CommunityClick: Capturing and Reporting Community Feedback from Town Halls to Improve Inclusivity

MAHMOOD JASIM, University of Massachusetts Amherst, USA
POOYA KHALOO, University of Massachusetts Amherst, USA
SOMIN WADHWA, University of Massachusetts Amherst, USA
AMY X. ZHANG, University of Washington, USA
ALI SARVGHAD, University of Massachusetts Amherst, USA
NARGES MAHYAR, University of Massachusetts Amherst, USA

Local governments still depend on traditional town halls for community consultation, despite problems such as a lack of inclusive participation for attendees and difficulty for civic organizers to capture attendees' feedback in reports. Building on a formative study with 66 town hall attendees and 20 organizers, we designed and developed CommunityClick, a communitysourcing system that captures attendees' feedback in an inclusive manner and enables organizers to author more comprehensive reports. During the meeting, in addition to recording meeting audio to capture vocal attendees' feedback, we modify iClickers to give voice to reticent attendees by allowing them to provide real-time feedback beyond a binary signal. This information then automatically feeds into a meeting transcript augmented with attendees' feedback and organizers' tags. The augmented transcript along with a feedback-weighted summary of the transcript generated from text analysis methods is incorporated into an interactive authoring tool for organizers to write reports. From a field experiment at a town hall meeting, we demonstrate how CommunityClick can improve inclusivity by providing multiple avenues for attendees to share opinions. Additionally, interviews with eight expert organizers demonstrate CommunityClick's utility in creating more comprehensive and accurate reports to inform critical civic decision-making. We discuss the possibility of integrating CommunityClick with town hall meetings in the future as well as expanding to other domains.

CCS Concepts: • **Human-centered computing** → **Human computer interaction (HCI)**.

Additional Key Words and Phrases: Town hall; automatic transcription; iClicker; community feedback

Authors' addresses: Mahmood Jasim, University of Massachusetts Amherst, Amherst, MA, USA, mjasim@cs.umass.edu; Pooya Khaloo, University of Massachusetts Amherst, Amherst, MA, USA, pkhaloo@cs.umass.edu; Somin Wadhwa, University of Massachusetts Amherst, Amherst, MA, USA, sominwadhwa@cs.umass.edu; Amy X. Zhang, University of Washington, Seattle, Washington, USA, axz@cs.uw.edu; Ali Sarvghad, University of Massachusetts Amherst, Amherst, MA, USA, asarv@cs.umass.edu; Narges Mahyar, University of Massachusetts Amherst, Amherst, MA, USA, nmahyar@cs.umass.edu.

1 INTRODUCTION

Traditional community consultation methods, such as town halls, public forums, and workshops are the *modus operandi* for public engagement [52, 94]. For fair and impartial civic decision-making, the inclusivity of community members' feedback is paramount [60, 94, 126]. However, traditional methods rarely provide opportunities for inclusive public participation [30, 87, 95]. For instance, reticent meeting attendees struggle to speak up and articulate their viewpoints due to fear of confronting outspoken and dominant individuals [27, 127]. This lack of inclusivity in traditional face-to-face meetings results in an uneven representation of community members and often fails to capture broader perspectives from attendees [70]. As a result, these methods often fall short in achieving the desired exchange of perspectives between government officials and the community [76, 118, 119]. Furthermore, meeting organizers grapple with simultaneously facilitating often contentious discussions and taking meeting notes to capture attendees' broader perspectives [70, 95]. These bottlenecks further obstruct inclusivity and may lead to biased decisions that can significantly impact people's lives [94, 95]. Advancements in computer-mediated technology can address this predicament by creating a communication channel between these entities.

Bryan [29] and Gastil [56] investigated the state of town halls[1] and demonstrated a steady decline in civic participation due to the growing disconnect between local government and the community. To reengage disconnected, reticent, or disenfranchised community members, researchers in HCI and digital civics[2] have offered novel strategies and technological interventions to increase engagement [60, 62, 94, 107, 130]. Researchers in this field have proposed several online technologies that made wider participation possible for community members (e.g., [4, 5, 7, 93]). Despite the introduction of such online platforms, government officials and decision-makers predominantly favor traditional face-to-face meetings to create relationships, foster discourse, and conduct follow-up conversations with community members [16, 40, 94] to understand their views and aspirations [52, 69, 131]. However, employing technology to capture attendees' feedback—in particular, silent attendees' feedback—in face-to-face meetings remains largely unexplored. Commonly, feedback in meetings is gathered using voting or polling attendees [25, 90, 103] or taking notes during the meeting [29, 91, 96]. However, voting often restricts attendees to only agreeing or disagreeing, which often does more harm to the richness of the captured feedback from attendees rather than promoting inclusivity [22, 94]. To help alleviate this problem, prior work mostly focused on automatic speech recognition [3, 128] and interactive annotations [19, 74] to help organizers take notes for creating reports. However, these methods rarely preserve the discussion context or improve the inclusivity of attendees' feedback.

To better understand the needs of attendees and organizers in community consultations and explore how technology can help to address these issues, we conducted a formative study by attending three town halls in a college town in the United States. We surveyed a total of 66 attendees to inquire about their ability to voice their opinions during town halls. 17% of the attendees (11 responses) expressed that despite being physically present, they could not voice their opinions. They attributed this to factors such as intimidation from other participants, lack of confidence, and fear of interruption. Moreover, we surveyed 20 organizers to identify what could help them to make better use of the town halls to ensure that the public feedback is prioritized in the reports they authored. We found that the organizers often relied on their memories or the notes taken during

---

[1]In this work, we use the term *Town Hall* to refer to various community consultation approaches where community members convene to meet and discuss civic issues with government officials or decision-makers. Town Halls can take many forms, but the main goal is to establish a communication channel between the officials and the community, to inform the community about civic issues, and often to receive community's feedback [29].

[2]Digital Civics is an emerging interdisciplinary area that explores novel ways to utilize technology for promoting broader public engagement and participation in the design and delivery of civic services [62, 107, 130].

the meeting to generate reports. However, they struggled to accurately capture or remember the attendees' feedback and important details after the meeting. In effect, these incomplete memories and meeting notes could potentially result in incomprehensive reports. These findings indicated a requirement for better capturing attendees' feedback and preserving meeting discussion details to support organizers in authoring more comprehensive reports.

Based on our formative study, we designed and developed CommunityClick, a system for town halls that captures more inclusive feedback from attendees during the meeting and enables organizers to author more comprehensive meeting reports. We modified iClickers [9] to allow attendees to silently and anonymously provide feedback and enable organizers to tag the meeting discussion at any time during the meeting. We chose to use iClickers due to their familiarity and ease of use [65, 102, 134] as an audience response system (ARS) [106]. Furthermore, we augmented the automatically-generated meeting transcript by synchronizing it with attendees' feedback and organizers' tags. We also provided an interactive interface where organizers can explore, analyze, and utilize the augmented meeting transcript in a data-driven way and a novel feedback-weighted summary of the meeting discussion to author meeting reports at their convenience. To evaluate the efficacy of our approach, we conducted a field experiment in the wild, followed by eight semi-structured interviews with experienced organizers. Our results demonstrate the efficacy of CommunityClick to give voice to reticent participants to increase their involvement in town halls, capture attendees' feedback, and enable organizers to compile more inclusive, comprehensive, and accurate meeting reports in a way that lends credibility to the report creation process.

Our key contributions in this work are as follows: 1) using a communitysourcing technology to enable attendees to share their feedback at any time during the meeting by modifying iClickers as a real-time response mechanism, 2) augmenting meeting transcripts by combining organizers' tags and attendees' feedback to preserve discussion context and feedback from a broader range of attendees, 3) applying a novel feedback-weighted summarization method to generate meeting summaries that prioritize community's feedback, 4) developing an interface for organizers to explore and utilize the augmented meeting transcript to author more comprehensive reports, and 5) insights from a field experiment in the wild that demonstrates how technology can be effective in capturing attendees' voices, authoring more inclusive reports, and future directions to enhance our approach.

2 BACKGROUND

In this section, we describe the current challenges that inhibit inclusivity in town hall meetings. We also discuss existing technologies designed to promote engagement in various meeting scenarios and to help analyze and utilize meeting data.

2.1 Current Challenges in Town Halls

Prior investigations by Bryan [29] and Gastil [56] showed a steady decline in civic participation in town halls due to the growing disconnect between local government and community members and the decline in social capital [43, 111, 113]. Despite the introduction of online methods to increase public engagement in the last decade [4, 5, 7, 37, 81, 93], government officials continue to prefer face-to-face meetings to engage the community in the decision-making process [32, 52, 94]. They believe face-to-face meetings facilitate two-way communications between decision-makers and community members that can foster discourse and help them understand the views and aspirations of the community members [32, 52, 69, 94, 131]. However, constraints such as fixed physical locations for co-located meetings and scarcity of time and resources, limit the efficacy of face-to-face processes and inhibit the ability of officials to make proper use of town halls [18, 70, 87, 95, 120]. Such constraints might repel or alienate citizens for whom traveling or dedicating time

for town halls may not be a viable option from the perspective of physical, economical, or intrinsic interest [38, 71, 100, 120]. While predictors such as education, income, and intrinsic interest help gauge civic engagement [45], social dynamics, such as shyness and tendency to avoid confrontation with dominant personalities can also hinder opinion sharing in town halls by favoring privileged individuals who are comfortable or trained to take part in contentious public discussions [27, 127]. Specifically, people with training in analytical and rhetorical reasoning often find themselves at an advantage when discussing complex and critical civic issues [119]. As a result, town halls inadvertently cater to a small number of privileged individuals, and silent participants often become disengaged despite physically attending the meetings [61]. Due to the lack of inclusivity, the outcome of such meetings often tends to feel unjust and opaque for the general public [39, 54].

2.2 Technological Interventions to Increase Engagements in Town Halls

To increase broader civic participation, researchers in HCI have proposed both online [4, 5, 7, 81, 93] and face-to-face [21, 80, 91, 125] technological interventions that use the communitysourcing[3] approach. For instance, to increase engagement in town halls, some researchers have experimented with audience response systems (ARS) [25, 77, 80, 103]. Murphy used such systems to promote democracy and community partnerships [103]. Similarly, Boulianne et al. deployed clicker devices in contentious public discussions about climate change to gauge public opinions [25]. Bergstrom et al. used a single button device where the attendees anonymously voted (agree/disagree) on issues during the meeting. They showed that back-channel voting helped underrepresented users get more involved in the meeting [22]. The America*Speaks*' public engagement platform, 21 Century Town Meeting®, also used audience response systems to collect feedback and perform straw polls and votes during town halls [90]. However, in these works, the audience response systems were used either for binary voting or polling [22, 90], or to receive feedback on specific questions that expected attendees' feedback on a Likert scale-like spectrum [88]. These restrictions limit when and how meeting attendees can share their feedback. Audience response systems have seen widespread success as a lightweight tool to engage participants, promote discussions, and invoke critical thinking in the education domain [14, 65, 79, 101, 102, 134]. As such, these devices have the potential to provide a communication channel for silent participants to express their opinions without the obligation to verbalize and risk confrontation. We build upon their success in the education domain by appropriating iClickers for the civic domain. We modify and utilize iClickers for town halls to create a mechanism that supports silent and real-time community feedback.

HCI researchers have proposed research solutions to increase participation in face-to-face meetings such as design charrettes, or group-meetings in general, by using various tools and methods [21, 22, 67, 92, 122]. Some researchers used interactive tabletop and large screen surfaces to engage attendees [67, 92, 122]. For example, UD Co-Spaces [92] used a tabletop-centered multi-display environment for engaging the public in complex urban design processes. Memtable [67] used a tabletop display to increase attendees' engagement by integrating annotations using a multi-user text editor on top of multimedia artifacts. IdeaWall visualized thematically grouped discussion contents in real-time on a large screen during the meeting to encourage attendees to discuss such topics [122]. However, large displays along with real-time visualizations might distract attendees from concentrating on meeting discussions [15], which might lead to less contribution from the participants irrespective of how vocal they are [61]. Furthermore, these innovative approaches might be overwhelming for meeting attendees, especially in the civic domain, due to the heterogeneity

---

[3]Communitysourcing leverages the specific knowledge of a targeted community. It follows the crowdsourcing model, but takes a more direct approach in harnessing collective ideas, input, and feedback from targeted community members to find solutions to complex problems such as civic decision-making. [28, 64]

of participants with a wide spectrum of expertise and familiarity with technology [13, 42, 127]. It might also be impractical and financially infeasible to use expensive tabletops and large interactive displays for town halls organized in a majority of cities.

