. https://imagine.boston.gov. Accessed: December 2020.
[9] 2020. Participatory Budget Project. https://www.participatorybudgeting.org. Accessed: March 2020.
[10] 2021. NextDoor. https://nextdoor.com. Accessed: January 2021.
[11] 2021. NVivo. https://www.qsrinternational.com/nvivo-qualitative-data-analysis-software/home. Accessed: January, 2021.
[12] 2021. Tableau. https://www.tableau.com/. Accessed: January 2021.
[13] Tanja Aitamurto, Kaiping Chen, Ahmed Cherif, Jorge Saldivar Galli, and Luis Santana. 2016. Civic CrowdAnalytics: Making sense of crowdsourced civic input with big data tools. In *Proceedings of the 20th International Academic Mindtrek Conference*. 86–94.
[14] Mohammad Alharbi and Robert S Laramee. 2019. Sos textvis: An extended survey of surveys on text visualization. *Computers* 8, 1 (2019), 17.
[15] JE Allen, Curry I Guinn, and Eric Horvtz. 1999. Mixed-initiative interaction. *IEEE Intelligent Systems and their Applications* 14, 5 (1999), 14–23.
[16] Jon Anson. 1991. Demographic indices as social indicators. *Environment and Planning A* 23, 3 (1991), 433–446.
[17] Mariam Asad, Christopher A Le Dantec, Becky Nielsen, and Kate Diedrick. 2017. Creating a Sociotechnical API: Designing City-Scale Community Engagement. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 2295–2306.
[18] Lyn Bartram, Abhisekh Patra, and Maureen Stone. 2017. Affective Color in Visualization. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) *(CHI '17)*. ACM, New York, NY, USA, 1364–1374.
[19] Eric PS Baumer and Micki McGee. 2019. Speaking on behalf of: Representation, delegation, and authority in computational text analysis. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. 163–169.
[20] Eric PS Baumer, David Mimno, Shion Guha, Emily Quan, and Geri K Gay. 2017. Comparing grounded theory and topic modeling: Extreme divergence or unlikely convergence? *Journal of the Association for Information Science and Technology* 68, 6 (2017), 1397–1410.
[21] Eric PS Baumer, Jordan Sinclair, David Hubin, and Bill Tomlinson. 2009. metaViz: Visualizing Computationally Identified Metaphors in Political Blogs. In *2009 International Conference on Computational Science and Engineering*, Vol. 4. IEEE, 389–394.
[22] Umang Bhatt, Alice Xiang, Shubham Sharma, Adrian Weller, Ankur Taly, Yunhan Jia, Joydeep Ghosh, Ruchir Puri, José MF Moura, and Peter Eckersley. 2020. Explainable machine learning in deployment. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 648–657.
[23] Patrick Biernacki and Dan Waldorf. 1981. Snowball sampling: Problems and techniques of chain referral sampling. *Sociological methods & research* 10, 2 (1981), 141–163.
[24] Mondher Bouazizi and Tomoaki Ohtsuki. 2017. A pattern-based approach for multi-class sentiment analysis in Twitter. *IEEE Access* 5 (2017), 20617–20639.
[25] Daren C Brabham. 2013. *Using crowdsourcing in government*. IBM Center for the Business of Government.
[26] Daren C Brabham, Thomas W Sanchez, and Keith Bartholomew. 2010. Crowdsourcing public participation in transit planning: preliminary results from the next stop design case. In *TRB 89th Annual Meeting Compendium*.
[27] Matthew Brehmer, Stephen Ingram, Jonathan Stray, and Tamara Munzner. 2014. Overview: The design, adoption, and analysis of a visual document mining tool for investigative journalists. *IEEE transactions on visualization and computer graphics* 20, 12 (2014), 2271–2280.
[28] Matthew Brehmer, Jocelyn Ng, Kevin Tate, and Tamara Munzner. 2015. Matches, mismatches, and methods: Multiple-view workflows for energy portfolio analysis. *IEEE transactions on visualization and computer graphics* 22, 1 (2015), 449–458.
[29] Volha Bryl, Paolo Giorgini, and John Mylopoulos. 2009. Designing sociotechnical systems: from stakeholder goals to social networks. *Requirements Engineering* 14, 1 (2009), 47–70.
[30] Richard Buchanan. 1992. Wicked problems in design thinking. *Design issues* 8, 2 (1992), 5–21.
[31] Philip Burnard. 1991. A method of analysing interview transcripts in qualitative research. *Nurse education today* 11, 6 (1991), 461–466.
[32] Alberto Cairo. 2019. *How Charts Lie: Getting Smarter about Visual Information*. WW Norton & Company.
[33] Giuseppe Carenini and Lucas Rizoli. 2009. A Multimedia Interface for Facilitating Comparisons of Opinions. In *Proceedings of the 14th International Conference on Intelligent User Interfaces* (Sanibel Island, Florida, USA) *(IUI '09)*. ACM, New York, NY, USA, 325–334.
[34] Chaomei Chen, Fidelia Ibekwe-SanJuan, Eric SanJuan, and Chris Weaver. 2006. Visual analysis of conflicting opinions. In *2006 IEEE Symposium On Visual Analytics Science And Technology*. IEEE, 59–66.
[35] Jaegul Choo and Shixia Liu. 2018. Visual analytics for explainable deep learning. *IEEE computer graphics and applications* 38, 4 (2018), 84–92.
[36] Jason Chuang, Daniel Ramage, Christopher Manning, and Jeffrey Heer. 2012. Interpretation and trust: Designing model-driven visualizations for text analysis. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 443–452.
[37] Andy Cockburn, Amy Karlson, and Benjamin B Bederson. 2009. A review of overview+ detail, zooming, and focus+ context interfaces. *ACM Computing Surveys (CSUR)* 41, 1 (2009), 2.
[38] Gregorio Convertino, Adam Westerski, Anna De Liddo, and Paloma D'iaz. 2015. Large-Scale Ideation & Deliberation: Tools and Studies in Organizations. *Journal*