2.3 Technologies to Analyze and Utilize Meeting Data

Prior works showed that decision-makers often relied on meeting reports generated by organizers to inform public policies that had potential long-term impacts on the community [91]. Commonly, organizers generate these reports based on the outcome of voting or polling attendees [25, 90, 103], and taking notes during the meeting [29, 91, 96]. They often use manual note-taking techniques such as pen-and-paper or text editors [47]. However, simultaneously taking notes from the meetings to capture attendees' feedback and facilitating the meeting to guide and encourage discussions is overwhelming and often lead to losing critical information and ideas [110, 117]. Sometimes the meetings are audio-recorded and transcribed for reviewing before writing the report. However, such reviewing processes require significant manpower and labor [90, 91]. Furthermore, the audio recording themselves do not capture the feedback of reticent participants who did not speak up.

In general meeting scenarios, researchers proposed improvements to manual note-taking for capturing meeting discussions [2, 35, 46, 75]. For example, LiteMinutes used a web interface that allowed creating and distributing meeting contents [35]. MeetingKing provided meeting agenda templates to help organizers create reports [2]. Similarly, LiveNotes facilitated note-taking using a shared whiteboard where users could write notes using a digital pen or a keyboard [20, 75]. Another group of researchers experimented with various tools and techniques to provide support for facilitating online synchronous group discussions. For example, SolutionChat provides moderation support to group discussion facilitators by visualizing group discussion stages and featured opinions from participants and suggesting appropriate moderator responses [86]. Bot in the Bunch takes a different approach and proposes a chatbot agent that aims to enhance goal-oriented online group discussions by managing time, encouraging participants to contribute to the discussion, and summarizing the discussion [82]. Similarly, Tilda synthesizes online chat conversation using structured summaries and enable annotation of chat conversations to improve recall [135].

Closer to our approach, some researchers proposed different techniques for automatically capturing meeting discussions without significant manual intervention [1, 3, 19, 128]. For example, CALO and Voicea are automated systems for annotating, transcribing, and analyzing multiparty meetings [3, 128]. These systems use automatic speech recognition [112] to transcribe the meeting audio and extract topics, question-answer pairs, and action items using natural language processing [66]. However, automatic approaches are often error-prone and susceptible to miscategorization of important discussion topics [36, 97]. To address this problem, some researchers suggested incorporating human intervention to control and compile reports without completely depending on automatically generated results [1, 19, 74]. SmartNotes [19] and commercial applications such as ICompassTech [1] use this approach where the users can add and revise notes and topics. Although these methods used a combination of automatic and interactive techniques, they enabled the utilization of meeting discussions without capturing sufficient circumstantial input. The automatically generated transcript might record discussions but it may not contain feedback shared by all attendees, especially silent ones. Furthermore, these tools are designed for small-scale group meetings and may not be practical in town halls in the civic domain where attendee numbers and requirements can vary significantly [29, 91]. The heterogeneity of attendees in town halls [13, 42], the lack of inclusivity in sharing their opinions [76, 119], the deficient methods to record meeting discussions [90], and limited financial and design pragmatism in existing technologies [92, 122] necessitate a closer investigation and innovation in designing technology to address both meeting attendees' and organizers' challenges regarding inclusivity in town halls.

3 FORMATIVE STUDY: SCHOOL BUILDING TOWN HALLS

To inform our work, we wanted to understand the perspectives and needs of organizers and attendees who participate in community consultations. To this end, we attended several town halls, including multiple public engagement sessions in a college town in the United States (U.S.). The agenda for these town halls included discussions regarding new public school buildings where the town authorities wanted the community's feedback on two new proposals that involved millions of dollars worth of renovation or reconstruction. The town authorities made careful considerations to ensure that the public had access to all the information about the proposals by arranging six community-wide engagement sessions organized by professional facilitators. We joined three out of six of these sessions in a span of three months. We decided to investigate this particular case due to the unique characteristics of the town, where there is relatively high citizen engagement on discussions around education and public schools. We also wanted to investigate how people discuss potentially contentious proposals in the town halls that pitted two contrasting ideas for future public school buildings.

3.1 Participants and Procedures

We approached both the meeting attendees, who were community members, and the meeting organizers, who included facilitators and town officials. The town halls began with the organizers presenting their proposal to attendees. Afterward, attendees engaged in discussions and shared their ideas, opinions, and concerns about the proposals. The facilitators were responsible for guiding the discussion and collecting the attendees' feedback. They played a neutral role in the discussion and did not influence the attendees' opinions in any way.

To understand the attendees' perspectives, we conducted surveys after each town hall. The attendees were all residents of the town and attended the meeting of their own volition. We surveyed a total of 66 attendees. We refer to the attendees from our formative study as **FA** and organizers from our formative study as **FO**. In the survey, we asked them open-ended questions regarding their motivation to join the town hall meetings, their experiences in these meetings, and their thoughts around what makes these face-to-face meetings successful. We also asked about their ability to voice their opinions in town halls and about their familiarity with audience response technologies. Furthermore, to gain an understanding of what type of semantic tags they wanted to provide during a town hall, we compiled a list of tags based on prior work on meeting discussions or public engagements that focused on characterizing effective participation during synchronous communication between organizers and meeting attendees in both online and face-to-face settings [68, 72, 104, 105, 135, 136]. The list of tags is presented in Fig. 1(B). We provided this list to the attendees and asked them to rate which of these tags would help them to express their thoughts in town halls. They rated each tag on a 5-point Likert scale [88].

During these town halls, we made contact with the organizers and explained our project goals. From these initial contacts, we used the snowball method [59] to reach other organizers working across the U.S. to learn more about their practices regarding town hall meeting organization, facilitation, and report creation. We conducted surveys with a total of 20 organizers with an average experience of 10.5 years (min 1 year, max 35 years) in conducting town halls across the U.S. Our survey participants consisted of town administrators, town clerks, senior planners, directors, professional facilitators, and chairs of different town committees. We asked them open-ended questions around their meeting data recording practices including what data do they prioritize when recording, the challenges they face while organizing the meeting simultaneously and recording meeting notes, and their post-meeting data analysis and meeting report generation processes. We also asked about their choice of tools and technologies for recording meeting notes and for
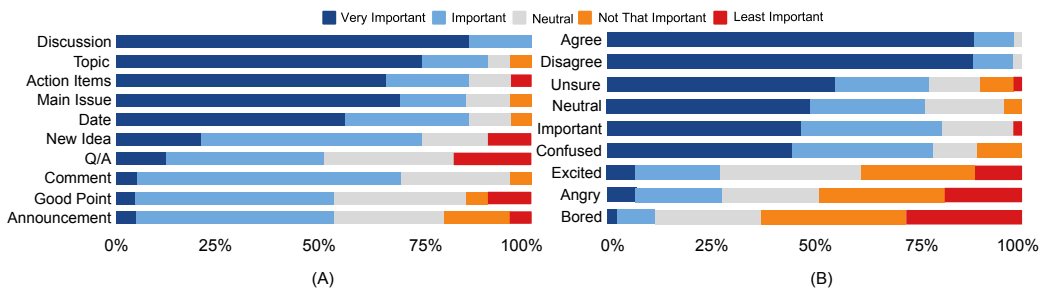
Fig. 1. This figure presents attendees' and organizers' perceived importance of two sets of tags that we compiled based on prior research: (A) shows the ratings of 20 organizers on a list of 10 tags for organizers; (B) shows the ratings of 66 Community members on a list of 9 tags for attendees' feedback.

creating reports, and what they think are the elements that constitute a good report. To identify representative tags that could help organizers to track and categorize meeting discussion, we compiled another list of tags based on prior works [68, 72, 104, 105, 135]. We used a different list from the one we provided to attendees (Fig. 1(A)) because prior works suggest that organizers and attendees need different tags to categorize meeting discussion based on their perspectives. We asked organizers to rate the tags in the order of importance on a five-point Likert scale [88]. The survey questions and list of tags are provided as supplementary materials.

3.2 Findings

Here, we report the findings from our surveys with both attendees and organizers. The 66 attendees we surveyed were highly motivated to attend town halls and the majority of them (64%, 42 responses) considered attending such meetings to provide their feedback on civic issues as their *civic duty*. Most of the attendees (88%, 58 responses) attended two or more town halls every year. Regarding their familiarity with technology, every meeting attendee (100%, 66 responses) mentioned having a computer, smartphone, and internet access, but 61% of them (40 responses) never used an audience response system before. We also found that 17% (11 responses) of meeting attendees felt they were not able to share their feedback during these meetings, and 23% (15 responses) were not satisfied with the way town halls were organized to discuss critical issues. It was surprising for us to find that 17% of people from a homogeneous, relatively wealthy, and educated community in a college town believed that they could not voice their opinions during town halls. Despite their unfamiliarity with audience response systems, the majority of the meeting attendees (87%, 57 responses) mentioned that they were willing to use such devices in town halls to share their feedback.

In response to the question regarding what makes face-to-face town hall meeting successful, a group of attendees mentioned that the success of town hall meetings hinges upon the attendees' ability to openly communicate their opinions and hear others' opinions. One attendee (FA-17) mentioned, "*Town halls need to provide opportunities for all people to be heard, having diversity of voices and opinions, and be present in the discussion.*" Another attendee (FA-36) mentioned, "*Being face-to-face means asking questions, listening to others' questions and ideas, and to be able to see their emotions, nuances, and expressions.*" Several attendees' mentioned facing challenges around sharing their opinions due to being dominated in the discussion. One such attendees' mentioned (FA-23), "*Town halls give us the chance to talk things out, but often this doesn't happen and people get shut down.*" Another attendee (FA-11) mentioned, "*You need to have ground rules so that we stay on track and no one dominates the discussion.*" Some attendees considered that facilitators play an important role in the success of town halls and skilled and well-equipped facilitators can

make a difference. One attendee (FA-57) emphasized the importance of skilled facilitator saying, "*Professional facilitators understand the context of our community, and the history of tension regarding the issues. They move it along and listen with open minds.*" Another attendee (FA-47) mentioned, "*Skilled facilitators make sure all voices are heard, organize the discussion around goals and keep everyone focused on tasks.*"

When rating the tags for sharing opinions during meetings (Fig. 1 (A)), the majority (90%) of the attendees considered *Agree* and *Disagree* to be the most important tags. 75% or more attendees thought *Unsure*, *Important*, *Confused*, and *Neutral* to be important.

The majority of the organizers (17 responses) mentioned that often manpower and budget constraints forced them to forego the appointment of designated note-takers and thus, they must shuffle between organizing and note-taking during the meetings. Some organizers also mentioned how context-switching between organizing the meeting and taking notes often led to missing critical evidence or information (8 responses). They employed a variety of methods for recording meeting data depending on their convenience and their operational abilities and experiences to utilize such methods, including pen-and-paper (5 responses), text editors (6 responses), audio or video recorders (3 responses), or a combination of these methods. To generate reports from the notes taken during the meetings, organizers usually used text editors (17 responses) with preferences towards Microsoft Word and Notepad (12 responses).

However, when asked about the time required to compile a report from meeting notes, their responses varied from 15 minutes to a few days. The variation, in part, can be attributed to the amount of notes and the format in which notes were captured. All organizers (20 responses) mentioned that the report generation process involves some form of summarization process of meeting records to retain the most important discussion components. One organizer (FO-7) explained, "*If I'm responsible for the report, I listen to the recording, review the notes, and translate them into coherent accounts of meeting attendees, topics, discussion, and decisions/outcomes.*". Another organizer (FO-1) mentioned, "*I start with the agenda, review the audio, then edit and summarize the meeting notes into report content.*"

Organizers also described high-level properties of what would constitute a *good report*. One organizer (FO-1) mentioned that, good meeting reports should be "*accurate, comprehensive, clear, and concise*". Another organizer (FO-4) emphasized that several components constitue a good report including, "*[the] main agenda items, relevant comments, consensus, action plans, and feedback*". One organizer (FO-17) thought meeting reports should be "*balanced, fair, and pertinent*" towards multiple perspectives, however, it is often challenging because they only "*listen to a few, while others stay silent*".