Social Media for Organizations 2, 1 (2015), 1.
[39] Eric Corbett and Christopher A. Le Dantec. 2018. Exploring Trust in Digital Civics. In *Proceedings of the 2018 Designing Interactive Systems Conference* (Hong Kong, China) *(DIS '18)*. Association for Computing Machinery, New York, NY, USA, 9–20.
[40] Eric Corbett and Christopher A Le Dantec. 2018. Going the Distance: Trust Work for Citizen Participation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 312.
[41] Eric Corbett and Christopher A Le Dantec. 2018. The Problem of Community Engagement: Disentangling the Practices of Municipal Government. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 574.
[42] Michael Correll. 2019. Ethical dimensions of visualization research. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–13.
[43] Lorrie F Cranor. 2008. A framework for reasoning about the human in the loop. (2008).
[44] Benedetto De Martino, Dharshan Kumaran, Ben Seymour, and Raymond J Dolan. 2006. Frames, biases, and rational decision-making in the human brain. *Science* 313, 5787 (2006), 684–687.
[45] Daantje Derks, Arjan ER Bos, and Jasper Von Grumbkow. 2007. Emoticons and social interaction on the Internet: the importance of social context. *Computers in human behavior* 23, 1 (2007), 842–849.
[46] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *CoRR* abs/1810.04805 (2018). arXiv:1810.04805 http://arxiv.org/abs/1810.04805
[47] Sidney D'mello and Arthur Graesser. 2007. Mind and body: Dialogue and posture for affect detection in learning environments. *Frontiers in Artificial Intelligence and Applications* 158 (2007), 161.
[48] Cynthia Dwork. 2008. Differential privacy: A survey of results. In *International conference on theory and applications of models of computation*. Springer, 1–19.
[49] Mennatallah El-Assady, Valentin Gold, Carmela Acevedo, Christopher Collins, and Daniel Keim. 2016. ConToVi: Multi-party conversation exploration using topic-space views. In *Computer Graphics Forum*, Vol. 35. Wiley Online Library, 431–440.
[50] Siamak Faridani, Ephrat Bitton, Kimiko Ryokai, and Ken Goldberg. 2010. Opinion Space: A Scalable Tool for Browsing Online Comments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA) *(CHI '10)*. ACM, New York, NY, USA, 1175–1184.
[51] Peter H Feindt and Angela Oels. 2005. Does discourse matter? Discourse analysis in environmental policy making. *Journal of Environmental Policy & Planning* 7, 3 (2005), 161–173.
[52] Shai Fine, Yoram Singer, and Naftali Tishby. 1998. The hierarchical hidden Markov model: Analysis and applications. *Machine learning* 32, 1 (1998), 41–62.
[53] Shawn Graham, Scott Weingart, and Ian Milligan. 2012. *Getting started with topic modeling and MALLET*. Technical Report. The Editorial Board of the Programming Historian.
[54] Michelle L. Gregory, Nancy Chinchor, Paul Whitney, Richard Carter, Elizabeth Hetzler, and Alan Turner. 2006. User-directed Sentiment Analysis: Visualizing the Affective Content of Documents. In *Proceedings of the Workshop on Sentiment and Subjectivity in Text* (Sydney, Australia) *(SST '06)*. Association for Computational Linguistics, Stroudsburg, PA, USA, 23–30.
[55] Miriam Greis, Jessica Hullman, Michael Correll, Matthew Kay, and Orit Shaer. 2017. Designing for Uncertainty in HCI: When Does Uncertainty Help?. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. 593–600.
[56] Emitza Guzman, Padma Bhuvanagiri, and Bernd Bruegge. 2014. Fave: Visualizing user feedback for software evolution. In *2014 Second IEEE Working Conference on Software Visualization*. IEEE, 167–171.
[57] Kenneth L. Hacker and Jan Van Dijk. 2001. *Digital Democracy: Issues of Theory and Practice*. Sage Publications, Inc., USA.
[58] Mike Harding, Bran Knowles, Nigel Davies, and Mark Rouncefield. 2015. HCI, civic engagement & trust. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 2833–2842.
[59] Brian W. Head. 2007. Community Engagement: Participation on Whose Terms? *Australian Journal of Political Science* 42, 3 (2007), 441–454.
[60] Ren Hongde, Wang Shuo, and Li Hui. 2014. Differential privacy data Aggregation Optimizing Method and application to data visualization. In *2014 IEEE Workshop on Electronics, Computer and Applications*. IEEE, 54–58.
[61] Enamul Hoque and Giuseppe Carenini. 2015. Convisit: Interactive topic modeling for exploring asynchronous online conversations. In *Proceedings of the 20th International Conference on Intelligent User Interfaces*. 169–180.
[62] Mengdie Hu, Krist Wongsuphasawat, and John Stasko. 2016. Visualizing social media content with sententree. *IEEE transactions on visualization and computer graphics* 23, 1 (2016), 621–630.
[63] Mengdie Hu, Huahai Yang, Michelle X. Zhou, Liang Gou, Yunyao Li, and Eben Haber. 2013. OpinionBlocks: A Crowd-Powered, Self-improving Interactive Visual Analytic System for Understanding Opinion Text. In *Human-Computer Interaction – INTERACT 2013*, Paula Kotzé, Gary Marsden, Gitte Lindgaard, Janet Wesson, and Marco Winckler (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 116–134.
[64] Shafiq Joty, Giuseppe Carenini, and Raymond T Ng. 2013. Topic Segmentation and Labeling in Asynchronous Conversations. *Journal of Artificial Intelligence Research* 47 (2013), 521–573.
[65] Yoon Kim. 2014. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882* (2014).
[66] Mark Klein. 2012. Enabling large-scale deliberation using attention-mediation metrics. *Computer Supported Cooperative Work (CSCW)* 21, 4-5 (2012), 449–473.
[67] Travis Kriplean, Jonathan Morgan, Deen Freelon, Alan Borning, and Lance Bennett. 2012. Supporting Reflective Public Thought with Considerit. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work* (Seattle, Washington, USA) *(CSCW '12)*. ACM, New York, NY, USA, 265–274.
[68] Klaus Krippendorff. 2011. Computing Krippendorff's alpha-reliability. (2011).
[69] Kostiantyn Kucher and Andreas Kerren. 2015. Text visualization techniques: Taxonomy, visual survey, and community insights. In *2015 IEEE Pacific Visualization Symposium (PacificVis)*. IEEE, 117–121.
[70] Jennifer S. Lerner, Ye Li, Piercarlo Valdesolo, and Karim S. Kassam. 2015. Emotion and Decision Making. *Annual Review of Psychology* 66, 1 (2015), 799–823. PMID: 25251484.
[71] Narges Mahyar, Kelly J. Burke, Jialiang Ernest Xiang, Siyi Cathy Meng, Kellogg S Booth, Cynthia L Girling, and Ronald W Kellett. 2016. UD Co-Spaces: A Table-Centred Multi-Display Environment for Public Engagement in Urban Design Charrettes. In *Proceedings of the 2016 ACM on Interactive Surfaces and Spaces*. ACM, 109–118.