When rating tags that would help them to take better notes to create good reports (Fig. 1 (B)), the meeting organizers unanimously (100%) endorsed the tag *Decision*. The tags *Topic*, *Action Items*, *Main Issue*, and *Date* were preferred by more than 70% organizers. However, all other tags except for *Q&A* were favored by more than 50% of organizers. These survey responses suggest that preferences for tagging meeting discussion varies among organizers, and different sets of tags might be required to be useful for generating reports from different meetings with diverse agendas.

3.3 Design Goals

Based on prior work and our formative study, we identified four design goals to guide our system design that address the requirements and challenges of both meeting attendees and organizers. First, we found that some attendees lacked a way to respond to ongoing discussions. Furthermore, many organizers struggled to keep track of the discussion and take notes simultaneously. Thus, we needed to provide a communication channel between attendees and organizers for sharing opinions and capturing feedback (G1). Second, the organizers often refer to several sources of
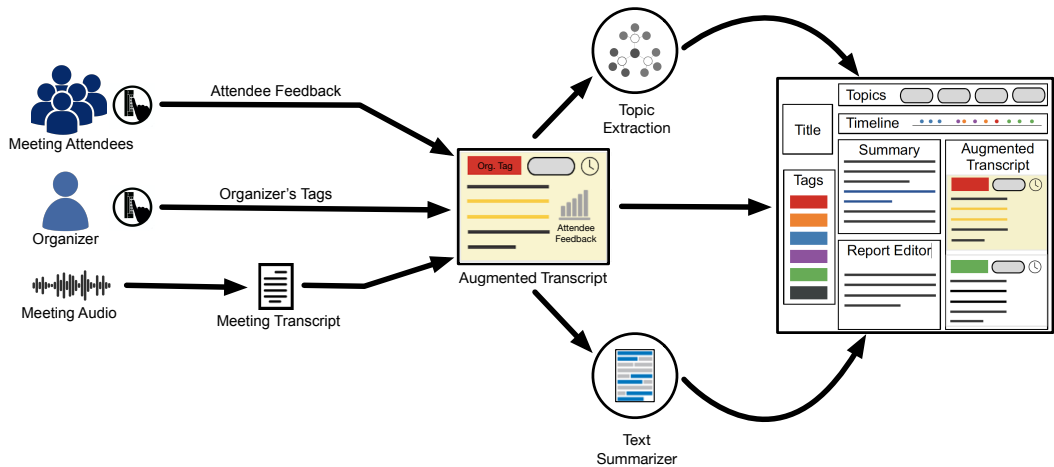
Fig. 2. A snapshot of CommunityClick's workflow. During the meeting, attendees and organizers can use iClickers to share feedback and tag the meeting respectively. The meeting is also audio-recorded for transcription. The audio recordings are transcribed automatically and then augmented with the organizer's tags and attendees' feedback. Furthermore, we generated the feedback-weighted discussion summary and extracted the most relevant topics. The interactive interface enables the exploration and utilization of augmented meeting discussions, which is available online for organizers to examine and author meeting reports.

meeting data including meeting notes and audio/video recording to compile meeting reports. Hence, the meeting discussion audio, organizers' tags, and attendees' feedback should be captured and combined together to provide organizers with a holistic view of the meeting data (G2). Third, the organizers perform some form of summarization to generate meeting reports. However, many organizers also struggled to account for attendees' feedback while summarizing meeting data. This challenge motivated a third goal to introduce summarization techniques that incorporate attendees' feedback to help organizers to get the gist of meeting discussions (G3). Finally, organizers needed to examine the meeting-generated data to capture more inclusive attendees' feedback so that they could write more comprehensive reports. To that end, our final design goal was to provide exploration and report authoring functionalities to help them investigate the meeting data and identify evidence to generate reports that included and reflected attendees' feedback (G4).

4 COMMUNITYCLICK

Guided by the design goals, we designed and developed CommunityClick, a system where we modified iClickers as a real-time response mechanism and augmented meeting transcripts to capture more inclusive feedback. We also introduced a novel feedback-weighted summarization to prioritize attendees' feedback and enabled exploration and utilization of the augmented transcript through an interactive interface for organizers to author meeting reports. Here, we provide a scenario where CommunityClick can be employed (Fig. 2), followed by the system description.

4.1 User Scenario

Michelle is an organizer who has been appointed by the local government officials to organize an important town hall. Given the importance of the meeting, she decides to deploy CommunityClick to focus on facilitating the meeting while using iClickers to capture the community's feedback.

Adam is a community member who is attending the town hall. He cares about the community and wants to share his opinions on the agenda. He prefers to avoid confrontations, especially in town halls, as he is worried about speaking up and running into arguments. In the meeting, he is given an iClicker and instructions for using it to share his opinions using five options. Adam engages in the discussion with other attendees in the meeting but whenever he feels hesitant to speak up, he uses the iClicker to provide feedback.

A week later, Michelle finally gets around to writing the report of the town hall. By now, she has forgotten a significant portion of the meeting discussion. She logs in to CommunityClick and selects the town hall. She uses the timeline and feedback-weighted summary to get an overview of the meeting discussion and jogs her memory by exploring the meeting discussion. She uses her own tags, timeline, and the interactive summary to investigate the augmented meeting transcript that contained attendees' feedback alongside the discussion segments. Finally, she authors the report by importing information into the text editor from the transcript.
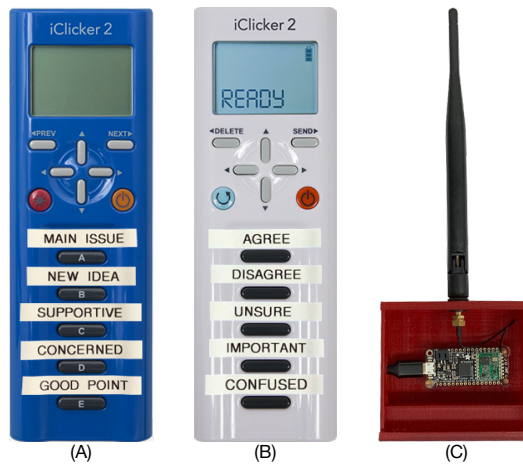

Fig. 3. The apparatus used to capture organizers' tags and attendees' feedback. (A) The iClicker for organizers to tag the meeting. (B) The iClicker for attendees to add their feedback. We used different sets of tags for organizers and attendees based on our formative study. Each iClicker was labeled with the respective set of tags to reduce the cognitive load of mapping options to the iClicker buttons. (C) The iClicker recorder. We used an Adafruit Feather M0 with the 900 MHz RFM69W/RFM69HW transceiver to capture iClicker clicks with timestamps in real-time to synchronize tags and feedback with meeting audio.

4.2 System Description

In the following, we describe how we addressed the design goals by using iClickers, augmenting meeting transcripts, performing text analysis, and developing an interactive interface.

4.2.1 Modifying iClickers to enable attendees and organizers to provide real-time responses (G1). We used iClickers for both organizers and attendees to enable them to respond to meeting discussions any time during the meeting without the need to manually take notes or speak up to share opinions. iClicker is a communication device that uses radio frequency to allow users to anonymously respond using its five buttons (Fig. 3). Despite the widespread usage of smartphones, we chose iClickers as an audience response system due to recent statistics that show 20% of U.S. citizens do not yet have access to smartphones [6]. Furthermore, they are often a major cause of distraction and hindrance to participation in meetings [17, 85]. There are also technical overheads involved including installation

of applications, maintenance, and issues regarding version compatibility based on operating systems which might disengage participants from using the application. In contrast, clicker devices have proven to be successful in town halls due to its familiarity and affordance of anonymity in sharing opinions in town halls and to receive attendees' feedback on specific questions [22, 25, 103]. The anonymous use of iClickers could ensure that silent participants can also share their opinions to the organizer about the ongoing discussion without engaging with a potentially heated debate. Moreover, we modified the iClickers to go beyond previous approaches by allowing meeting organizers and attendees to respond to the ongoing discussion using all five different options instead of only binary agree or disagree. The tag options are customizable and depending on their meeting agendas, the organizers can set up an appropriate list of tags and feedback before the meeting. We used different types of iClickers for organizers and attendees. The organizers used instructor iClickers and attendees used regular ones (Fig. 3). It helped us to effectively separate organizers' and attendees' responses. To reduce the cognitive load of iClicker users to map and remember the options, we labeled each iClicker with organizers' tags and attendees' feedback.

4.2.2 Combining automatically generated transcripts with tags and feedback to capture more inclusive feedback (G2). We recorded and synchronized three different sets of data generated simultaneously in the meeting—the discussion audio, the organizers' tags, and the attendees' feedback via iClickers. We recorded the meeting audio using a regular and commonly used omnidirectional microphone. To remove the noise from the audio recording, we used an open-source freeware named Audacity®. However, capturing organizers' tags and attendees' feedback from iClickers was non-trivial due to the limitations in hardware access and the API provided by the iClicker manufacturer. The original software and hardware in factory settings did not provide timestamped data on each click. As a result, we customized the hardware and API to record organizers' tags and attendees' feedback (Fig. 3). We used an Adafruit Feather M0 with the 900 MHz RFM69W/RFM69HW transceiver to collect an iClicker's clicks and timestamps that were transmitted through radio frequency on the same bandwidth. This allowed us to accurately and precisely capture and synchronize iClicker interactions to match the time of discussion.

To transcribe the meeting audio, we used automatic speech recognition techniques from Assembly AI [8]. We assessed the quality of this method by comparing them to human-generated reference transcripts. We found our approach to be on-par with human-generated transcripts. The results of these analyses are presented in full in the Appendix section. We combined the transcript with the timestamped tags and feedback to transform the recorded meeting audio into timestamped text. Furthermore, we used the organizers' tags to divide the meeting transcript into manageable and consumable segments. Previous work showed that there is a gap (2 seconds on average) between hearing something, registering it, and taking actions upon it, such as clicking a button for annotation [115]. Based on prior work and our early pilot experiments, for each organizer's tag, we created a 30 second time window around the tag (2 seconds before the tag and 28 seconds after the tag). The complete meeting transcript is divided into similar 30-second segments. For each segment, we collected the attendees' feedback provided within that time-window. Consequently, the meeting audio is transformed into timestamped segments, each containing transcribed conversation, organizers' tags, and a set of attendees' feedback (Fig. 4(E)). We also extracted the main discussion points from the transcript segments using Topic Rank [24] (Fig. 4(B)). We chose this topic modeling method to have better multi-word topics which are useful for better topic representation [23].

4.2.3 Applying text summarization method to incorporate attendee's feedback into meeting discussion summaries (G3). To summarize the meeting transcript, we used the graph-based TextRank algorithm [99], which is based on a variation of the PageRank [108] algorithm. TextRank is known to deliver reasonable results in summarizing meeting transcripts [55] in unsupervised settings.

However, it treats all input text the same way without any domain-specific consideration. Our goal was to incorporate attendees' feedback in the summarization process so that the resultant summary was weighted by attendees' feedback. To that end, we added two critical modifications to the original methodology: 1) by incorporating attendees' feedback while computing relative importance of sentences, and 2) replacing the vanilla similarity function used in TextRank with a bag-of-words ranking function called BM25, which is proven to work well in information extraction tasks [116].

Each individual transcript is treated as a set of sentences $(s_1, s_2...s_n)$. Each sentence is considered as an independent node. We used a function to compute the similarity between these sentences to construct edges between them. The higher the similarity between the sentences, the more important the edge between them will be in the graph. The original TextRank algorithm considers the relation between two sentences based on the content (tokens) they share. This relationship is between two sentences S_i, S_j, for every w_k common token, is given by equation 1.

$$sim(S_i, S_j) = \frac{|\{w_k | w_k \in S_i \& w_k \in S_j\}|}{\log(|S_j|) + \log(|S_i|)} \tag{1}$$

We replaced the above similarity function by a BM25 ranking function utility defined by equation 2.

$$sim(S_i, S_j) = \sum_{k=1}^{n} IDF(w_k \in S_i) \frac{f(w_k \in S_i, S_j).(a+1)}{f(w_k \in S_i, S_j) + a.(1 - b + b.\frac{|P|}{\mu_{DL}})} \tag{2}$$

where a, b are function parameters ($a = 1.2, b = 0.75$), $f(w_k, S_j)$ is w_k's term frequency in S_j, IDF is the inverse document frequency, and μ_{DL} is the average length of the sentences in our collection. More importantly, since we timestamped and tagged our transcripts, we knew which instances were potentially more important in terms of garnering attendees' feedback. To incorporate those, we augmented every edge weight $w_{i,j}$ by a factor of ϵ determined experimentally (set to 1.10 if either of the sentences being compared prompted feedback, or 0.90 otherwise). From the constructed graph of sentences, we computed the final relative importance of every vertex (sentence), and selected the top-n sentences corresponding to about 30% of the total length of the transcript which were then presented, in their original order, as the summary of the meeting discussion.

We performed a series of ablation tests to evaluate the robustness of our summarization approach. The results of these tests are presented in full in the Appendix section. For three different meeting transcripts that we had generated using CommunityClick, we quantitatively evaluated the auto-generated summaries against human-annotated reference summaries using the widely used ROUGE [89] metrics. Across different meeting transcripts, we observed similar ROUGE scores, indicating that we produced summaries of consistent quality of summaries across different meetings. Additionally, we evaluated our algorithmic approach on the popular AMI meeting corpus [33]. Our results were found to be comparable to the current state-of-the-art methods [121, 137] employed on the AMI dataset under unsupervised settings.

4.2.4 Enabling exploration of an augmented meeting transcript through an interactive interface to help organizers author meeting reports (G4). We developed CommunityClick's interface as a web application. It allows multi-faceted exploration and analysis of meeting data by providing components including the title, filters, discussion topics, timeline, summary, text editor, and finally the augmented meeting transcript segments (Fig. 4). The title contains the metadata about the meeting, including the meeting title, date, and location (Fig. 4(A)). The filters allow organizers to explore the transcript segments according to the selected feedback or tags of interest (Fig. 4(F, H)). These options are customizable, and organizers may customize the tags to suit their purpose before the meeting. We chose to visually collapse transcript segments that are filtered as opposed
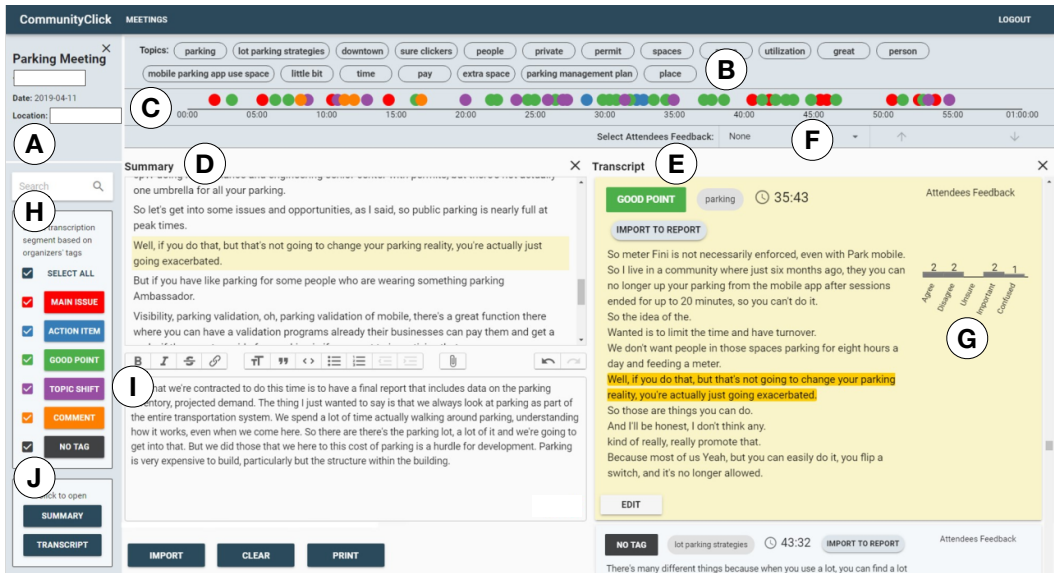
Fig. 4. A snapshot of CommunityClick's interface. A) The title provides useful metadata about the meeting such as date and location. B) The main topics extracted from the meeting transcript. C) The timeline visualizes organizers' tags in chronological order. Each circle represents a tag. Clicking on a circle brings organizers to the corresponding transcript segment. D) The interactive feedback-weighted summary. E) Transcript view displays the transcript text alongside organizers' assigned tag, the main topic, and aggregated attendees' feedback in that time interval for each segment. F) Filters for attendee's feedback (based on what was provided on iClicker). G) The bar chart displays attendees' feedback. H) Filters for organizer's tags. I) Rich text editor for organizers to author the report. J) Options to view or collapse the summary or the transcript view.

to completely removing them from the view to communicate to the users that there are additional conversations that transpired between the segments currently visible.

In the topic and timeline component, we provide the list of most relevant topics and the timeline of the meeting discussion (Fig. 4(B, C)). The organizers can filter the transcript segments based on any topic. The timeline displays the organizers' tags using circles in a chronological manner, where each circle represents a tag, and the color corresponds to organizers' tags (Fig. 4(C)). This provides the organizers with a temporal distribution of tags that demonstrates how the conversation progressed during the meeting. When a circle is selected, the transcript is scrolled to the corresponding segment and highlights the background to distinguish it from other segments.

The feedback-weighted extractive summary is presented in a textbox (Fig. 4(D)). Each of these sentences is interactive, and upon selection, they navigate to the transcript segment it was extracted from. This can enable organizers to explore the transcript and get a better understanding of why the sentence was added to the summary. Below the summary, we added a rich text editor for authoring the meeting report with rich formatting options (Fig. 4(I)). We also added options for attaching additional files or images. Once the report is created, it can be printed in PDF format directly, without switching to other external printing applications.

Finally, we present the augmented transcript divided into transcript segments (Fig. 4(E)). The segments are ordered chronologically. Each transcript segment contains the transcript text, associated organizer's tag, the most relevant extracted topic, time of the segment, option to import the summary of the selected transcript to the text editor, and aggregated attendees' feedback in

the form of a bar chart. For easy tracking, we highlight the transcript text that are added to the summary. Organizers can edit the segments to assign or change tags and topics. However, they do not have control over attendees' feedback to mitigate bias injection. To reduce clutter on the screen, we added two additional filters to collapse the summary or augmented transcript (Fig. 4(J)).

4.3 Pilot Study

To explore whether CommunityClick could be effectively deployed in a real-world town hall meeting, we performed a pilot study where we simulated a town hall with nine participants. We recruited eight participants as meeting attendees and one participant as the organizer, who had previous experience with organizing meetings. We refer to the attendees who participated in our pilot study as **PA**. We recruited all participants using word of mouth from a public university in the U.S. For the discussion topic, we selected two contentious issues regarding the university that were popular at the time of our pilot study. The topics included discussions around building a new common room for graduate students and the rise of racist activities across the campus. All participants were graduate students (6 males and 2 females with a average age of 27.25). The goal of the pilot study was to assess the system workflow for potential deployment and whether the attendees could share their feedback silently using iClickers without interrupting others. Furthermore, we used the augmented transcript from the meetings to enable the pilot study organizer we recruited to explore the meeting discussions to identify potential interface issues.

The meeting took 60 minutes which is similar to a traditional town hall. We collected 292 items of feedback from attendees (avg 36.5 per attendee ± 8.23) and 56 tags from the organizer. After the meeting, we asked attendees to share their experiences of using iClickers to voice their opinions using open-ended questions. The findings suggested that attendees found iClickers easy to get used to. They were able to share their feedback silently using iClickers, managed to avoid potential confrontations, and thought they could contribute more compared to their experiences in other meetings. However, one attendee mentioned about difficulties around remembering which iClicker button mapped to which attendees' tag. Another attendee mentioned that while expressing strong agreement or disagreement with the ongoing discussion, some attendees might spam the iClicker button which might "*dilute the opinion values*" (PA-4). The organizer mentioned that using iClickers enabled him to focus more on the discussion and the interface allowed him to better capture attendees' feedback. He recalled that some attendees were silent, but the bar charts showed feedback from all eight participants, meaning they were participating silently. He also identified the flow of the discussion, and important discussion points using the interface.

This pilot study helped us to better understand and solidify operational procedures to perform real-world deployment of our system. Based on the feedback we received, we modified the system and user interface. For instance, we added a spamming prevention technique [123] by calibrating the system to capture one click from each attendees' iClicker in a 30-second window to negate the possibility of diluting the values of specific feedback options. Furthermore, we added an option to collapse the summary or transcript to reduce interface clutter. Finally, we used written labels on iClickers to reduce attendees' cognitive load in remembering the mapping of response options.

5 EVALUATION

We evaluated the application of iClickers as a real-time response system, particularly for the ability of silent attendees to share feedback, and the efficacy of our approach in enabling organizers to explore, capture, and incorporate attendees' feedback to author more comprehensive reports. To that end, we conducted a field experiment to examine if the attendees could effectively use iClickers to voice their feedback. In addition, we followed up by conducting semi-structured interviews

with 8 expert organizers to evaluate if CommunityClick could enable them to capture attendees' feedback and generate more comprehensive reports.

5.1 Field Experiment: Parking Town Hall

We deployed CommunityClick at a town hall in a college town in the U.S. The meeting focused on a new set of proposals to improve the parking condition of the town. We reached out to the town officials a month before the meeting took place. They allowed us to deploy our system, record the discussion, and use the data. They also introduced us to the organizer who facilitated the meeting. We explained the procedure to them and discussed the tags to be used for both the organizer and attendees. The town hall took place on a Thursday evening and was open for all to attend.

5.1.1 Meeting Participants. There were 31 attendees and 1 organizer present in the meeting. We provided attendees' iClickers with Agree, Disagree, Unsure, Important, and Confused tags to 31 attendees. For the organizers, we provided them iClickers labeled with Main Issue, Concern, Supportive, New Idea, and Good Point tags as per our pre-meeting discussion.

5.1.2 Procedure. At the beginning of the town hall, we provided a brief tutorial for five minutes on how to use the iClickers to the meeting attendees. We also received consent from the attendees about recording their discussions. The meeting began with the organizer presenting the meeting agenda and the new parking proposals to the meeting attendees. The attendees and the organizer used iClickers to tag the conversations throughout the meeting. After the presentation, the attendees engaged in discussing the proposals. The meeting lasted for 76 minutes. At the end of the meeting, we provided post-study questionnaires to attendees that asked various questions, such as their reasons behind attending the meeting, their usual experience during town halls, whether they could share their opinions by speaking up, and how did using iClickers compare to such experiences. They responded on a five-point Likert scale. We also asked them open-ended questions around their experience of working with iClickers, whether they faced any issues or challenges, and suggestions to improve their experiences and our approach. The post-study questionnaire is provided as a supplementary material.

5.1.3 Data Collection and Analysis. We were given permission to collect and use the data from the town hall by government officials. We collected 61 minutes of meeting audio for transcription. We also collected organizer's tags and attendees' feedback from the meeting. In total, we captured 56 tags from the organizer. Out of 31 meeting attendees, 22 used the iClickers we provided to share their feedback. Out of these 22 attendees, 20 of them filled up the post-study questionnaire. We report the statistics based only on these 20 attendees' responses. We captured a total of 492 attendees' feedback with an average of 24.6 feedback items per attendee with a standard deviation of 6.44. This data was later used to populate CommunityClick's interface for demonstrating its various functionalities to meeting organizers, which we describe in 5.2. We also collected the post-study questionnaire responses and entered them into spreadsheets for creating charts (Fig. 5) and statistics for analysis.