[72] Narges Mahyar, Michael R. James, Michelle M. Ng, Reginald A. Wu, and Steven P. Dow. 2018. CommunityCrit: Inviting the Public to Improve and Evaluate Urban Design Ideas Through Micro-Activities. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI '18)*. ACM, New York, NY, USA, Article 195, 14 pages.
[73] Narges Mahyar, Diana V Nguyen, Maggie Chan, Jiayi Zheng, and Steven P Dow. 2019. The Civic Data Deluge: Understanding the Challenges of Analyzing Large-Scale Community Input. In *Proceedings of the 2019 on Designing Interactive Systems Conference*. ACM, 1171–1181.
[74] Adam Marcus, Michael S Bernstein, Osama Badar, David R Karger, Samuel Madden, and Robert C Miller. 2011. Twitinfo: aggregating and visualizing microblogs for event exploration. In *Proc. CHI*. ACM, 227–236.
[75] Fahd Saud Nawwab, Paul E Dunne, and Trevor JM Bench-Capon. 2010. Exploring the Role of Emotions in Rational Decision Making.. In *COMMA*. 367–378.
[76] Matti Nelimarkka. 2019. A Review of Research on Participation in Democratic Decision-Making Presented at SIGCHI Conferences. Toward an Improved Trading Zone Between Political Science and HCI. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–29.
[77] Mary J Newhart and Joshua D Brooks. 2017. Barriers to Participatory eRulemaking Platform Adoption: Lessons Learned from RegulationRoom. (2017).
[78] Donald A Norman and Pieter Jan Stappers. 2016. DesignX: Complex Sociotechnical Systems. *She Ji: The Journal of Design, Economics, and Innovation* 1, 2 (2016), 83–106.
[79] Daniela Oelke, Ming Hao, Christian Rohrdantz, Daniel A Keim, Umeshwar Dayal, Lars-Erik Haug, and Halldór Janetzko. 2009. Visual opinion analysis of customer feedback data. In *2009 IEEE symposium on visual analytics science and technology*. IEEE, 187–194.
[80] Patrick Olivier and Peter Wright. 2015. Digital Civics: Taking a Local Turn. *interactions* 22, 4 (June 2015), 61–63.
[81] Huamin Qu and Qing Chen. 2015. Visual analytics for MOOC data. *IEEE computer graphics and applications* 35, 6 (2015), 69–75.
[82] Wojciech Samek, Grégoire Montavon, Andrea Vedaldi, Lars Kai Hansen, and Klaus-Robert Müller. 2019. *Explainable AI: interpreting, explaining and visualizing deep learning*. Vol. 11700. Springer Nature.
[83] Ali Sarvghad and Melanie Tory. 2015. Exploiting Analysis History to Support Collaborative Data Analysis. In *Proceedings of the 41st Graphics Interface Conference* (Halifax, Nova Scotia, Canada) *(GI '15)*. Canadian Information Processing Society, Toronto, Ont., Canada, Canada, 123–130.
[84] Reijo Savolainen. 2015. Expressing emotions in information sharing: a study of online discussion about immigration. (2015).
[85] Farhana Shahid, Shahinul Hoque Ony, Takrim Rahman Albi, Sriram Chellappan, Aditya Vashistha, and ABM Alim Al Islam. 2020. Learning from Tweets: Opportunities and Challenges to Inform Policy Making During Dengue Epidemic. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW1 (2020), 1–27.
[86] Ben Shneiderman. 1996. The eyes have it: A task by data type taxonomy for information visualizations. In *IN IEEE SYMPOSIUM ON VISUAL LANGUAGES*. 336–343.
[87] Ben Shneiderman. 2016. *The new ABCs of research: Achieving breakthrough collaborations*. Oxford University Press.
[88] Julie Simon, Theo Bass, Victoria Boelman, and Geoff Mulgan. 2017. Digital democracy: the tools transforming political engagement. *NESTA, UK, England and Wales* 1144091 (2017).