5.1.4 Findings. From the analysis of the attendees' iClicker usage patterns, we found that the attendees' used the tag *Agree* the most with 187 clicks (38%), followed by *Important* with 103 clicks (21%), *Disagree* with 93 clicks (19%), *Confused* with 79 clicks (16%), and finally *Unsure* with the least amount of 30 clicks (6%) only. Initially, we were surprise to see the large gap between agreement and disagreement. However, upon closer inspection, we found that on several occasions, the attendees who were using iClickers were clicking *Agree* when other vocal attendees were verbally expressing their disagreement to a discussion topic. This behavior pattern indicates that the silent attendees used iClickers to provide their support for an ongoing argument alongside
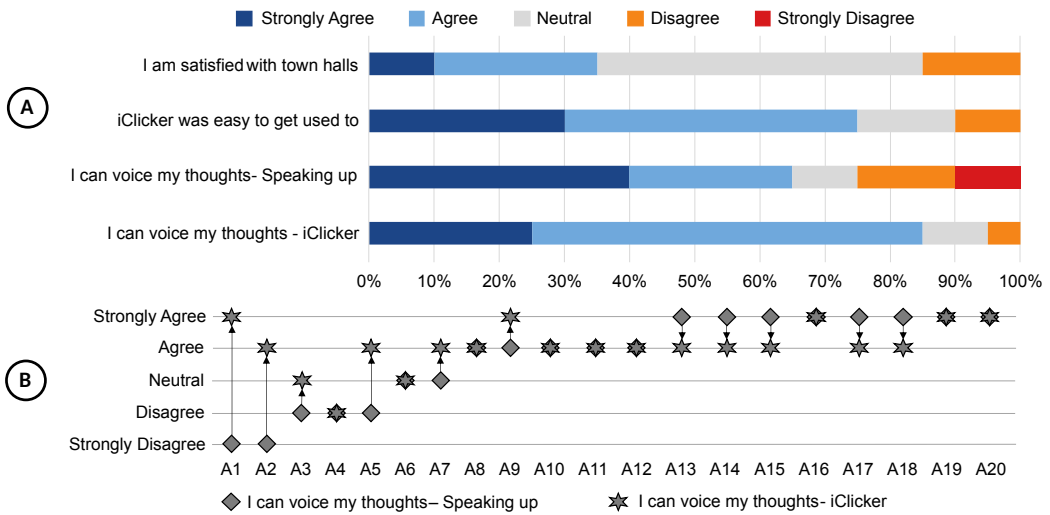
Fig. 5. The results from the field experiment. A) Attendees' responses show that the majority of meeting attendees were not satisfied with the status quo of town halls but found iClickers easy to get used to. It also displays the number of attendees who thought they could share their voices by speaking up or using iClickers. B) A deeper comparison between speaking up and using iClickers to share attendees' feedback. The diamonds (◆) and stars (✳) represent 20 attendees' (A1-A20) responses to questions that asked them to rate their experiences of sharing opinions by speaking up and using iClickers respectively during town halls. The arrows show the difference and increase or decrease in their ratings. The arrows demonstrate that the majority of the participants who were not satisfied with the current methods of sharing opinions (A1-A3, A5) during town halls found iClickers to be a better alternative. They also show the participants who were comfortable with speaking up during meetings, did not endorse iClickers as strongly to share their voices.

sharing their own opinions. We also found that the attendees did not press any iClicker options during the introduction when the organizers were setting up the discussion and conclusion of the meeting when the organizer expressed his gratitude for attending the meeting and continued with other social conversations. This suggests that the attendees took their opinion sharing using iClickers seriously and did not randomly click different options during the meeting.

From the analysis of the post-study questionnaires, we found that all of the 20 attendees either lived or worked in the town of Amherst, where the town hall was organized. 95% of these attendees (19 responses) were well-accustomed with such meetings and they mentioned attending similar town halls twice a year, while 50% (10 responses) attended these meetings more than five times per year. When asked about their exposure to technology, all meeting attendees responded to owning at least one computer and one smartphone with an internet connection, which they were comfortable using. However, their responses to the experiences in town halls varied as presented in Fig. 5(A). 25% of attendees (5 responses) mentioned that they did not feel that they are able to voice their thoughts in town halls. Only 35% (7 responses) were pleased with the way such town halls are organized while 50% of the attendees (10 responses) were neutral in their responses. 75% attendees (15 responses) responded that they got used to iClickers quickly. 85% mentioned (17 responses) they were able to share their thoughts using iClickers compared to only 65% (13 responses) attendees who are comfortable with speaking up to share opinions. The majority of the attendees (90%) were positive about their experiences of using iClickers to share their opinions (18 responses). One

attendee (A9) mentioned, "*I feel like I could contribute more than usual and I would definitely like use it in future meetings.*"

We further compared the attendees' responses on their ability to share their thoughts in town halls by voicing their opinions against using iClickers (Fig. 5(B)). The data shows that almost all of the attendees who did not think they could share their thoughts by speaking up, except for one (A4) thought they could voice their opinions using iClickers. One such attendee (A2) mentioned, "*I didn't like what others were saying. But instead of interrupting, I just used the clicker to say that I didn't agree with them.*" We also found that, while agreeing that iClickers' could provide a way to voice opinions in town halls, the attendees who strongly preferred speaking up did not rate their experience of using iClickers to share opinions as high. One of these attendees (A17) mentioned, "*I was distracted, so I didn't use it that much.*"

We identified two important insights from this field experiment (Fig. 5). Firstly, it demonstrated that the silent attendees' who could not speak their minds found a way to voice their opinions without the apprehension of confrontation in town halls using iClickers (Fig. 5(A) and (B), attendees A1-A3, A5). However, we also found that some attendees who strongly agreed that they were satisfied with the current method of sharing opinions by speaking up, did not as strongly endorse iClickers as a way to share their opinions (Fig. 5(B), attendees A13-A15, A17, A18). We speculate two reasons for such reduced ratings for iClickers. First, for the attendees who are already comfortable with speaking up, iClickers might seem like an additional step to share opinions which might lead to distractions, as mentioned by one of the attendees (A17). The second reason might be a reluctance to deviate from the norm and use technology within the established albeit impaired town hall proceedings and customs. Nevertheless, our results suggest that the addition of iClickers could be an acceptable trade-off between providing the silent attendees a way to communicate their opinions and mildly inconveniencing the adept and vocal meeting attendees.

5.2 Semi-structured Interviews with Meeting Organizers

We conducted semi-structured interviews with 8 expert meeting organizers, who were experienced with organizing or facilitating town halls to gather data on community's needs, issues, and ideas. They were also adept at compiling meeting reports that play a pivotal role in informing civic decision-making. Our objective was to examine if CommunityClick's interactive interface could help the organizers to better capture the attendees' feedback to author more comprehensive reports that preserve the equity and inclusivity of voiced opinions in town halls.

5.2.1 Participants. We reached out to a total of 29 expert organizers from across the U.S. 8 of them responded by agreeing to help us with evaluating CommunityClick. Our interviewees were experts in their fields with intimate knowledge of town hall organization and decision-making. We refer to our semi-structured interview participants as **P**. On average, they had over 20 years of experience. One interviewee (P1) was the organizer from our field experiment—the town hall on parking. We made connections with the others (P2-P8) during our formative study. All of our interviewees were based in the U.S.

5.2.2 Procedure. Several experts we engaged with to evaluate CommunityClick were excited about its potential and agreed to deploy the system in their then upcoming town halls. Our original evaluation plan involved several deployments in the wild, then providing organizers with the meeting audio from these deployments and ask them to write two reports, one using their current method of writing reports, and the other using CommunityClick's augmented meeting transcript and interactive interface. We wanted to study the differences between these reports to investigate the efficacy of our system. However, the recent COVID-19 pandemic forced the organizers to cancel

Table 1. This table shows themes that emerged from analyzing the interviews with the organizers. The codes associated with the themes and their description is also presented in the table.

Themes	Codes	Descriptions
Enabling inclusivity	Equitable platform, problem speaking up, opinion sharing, understanding others, inclusive opinions	CommunityClick's impact on inclusivity in town halls
Diverse perspectives	Shared narrative, honest reflections, attendee's reactions, identifying conversation flow	Different perspectives and opinion shared in town halls
Report quality	meeting summarization, missing information, credible process, comprehensiveness, accurate representation	CommunityClick's utility in creating reports
Meeting organization	Unstructured discussions, real-time attendee's response, tracking response, customized tags, measuring consensus	Organizing meeting-generated data
Interface learnability	Intuitiveness, easy-to-use, formatting, data exploration	Users' ability to learn and use interface features
Concerns and caveats	Technology as a barrier, tech-savvy, distraction factors, young generation	Concerns regarding CommunityClick's usage in town halls
Improvement suggestions	Real-time feedback, opinion statistics, organizers' input	Suggestions to improve our approach

all town halls until further notice, and we were compelled to cut our evaluation short. Due to this setback, we revised and modified our evaluation procedure as follows.

We deployed the CommunityClick interface on a public domain and shared it with our interviewees via emails at least two weeks before our interviews. To maintain privacy of usage, we provided each organizer with their own user account and login credentials. We also provided detailed instructions on how to use CommunityClick's various features and encouraged the interviewees to explore the interface at their own convenience. We populated the interface with the data collected from the simulated meetings from our pilot study as well as the meeting from our field experiment for the interviewees to explore. During the interview sessions, we asked them open-ended questions focusing around their current practices towards town hall organization, how using CommunityClick differed from these practices, how useful could CommunityClick be to capture silent attendees' feedback and marginalized perspectives, could the interface allow them to author better reports, and finally, suggestions to improve CommunityClick. We also allowed them to ask any questions they might have about CommunityClick. The interview questions are provided as a supplementary material. We conducted the meetings over video conferencing via Zoom [12]. The interviews lasted between 45-60 minutes. All participation was voluntary. Each interview was conducted by an interviewer and a dedicated note-taker from our research team.

5.2.3 Data Collection and Analysis. We transcribed over 400 minutes of audio recording from our interviews with organizers. We also took extensive notes throughout the meeting. Finally, we thematically analyzed the interview transcripts and notes taken 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 0.89 which was measured using Krippendorff's alpha [84]. We had several iterations of discussions among the research team to condense the codes into the themes presented in Table 1.

5.2.4 Findings. Our analysis of interview transcripts and notes surfaced critical insights on how CommunityClick could enable attendees to share opinions and help organizers to capture inclusive feedback to author more comprehensive reports. We elaborate on these insights in the following and discuss possible room for improvement.

CommunityClick can create a more inclusive platform to share opinions. Our intervie-wees were unanimous (8 out of 8) in acknowledging CommunityClick's potential to create an inclusive platform for community to share their opinions. They shared with us several example scenarios they experienced where CommunityClick could have provided silent attendees a way to speak their minds. P1 mentioned, "*People want to share their opinions, but sometimes they just can't express themselves because they're not comfortable talking, or they're nervous about how they'll appear to or who is around the table. Often they are intimidated. Here,* "intimidated" *is a strong word. But I don't think it's the wrong word.*" P2 mentioned,

> *There was an individual who attended several meetings, it was clear that their presence had an impact on people's willingness to speak at all, or the opposite effect, where people escalated in reaction to that person. Giving them the ability to click help both ways. They can avoid confrontation or avoid escalation by just clicking.*

Similarly, P3 drew examples from his experiences, saying,

> *Even if the attendees are from the U.S., [people with] different upbringings or cultural backgrounds have a disadvantage to those who are quite familiar with the moors of group dynamics. In our town halls, we only take votes on questions or get agreements, but in a conversation, there are so many euphemisms, colloquialisms, and metaphors that make it difficult for someone unfamiliar with them to understand others' reactions. There is real value in using options like "confused" and "unsure" to allow them to record that they didn't understand the conversation instead of forcing them to agree or disagree.*

P6 found further value in separating the attendees' tags and the organizers' tags to establish organizers' credibility. She mentioned, "*The organizers cannot unintentionally skew the attendees' feedback because [their tags] are separate. That way, we know the recorded feedback is unbiased.*"

The augmented transcripts provide evidence of attendees' reflections. One of our primary goals was to enable organizers to have access to attendees' perspectives to form a shared narrative of the meeting discussions. After exploring CommunityClick's interface, the majority of interviewees (7 out of 8) mentioned how it enabled them to capture meeting attendees' reflections on the meeting agenda. P6 mentioned, "*It provides a way of ensuring that voices and reactions are reflected as people speak and click. It is a huge step towards having a more honest reflection of what really went on in the meeting.*" P3 further emphasized how CommunityClick not only captured the attendees' feedback but also allowed navigation of the conversation flow using the Timeline,

> *This tool allows me to see both how many ideas have traction or agreement and how many don't, but just as importantly, how the flow went. The facilitators are concerned with the way topics are discussed in town halls. These topics are influenced by the surrounding conversations. It [Timeline] allows me to see reactions that might or might not be intended because of the sequence of conversations. Having a way to track that has a huge value.*

Regarding the interactive augmented transcript, P4 specifically preferred the way it enabled her to track attendees' responses. She drew a comparison with her usual methods for note-taking during town halls saying,

> *We usually have a table facilitator and then a table observer. The table observer takes detailed notes, but it adds to the number of staff we have to have. So that creates an additional challenge, but the speech to text transcription makes a big difference in recording people's reactions. With [CommunityClick], maybe we won't need a table observer.*

P5 also mentioned how CommunityClick gave credence to attendees' reactions during the meeting discussions through the feedback bar charts. She said,

It makes a lot of sense to see where people are aligned, where the challenges are, and giving information from their reactions. When changing policies, we hear from only a few voices who are either for or against something, and they tend to dominate the conversation. Having a visual and data-driven way to show what was the actual participation is gold. Sometimes people feel that a proposal is not aligned with their ideas. With the bar chart, you can show them that maybe you are the outlier and others agree with the proposal.

CommunityClick can help create more comprehensive and accurate reports. All of our interviewees had prior experiences of writing reports by summarizing the meeting discussions and identifying key discussion points. They found various aspects of CommunityClick useful to not only author more comprehensive reports but also more accurate ones that lend credibility to the report creation process. P1 drew parallels with his experience of working in scenarios where designated note-takers took notes and his team generated the reports from those notes. He mentioned, "*People who take notes have varying abilities and the notes vary in quality. Instead, as you are writing reports, you have [CommunityClick], where you can see and add the reactions to what [attendees] discussed right away, it builds credibility for the process.*" P3 echoed similar sentiments, saying,

You are usually distracted by the conversation while taking notes, which means you might miss every third word at a particular moment, which could be the difference between agreement and disagreement. Having it transcribed and summarized will remind a facilitator of some things that he or she may not have remembered or make it more accurate, because they may have remembered it differently. I love the fact that the [text analysis methods] can capture that objectively for us.

P4 also emphasized the usefulness of importing a summary to the text editor. She mentioned,

Having the text editor where you can start writing the report and pull in pieces from the transcript could be really helpful, because then as you read through the transcript and you're writing about some themes, you can pull characteristic quotes that would really help bring in more evidence for claims for those themes.

Furthermore, we found that the report creation process can take a few hours to a few days depending on variables such as the way notes were taken, the length of the meeting, report creators' skills, etc. P7 highlighted the reduced workload and efficiency that CommunityClick could provide, saying, "*There is a physical component of getting into it, typing it up, theming, organizing, and editing which always takes longer than anticipated, I can see some of those issues can be fixed with this.*"

CommunityClick preserves the flow of meeting discussions by establishing an implicit structure. The majority of our interviewees (6 out of 8) thought CommunityClick could be best utilized to organize unstructured meeting discussions. They emphasized that contrary to asking meeting attendees to respond to specific questions in town halls, CommunityClick allowed attendees to respond whenever they wanted, creating an implicit structure to the meeting while preserving the flow of discussion. One interviewee (P1) mentioned, "*[CommunityClick] would provide the biggest benefit in more unstructured kind of discussions. If you have a town hall, where people are less likely to speak up, [tags] would be helpful to understand their reactions and help with the theming.*" Another interviewee (P5) mentioned, "*It's hard to keep track of many ideas, but the visual organization of information helps to gauge reactions and figuring out if we reached consensus. But most importantly, it helps me to see if there are any social or racial injustice components into the proposals where there can be negative reactions.*" They also found the option to customize the attendees' feedback and organizers' tags useful to adapt to different meeting scenarios. One interviewee (P3) mentioned, "*Words may mean different things in different meetings. Having the ability to label and*

customize [tags] individually would be a way for different organizations to adjust to that. Sometimes we want to know [attendees'] hopes and concerns, but other times, we just want to know if they agree or disagree." However, P7 raised a concern about larger meetings, saying, "I think [CommunityClick] will be useful for smaller meetings, but if there are hundreds of people, and everyone is speaking over each other, I'm not sure if you will be able to cope with that."

The simplicity and learnability of the interface affords intuitive exploration. From a usability standpoint, all of our interviewees (8 out of 8) found CommunityClick's interface to be simple and straightforward to work with. P3 extolled the interface saying, "It's very intuitive, simple, easy to use, and navigate after the fact, edit, and update. All user interface features look well-thought-out to me considering the inner workings are extremely delicate and complicated." P4 valued the rich editing options of the text editor. She said, "The automatic summaries can be used as a starting point of the report, as an initial cut, and then I can delete the things that might not be very useful and build up the report by adding more to it and formatting it. I can clearly see a lot of thought was put into designing the interface." P5 thought that the interactivity of the timeline was useful for navigating the augmented transcript. She mentioned, "When I started clicking on the buttons on the circles at the top [timeline], it was very intuitive, like, it just automatically brings you to the places where that correlates with the statements, so you understand what it's connected to." P6 further emphasized CommunityClick's potential as a record-keeping and meeting exploration system, saying, "Everything is linked together. So in that sense, it makes intuitive and logical sense when I'm looking at the data. It will be a total game-changer for policymakers and community organizers."

Concerns around technology in town halls. Although our interviewees praised CommunityClick, some of them raised a few important concerns. P8 mentioned how technology usage could be troublesome in town halls, saying, "It feels like the technology itself could be seen as a barrier, because a lot of people might not feel quite as comfortable, clicking on things and reacting to." On a different note, P4 raised concerns about the sense of urgency such technology might impose on the meeting attendees. She said, "[Attendees'] reactions, they are decisions that are being made in the spur of the moment as they're hearing information. And it's such a complex, sociological, and psychological response to information." Her concern was whether the urge to immediately respond to someone's perspectives could inhibit the ability to contemplate and deliver a measured response. Further concerns arose from P5, who mentioned how younger meeting attendees might have an advantage in the town halls if the technology is heavily used, saying, "Younger generations tend to use technology so much more easily. And they turn to it so much more easily than older generations."

Possible room for improvements. We received some feedback from our interviewees on how to improve CommunityClick. Some of these suggestions focused on adding more real-time components to CommunityClick that can further augment the ongoing discussions. For example, P3 mentioned, "Right now, [CommunityClick] helps me to understand attendees' reactions after the fact. I think if the facilitators could see them in real-time as the attendees are clicking away, they might be able to steer the conversation better to be more fruitful and fair." Other suggestions involved adding functionalities to add the organizers' own topics on top of the automatically extracted topics. In that regard, P4 mentioned, "I guess it depends on who is using the system, but we look to dive a little bit more and would want to maybe customize the topics and themes".

6 DISCUSSION

From our field experiment and interviews with experts, we found that CommunityClick could be used to create an equitable platform to share more inclusive feedback in town halls. Compared to

prior approaches to utilize technology in town halls [25, 103], CommunityClick enabled attendees to utilize iClickers to silently and anonymously voice their opinions any time during the meeting without interrupting the discussion or the apprehension of being shut out by dominant personalities. We extended the use of audience response systems in prior works by adding five customizable options that go beyond binary agreement or disagreement and modifying iClickers as a real-time response system [22, 25]. The experts we interviewed valued options such as *confused*, or *unsure* to identify if attendees were disengaged or did not understand the conversation without forcing them to agree or disagree [76, 119]. The customizability of both organizers' tags and attendees' feedback further added to the flexibility of our approach that could potentially increase adaptability in meetings with diverse agendas in both civic and other domains. Moreover, the automation and augmentation of speech-to-text transcription, extraction of most relevant discussion topics, and feedback-weighted summarization of meeting discussion could potentially eliminate the need for separate note-takers. According to the organizers we interviewed, it could reduce the manpower requirement significantly compared to established approaches [90, 91].

From organizers' perspectives, CommunityClick could help create more comprehensive and accurate reports that could provide evidence of attendees' reflections. Prior work experimented with annotations during face-to-face meetings [74, 78] for memory recall. In our work, we empowered organizers to go beyond recollection of events during meetings by integrating their own customized tags and enabled them to capture a more comprehensive picture of the meeting discussion. Furthermore, our interactive summary and attendees' feedback visualization provided a visual and data-driven way to highlight attendees' viewpoints, and outliers on critical points of discussions. This could enable organizers to receive a more clear reflection of what occurred in the meeting and author more accurate reports based on tangible evidence rather than incomprehensive interpretation [94], that could further lend credibility to the report creation process. Prior work highlighted concerns about accuracy in computational approaches to analyze meeting data [97]. However, from our interviews with experts, we found that the accuracy and comprehensiveness of meeting reports often depends on meeting length, method of taking notes during the meetings, and note-takes' skills. We posit that synchronization of meeting data and addition of inclusive attendees' feedback into the summary and interface will enable organizers to author more accurate reports with the added benefit of reduced manpower requirement. Although during our interviews with meeting organizers, some of them highlighted that the augmented transcripts, discussion topics, and summaries could only be accessed after the meeting is finished, we argue that the latency is an acceptable tradeoff for increased comprehensiveness, credibility, and accuracy in generating reports. Furthermore, enabling access to the variety of meeting generated data could also help to reduce both organizers' and attendees' distraction during the meeting [15, 61] and allow them to engage with the ongoing discussion. In particular, the separation of attendees' feedback and organizers' tags along with the evidence of attendees' feedback in meeting reports could pave the way to instill trust between the decision-makers in the local government and the community members, which is considered to be a wicked problem in the civic domain [39, 41, 63, 94].

6.1 Marginalization in Civic Participation and the Role of Technology

Marginalization can be broadly defined as the exclusion of a population from mainstream social, economic, cultural, or political life [58], which still stands as a barrier to inclusive participation in the civic domain [48, 94]. Researchers in HCI and CSCW have explored various communitysourcing approaches to include marginalized populations in community activities, proceedings, and designs [48, 53, 81, 93, 132]. In this work, we added to this body of work by designing technology that included silent voices in civic participation to increase the visibility of marginalized opinions. Our field experiment and interviews with experts demonstrated the efficacy of our approach to

enable, capture, and include silent attendees' participation in town halls, regardless of their reasons behind keeping silent (social dynamics, fear of confrontation, cultural background, etc.). Our work also answers the call for more inclusive data analysis processes and practices to augment computational approaches [50, 114] by proposing a text summarization method that included and prioritized attendees' feedback when generating meeting discussion summaries. Such inclusive data analysis techniques could reflect the community's opinions, identify social and political injustice, and support accountability in the outcome of civic engagements [53]. However, designing communitysourcing technologies to include marginalized opinions and amplify participation alone may not be enough to solve inequality of sharing opinions in the civic domain [26, 126]. Despite the success of previous works [25, 53, 90], technology is rarely integrated with existing manual practices and follow-ups of engagements between government officials and community members are seldom propagated to the community. This lack of communication might lead to the uncertainty of attendees on whether actions will be taken to reflect their opinions. As a result, the power dynamics between the government officials and the community remain unbalanced, especially for the marginalized populations [48, 49, 53, 125]. One way to establish the practicality of using technology to include marginalized opinions is to integrate them into existing processes to convince the officials of its efficacy. Our work provides first steps towards integration of marginalized perspectives, however, long-term studies are required to assess the possibility and feasibility of using such technologies to integrate public perspectives into the actual civic decisions.