[89] Nicole Sultanum, Michael Brudno, Daniel Wigdor, and Fanny Chevalier. 2018. More text please! understanding and supporting the use of visualization for clinical text overview. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 422.

[90] James J. Thomas and Kristin A. Cook. 2006. A Visual Analytics Agenda. *IEEE Comput. Graph. Appl.* 26, 1 (Jan. 2006), 10–13.

[91] Nina Valkanova, Robert Walter, Andrew Vande Moere, and Jörg Müller. 2014. MyPosition: Sparking Civic Discourse by a Public Interactive Poll Visualization. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing* (Baltimore, Maryland, USA) *(CSCW '14)*. ACM, New York, NY, USA, 1323–1332.

[92] Vasillis Vlachokyriakos, Clara Crivellaro, Christopher A. Le Dantec, Eric Gordon, Pete Wright, and Patrick Olivier. 2016. Digital Civics: Citizen Empowerment With and Through Technology. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI EA '16)*. ACM, New York, NY, USA, 1096–1099.

[93] Franz Wanner, Christian Rohrdantz, Florian Mansmann, Daniela Oelke, and Daniel A. Keim. 2009. Visual Sentiment Analysis of RSS News Feeds Featuring the US Presidential Election in 2008.

[94] Karl E Weick. 1995. *Sensemaking in organizations*. Vol. 3. Sage.

[95] Yingcai Wu, Shixia Liu, Kai Yan, Mengchen Liu, and Fangzhao Wu. 2014. Opinionflow: Visual analysis of opinion diffusion on social media. *IEEE transactions on visualization and computer graphics* 20, 12 (2014), 1763–1772.

[96] Yingcai Wu, Furu Wei, Shixia Liu, Norman Au, Weiwei Cui, Hong Zhou, and Huamin Qu. 2010. OpinionSeer: Interactive Visualization of Hotel Customer Feedback. *IEEE Transactions on Visualization and Computer Graphics* 16, 6 (Nov. 2010), 1109–1118.

[97] Baoxun Xu, Xiufeng Guo, Yunming Ye, and Jiefeng Cheng. 2012. An Improved Random Forest Classifier for Text Categorization. *JCP* 7, 12 (2012), 2913–2920.

[98] Koji Yatani, Michael Novati, Andrew Trusty, and Khai N. Truong. 2011. Review Spotlight: A User Interface for Summarizing User-generated Reviews Using Adjective-noun Word Pairs. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vancouver, BC, Canada) *(CHI '11)*. ACM, New York, NY, USA, 1541–1550.

[99] Amy X Zhang, Lea Verou, and David Karger. 2017. Wikum: Bridging discussion forums and wikis using recursive summarization. (2017).

[100] Wen Zhang, Taketoshi Yoshida, and Xijin Tang. 2008. Text classification based on multi-word with support vector machine. *Knowledge-Based Systems* 21, 8 (2008), 879–886.

[101] Jian Zhao, Nan Cao, Zhen Wen, Yale Song, Yu-Ru Lin, and Christopher Collins. 2014. # FluxFlow: Visual analysis of anomalous information spreading on social media. *IEEE transactions on visualization and computer graphics* 20, 12 (2014), 1773–1782.

[102] Roshanak Zilouchian Moghaddam, Zane Nicholson, and Brian P Bailey. 2015. Procid: Bridging consensus building theory with the practice of distributed design discussions. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. ACM, 686–699.