6.2 Integrating CommunityClick in Today's Town Halls

Our formative study with 66 attendees and field experiment with 20 attendees bore a striking resemblance regarding the attendees' ability to share their opinions in town halls. We found that 17% (11 out of 66) of attendees from the formative study and 25% (5 out of 20) of attendees from the field experiment were not comfortable to speak up in town halls and needed a way to share their opinions. However, similar to previous works [56, 90], some of the meeting organizers we interviewed were apprehensive towards depending solely on technology due to concerns around logistics involved with the procurement and maintenance of electronic devices such as iClickers, reliability concerns of using technology that require proper management, and unfair advantage towards newer generation who are more receptive to novel technologies. We argue that renewed motivation, careful design choices, and detailed instructions could help overcome the novelty barrier of technology even for people from older generations [44, 129]. Based on our experiences from this study, we advocate for integrating technology with current face-to-face methods. We do not suggest complete replacement of traditional methods with technology, rather we suggest augmenting them with technology to address some of their limitations.

CommunityClick could be gradually integrated as a fail-safe or an auxiliary method to complement the current process. All the functionality of CommunityClick would remain operational while the organizers could take personal notes in parallel alongside their iClicker interactions. This would allow organizers to retain their current practices while taking advantage of augmented meeting transcripts, discussion topics, and summaries from CommunityClick to better capture and understand attendees' perspectives when compiling reports. Another way to integrate CommunityClick into current processes would be to provide statistics of attendees' feedback so that the organizers could track the discussion dynamics and facilitate the conversation accordingly to keep attendees engaged. However, prior works suggested that real-time visualization or displays can add to attendees' distraction, causing them to disengage from the ongoing discussion [15, 61]. To circumvent this issue, optional companion mobile applications could be introduced only for organizers to receive real-time feedback on attendees' iClicker responses so that they can make course corrections accordingly without distracting the attendees or influencing their opinions.

6.3 Expanding CommunityClick to other Domains

From our field experiment and interviews with the experts, we demonstrated CommunityClick's potential in elevating traditional town halls. We argue that our proposed data pipeline can be expanded to other domains where opinion sharing is important for successful operation. For example, in education, whether in classroom settings or in massive open online courses [51, 73, 83], it could enable students to share their feedback anytime without interrupting the class on whether they could understand a certain concept, or if they were confused and needed more elaboration, especially when the class size is large. For educators, it could allow them to track the effectiveness of their content delivery, class progress, and student motivation, which might help them to adjust the course curriculum more effectively and efficiently, instead of receiving feedback from students once every semester. More importantly, familiarity with iClickers in education eliminates the entry point barrier for technology making the system readily adoptable [14, 134].

Similarly, CommunityClick could be used as a possible alternative to expensive commercial applications in meeting domains within the business and other corporate organizations [1, 2]. Similar to town halls, in a corporate setting, CommunityClick could provide text summary and direct evidence of attendees' feedback from the meeting transcripts for better decision-making. Automatic text summarization remains a challenging problem [124], especially when it comes to automatically summarize meeting discussions [57]. Our summarization approach emphasized the importance of meeting attendees' feedback instead of purely text-based summarization approach that treated all sorts of text documents similarly [99, 124]. We argue that this *feedback-weighted* summary could be valuable in generating domain-specific contextual text summarization in other meeting genres. Furthermore, there are potential applications of CommunityClick as a record-keeping tool which might be particularly applicable in journalism where the journalists can utilize the iClickers to annotate the interview conversation with important tags and later review the augmented interview conversation using CommunityClick's interface to better organize and write news articles or interview reports.

7 LIMITATIONS AND FUTURE WORK

Our evaluations suggested the efficacy of CommunityClick in providing a voice to reticent participants and enabling organizers to better capture the community's perspectives. However, we had only one real-world deployment of CommunnityClick at a town hall due to the pandemic. We will continue to engage with meeting organizers and deploy CommunityClick in town halls to study the long-term impact of CommunityClick and attempt to gradually integrate our approach in the predominantly manual town hall ecosystem. Also, we found that iClickers might be distracting as a new technology for some attendees in town halls as the findings from our field experiment suggested. Furthermore, the logistical issues associated with both hardware procurement and the unavailability of software APIs might be a hindrance to some communities. To circumvent these issues, low-cost alternatives to iClickers or fully software-based substitute audience response system applications could be utilized [10, 11]. A fully software based solution could also enable attendees to provide open-ended textual feedback which CommunityClick does not allow in its current state. However, further comparative studies are required to assess the cost, efficacy, and applicability of such alternatives to replace iClickers that would provide the same benefit without incurring additional financial, computational, or cognitive overhead.

There are several avenues to improve CommunityClick in the future. For example, it could be augmented with more real-time components including a companion application to deliver dynamic feedback statistics for organizers to access and utilize during the meeting. We will also explore novel methodologies to speed up the automatic speech to text transcription process to further

reduce the time required for data processing. One approach could be to utilize parallel pipeline processing [34] where audio signals from the meeting will be processed concurrently, which might reduce processing time significantly.

To further provide evidence from attendees' feedback to help organizers when authoring reports, the audio of discussions could be added and synchronized with the transcript segments. It could enable the organizers to identify vocal cues and impressions from the attendees who spoke during the town hall to further contextualize the discussion segment [98]. In addition, we will investigate the possibility of tracking individual iClickers IDs without risking the privacy of attendees to anonymously identify potentially contentious individuals who might be pushing specific agendas or marginalize minority viewpoints in a discussion. However, further studies are required to understand the potential computational challenges that might arise with such extensions.

To further improve the report generation, we will explore novel technologies in natural language generation [133] to automatically write meeting reports that could further reduce meeting organizers' workload. In addition to exporting the created reports, we will further enable exporting various statistics around attendees' feedback, organizers' tags, discussion topics, etc. to be added to the report or used separately in presenting the outcome of the town halls to decision-makers.

8   CONCLUSION

In this study, we investigated the practices and issues around the inequality of opportunity in providing feedback in town halls, especially for reticent participants. To inform our work, we attended several town halls and surveyed 66 attendees and 20 organizers. Based on our findings, we designed and developed CommunityClick, where we modified iClickers as a real-time response mechanism to give voice to silent meeting attendees and reflect on their feedback by augmenting automatically generated meeting transcripts with organizers' tags and attendees' feedback. We proposed a novel feedback-weighted text summarization method along with extracting the most relevant discussion topics to better capture community's perspectives. We also designed an interactive interface to enable multi-faceted exploration of the summary, main discussion topics, and augmented meeting-generated data to enable organizers to author more inclusive reports. We deployed CommunityClick in-the-wild to conduct a field experiment and interviewed 8 expert organizers to evaluate our system. Our evaluation demonstrated CommunityClick's efficacy in creating a more inclusive communication channel to capture attendees' opinions from town halls and provide evidence of attendees' feedback that could help organizers to author more comprehensive and accurate reports to inform critical civic decision-making. We discussed how CommunityClick could be integrated into the current town hall ecosystem and possibly expanded to other domains.

REFERENCES

[1]   2015. *ICompassTech*.   https://www.icompasstech.com/
[2]   2016. *MeetingKing*.   https://meetingking.com
[3]   2017. *Voicea*.   https://www.voicea.com/
[4]   2019. DemocracyOS. http://democracyos.org/.
[5]   2019. Free open-source participatory democracy for cities and organizations. https://decidim.org/.
[6]   2019. *Mobile Fact Sheet*.   Retrieved August 28, 2020 from https://www.pewresearch.org/internet/fact-sheet/mobile/
[7]   2019. Pol.is. https://pol.is/gov.
[8]   2020. *Assembly AI*.   https://www.assemblyai.com/
[9]   2020. *iClicker*.   https://www.iclicker.com/
[10] 2020. *Slido*.   https://www.sli.do/
[11] 2020. *VoxVote*.   https://www.voxvote.com/
[12] 2020. *Zoom*.   https://zoom.us/
[13] Brian Adams. 2004. Public meetings and the democratic process. *Public Administration Review* 64, 1 (2004), 43–54.

[14] Stephen Addison, Adrienne Wright, and Rachel Milner. 2009. Using clickers to improve student engagement and performance in an introductory biochemistry class. *Biochemistry and Molecular Biology Education* 37, 2 (2009), 84–91.

[15] Katy Appleton and Andrew Lovett. 2005. GIS-based visualisation of development proposals: reactions from planning and related professionals. *Computers, Environment and Urban Systems* 29, 3 (2005), 321–339.

[16] 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.

[17] Robert Bajko and Deborah I Fels. 2016. Prevalence of Mobile Phone Interaction in Workplace Meetings. In *International Conference on HCI in Business, Government, and Organizations*. Springer, 273–280.

[18] Mark Baker, Jon Coaffee, and Graeme Sherriff. 2007. Achieving successful participation in the new UK spatial planning system. *Planning, Practice & Research* 22, 1 (2007), 79–93.

[19] Satanjeev Banerjee and Alexander I Rudnicky. 2006. Smartnotes: Implicit labeling of meeting data through user note-taking and browsing. In *Proceedings of the 2006 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology: companion volume: demonstrations*. Association for Computational Linguistics, 261–264.

[20] William E Bennett, Stephen J Boies, Anthony R Davies, Karl-Friedrich Etzold, and Todd K Rodgers. 1991. Optical stylus and passive digitizing tablet data input system. US Patent 5,051,736.

[21] Tony Bergstrom and Karrie Karahalios. 2007. Conversation Clock: Visualizing audio patterns in co-located groups. In *System Sciences, 2007. HICSS 2007. 40th Annual Hawaii International Conference on*. IEEE, 78–78.

[22] Tony Bergstrom and Karrie Karahalios. 2009. Vote and be heard: Adding back-channel signals to social mirrors. In *IFIP Conference on Human-Computer Interaction*. Springer, 546–559.

[23] David M Blei and John D Lafferty. 2009. Visualizing topics with multi-word expressions. *arXiv preprint arXiv:0907.1013* (2009).

[24] Adrien Bougouin, Florian Boudin, and Béatrice Daille. 2013. Topicrank: Graph-based topic ranking for keyphrase extraction. In *International Joint Conference on Natural Language Processing (IJCNLP)*. 543–551.

[25] Shelley Boulianne, Kristjana Loptson, and David Kahane. 2018. Citizen Panels and Opinion Polls: Convergence and Divergence in Policy Preferences. *Journal of Public Deliberation* 14, 1 (2018), 4.

[26] Tony Bovaird. 2007. Beyond engagement and participation: User and community coproduction of public services. *Public administration review* 67, 5 (2007), 846–860.

[27] Daren C Brabham. 2009. Crowdsourcing the public participation process for planning projects. *Planning Theory* 8, 3 (2009), 242–262.

[28] 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*.

[29] Frank M Bryan. 2010. *Real democracy: The New England town meeting and how it works*. University of Chicago Press.

[30] John M Bryson, Kathryn S Quick, Carissa Schively Slotterback, and Barbara C Crosby. 2013. Designing public participation processes. *Public administration review* 73, 1 (2013), 23–34.

[31] Philip Burnard. 1991. A method of analysing interview transcripts in qualitative research. *Nurse education today* 11, 6 (1991), 461–466.

[32] Mark Button and Kevin Mattson. 1999. Deliberative democracy in practice: Challenges and prospects for civic deliberation. *Polity* 31, 4 (1999), 609–637.

[33] Jean Carletta, Simone Ashby, Sebastien Bourban, Mike Flynn, Mael Guillemot, Thomas Hain, Jaroslav Kadlec, Vasilis Karaiskos, Wessel Kraaij, Melissa Kronenthal, et al. 2005. The AMI meeting corpus: A pre-announcement. In *International workshop on machine learning for multimodal interaction*. Springer, 28–39.

[34] Ronnie Chaiken, Bob Jenkins, Per-Åke Larson, Bill Ramsey, Darren Shakib, Simon Weaver, and Jingren Zhou. 2008. SCOPE: easy and efficient parallel processing of massive data sets. *Proceedings of the VLDB Endowment* 1, 2 (2008), 1265–1276.

[35] Patrick Chiu, John Boreczky, Andreas Girgensohn, and Don Kimber. 2001. LiteMinutes: an Internet-based system for multimedia meeting minutes. In *Proceedings of the 10th international conference on World Wide Web*. ACM, 140–149.

[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] Stephen Coleman and John Gotze. 2001. *Bowling together: Online public engagement in policy deliberation*. Hansard Society London.

[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] Dora L Costa and Matthew E Kahn. 2003. Civic engagement and community heterogeneity: An economist's perspective. *Perspectives on politics* 1, 1 (2003), 103–111.

[43] Dora L Costa and Matthew E Kahn. 2003. Understanding the American decline in social capital, 1952–1998. *Kyklos* 56, 1 (2003), 17–46.

[44] Lisa A D'Ambrosio and Alea C Mehler. 2014. Three things policymakers should know about technology and older adults. *Public Policy & Aging Report* 24, 1 (2014), 10–13.

[45] Rodrigo Davies. 2014. Civic crowdfunding: participatory communities, entrepreneurs and the political economy of place. *Entrepreneurs and the Political Economy of Place (May 9, 2014)* (2014).

[46] Richard C Davis, Jason A Brotherton, James A Landay, Morgan N Price, and Bill N Schilit. 1998. NotePals: Lightweight Note Taking by the Group, for the Group. *University of California, Berkeley, Computer Science Division* 8 (1998).

[47] Francis J Di Vesta and G Susan Gray. 1972. Listening and note taking. *Journal of educational psychology* 63, 1 (1972), 8.

[48] Jessa Dickinson, Mark Díaz, Christopher A Le Dantec, and Sheena Erete. 2019. " The cavalry ain't coming in to save us" Supporting Capacities and Relationships through Civic Tech. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–21.

[49] Jessa Dickinson, Sheena Erete, Mark Diaz, and Denise Linn Riedl. 2018. Inclusion of Underserved Residents in City Technology Planning. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–6.

[50] Catherine D'Ignazio and Lauren F Klein. 2020. *Data feminism*. MIT Press.

[51] 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.

[52] Amir Ehsaei, Thomas Sweet, Raphael Garcia, Laura Adleman, and Jean M Walsh. 2015. Successful Public Outreach Programs for Green Infrastructure Projects. In *International Low Impact Development Conference 2015: LID: It Works in All Climates and Soils*. 74–92.

[53] Sheena Erete and Jennifer O Burrell. 2017. Empowered participation: How citizens use technology in local governance. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 2307–2319.

[54] Archon Fung. 2006. Varieties of participation in complex governance. *Public administration review* 66 (2006), 66–75.

[55] Nikhil Garg, Benoit Favre, Korbinian Reidhammer, and Dilek Hakkani-Tür. 2009. Clusterrank: a graph based method for meeting summarization. In *Tenth Annual Conference of the International Speech Communication Association*.

[56] John Gastil. 2008. *Political communication and deliberation*. Sage.

[57] Dan Gillick, Korbinian Riedhammer, Benoit Favre, and Dilek Hakkani-Tur. 2009. A global optimization framework for meeting summarization. In *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 4769–4772.

[58] Lisa M Given. 2008. *The Sage encyclopedia of qualitative research methods*. Sage publications.

[59] Leo A Goodman. 1961. Snowball sampling. *The annals of mathematical statistics* (1961), 148–170.

[60] Eric Gordon and Paul Mihailidis. 2016. *Civic Media: Technology, Design, Practice*. MIT Press.

[61] Eric Gordon, Steven Schirra, and Justin Hollander. 2011. Immersive planning: a conceptual model for designing public participation with new technologies. *Environment and Planning B: Planning and Design* 38, 3 (2011), 505–519.

[62] Kenneth L. Hacker and Jan Van Dijk. 2001. *Digital Democracy: Issues of Theory and Practice*. Sage Publications, Inc., USA.

[63] 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* (Seoul, Republic of Korea) *(CHI '15)*. ACM, New York, NY, USA, 2833–2842.

[64] Kurtis Heimerl, Brian Gawalt, Kuang Chen, Tapan Parikh, and Björn Hartmann. 2012. CommunitySourcing: engaging local crowds to perform expert work via physical kiosks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1539–1548.

[65] Clyde Freeman Herreid. 2006. " Clicker" Cases. *Journal of College Science Teaching* 36, 2 (2006), 43.

[66] Julia Hirschberg and Christopher D Manning. 2015. Advances in natural language processing. *Science* 349, 6245 (2015), 261–266.

[67] Seth Hunter, Pattie Maes, Stacey Scott, and Henry Kaufman. 2011. MemTable: an integrated system for capture and recall of shared histories in group workspaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 3305–3314.

[68] Jane Im, Amy X Zhang, Christopher J Schilling, and David Karger. 2018. Deliberation and Resolution on Wikipedia: A Case Study of Requests for Comments. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 74.

[69] Judith E Innes and David E Booher. 1999. Consensus building and complex adaptive systems: A framework for evaluating collaborative planning. *Journal of the American planning association* 65, 4 (1999), 412–423.

[70] Judith E Innes and David E Booher. 2004. Reframing public participation: strategies for the 21st century. *Planning theory & practice* 5, 4 (2004), 419–436.

[71] Renee A Irvin and John Stansbury. 2004. Citizen participation in decision making: is it worth the effort? *Public administration review* 64, 1 (2004), 55–65.

[72] James Derek Jacoby. 2010. *Collaborative and automatic annotations to improve the utility of recorded meetings*. Ph.D. Dissertation.

[73] Sara J Jones, Jason Crandall, Jane S Vogler, and Daniel H Robinson. 2013. Classroom response systems facilitate student accountability, readiness, and learning. *Journal of Educational Computing Research* 49, 2 (2013), 155–171.

[74] Vaiva Kalnikaitė, Patrick Ehlen, and Steve Whittaker. 2012. Markup as you talk: establishing effective memory cues while still contributing to a meeting. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*. ACM, 349–358.

[75] Matthew Kam, Jingtao Wang, Alastair Iles, Eric Tse, Jane Chiu, Daniel Glaser, Orna Tarshish, and John Canny. 2005. Livenotes: a system for cooperative and augmented note-taking in lectures. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 531–540.

[76] Christopher F Karpowitz and Jane Mansbridge. 2005. Disagreement and consensus: The need for dynamic updating in public deliberation. *Journal of Public deliberation* 1, 1 (2005).

[77] Robin H Kay and Ann LeSage. 2009. Examining the benefits and challenges of using audience response systems: A review of the literature. *Computers & Education* 53, 3 (2009), 819–827.

[78] Shreeharsh Kelkar, Ajita John, and Doree Duncan Seligmann. 2010. Some observations on the live collaborative tagging of audio conferences in the enterprise. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 995–998.

[79] Christopher Keller, N Finkelstein, K Perkins, S Pollock, C Turpen, and M Dubson. 2007. Research-based Practices For Effective Clicker Use. In *AIP Conference Proceedings*, Vol. 951. AIP, 128–131.

[80] Catherine Keske and Steve Smutko. 2010. Consulting communities: using audience response system (ARS) technology to assess community preferences for sustainable recreation and tourism development. *Journal of Sustainable Tourism* 18, 8 (2010), 951–970.

[81] Nam Wook Kim, Jonghyuk Jung, Eun-Young Ko, Songyi Han, Chang Won Lee, Juho Kim, and Jihee Kim. 2016. Budgetmap: Engaging taxpayers in the issue-driven classification of a government budget. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*. 1028–1039.

[82] Soomin Kim, Jinsu Eun, Changhoon Oh, Bongwon Suh, and Joonhwan Lee. 2020. Bot in the Bunch: Facilitating Group Chat Discussion by Improving Efficiency and Participation with a Chatbot. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.

[83] René F Kizilcec, Chris Piech, and Emily Schneider. 2013. Deconstructing disengagement: analyzing learner subpopulations in massive open online courses. In *Proceedings of the third international conference on learning analytics and knowledge*. 170–179.

[84] Klaus Krippendorff. 2011. Computing Krippendorff's alpha-reliability. (2011).

[85] Kostadin Kushlev, Jason Proulx, and Elizabeth W Dunn. 2016. " Silence Your Phones" Smartphone Notifications Increase Inattention and Hyperactivity Symptoms. In *Proceedings of the 2016 CHI conference on human factors in computing systems*. 1011–1020.

[86] Sung-Chul Lee, Jaeyoon Song, Eun-Young Ko, Seongho Park, Jihee Kim, and Juho Kim. 2020. SolutionChat: Real-time Moderator Support for Chat-based Structured Discussion. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–12.

[87] Peter Levine, Archon Fung, and John Gastil. 2005. Future directions for public deliberation. *Journal of Public Deliberation* 1, 1 (2005).

[88] Rensis Likert. 1932. A technique for the measurement of attitudes. *Archives of psychology* (1932).

[89] Chin-Yew Lin. 2004. ROUGE: A Package for Automatic Evaluation of Summaries. In *Text Summarization Branches Out*. Association for Computational Linguistics, Barcelona, Spain, 74–81. https://www.aclweb.org/anthology/W04-1013

[90] Carolyn Lukensmeyer, Susanna Haas Lyons, and America Speaks. [n.d.]. 21 st Century Town Meeting®.

[91] Carolyn J Lukensmeyer and Steve Brigham. 2002. Taking democracy to scale: Creating a town hall meeting for the twenty-first century. *National Civic Review* 91, 4 (2002), 351–366.

[92] 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.

[93] 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. (2018).

[94] Narges Mahyar, Diana V Nguyen, Maggie Chan, Jiayi Zheng, and Steven Dow. 2019. The Civic Data Deluge: Understanding the Challenges of Analyzing Large-Scale Community Input. (2019), 11 pages pages. ACM Designing Interactive Systems (DIS) (to appear).

[95] Jane Mansbridge, Janette Hartz-Karp, Matthew Amengual, and John Gastil. 2006. Norms of deliberation: An inductive study. (2006).

[96] Jennifer Manuel, Geoff Vigar, Tom Bartindale, and Rob Comber. 2017. Participatory Media: Creating Spaces for Storytelling in Neighbourhood Planning. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 1688–1701.

[97] Moira McGregor and John C Tang. 2017. More to Meetings: Challenges in Using Speech-Based Technology to Support Meetings.. In *CSCW*. 2208–2220.

[98] Albert Mehrabian. 2017. *Nonverbal communication*. Routledge.

[99] Rada Mihalcea and Paul Tarau. 2004. Textrank: Bringing order into text. In *Proceedings of the 2004 conference on empirical methods in natural language processing*. 404–411.

[100] Aditi Misra, Aaron Gooze, Kari Watkins, Mariam Asad, and Christopher A Le Dantec. 2014. Crowdsourcing and its application to transportation data collection and management. *Transportation Research Record* 2414, 1 (2014), 1–8.

[101] Stefanie Mollborn and Angel Hoekstra. 2010. "A meeting of minds" using clickers for critical thinking and discussion in large sociology classes. *Teaching sociology* 38, 1 (2010), 18–27.

[102] Judith Morse, Margaret Ruggieri, and Karen Whelan-Berry. 2010. Clicking Our Way to Class Discussion. *American Journal of Business Education* 3, 3 (2010), 99–108.

[103] Mary Ann Murphy. 2009. Promotion of Engaged Democracy and Community Partnerships Through Audience Response System Technology. (2009).

[104] Gabriel Murray, Steve Renals, Jean Carletta, and Johanna Moore. 2006. Incorporating speaker and discourse features into speech summarization. In *Proceedings of the main conference on Human Language Technology Conference of the North American Chapter of the Association of Computational Linguistics*. Association for Computational Linguistics, 367–374.

[105] Mukesh Nathan, Mercan Topkara, Jennifer Lai, Shimei Pan, Steven Wood, Jeff Boston, and Loren Terveen. 2012. In case you missed it: benefits of attendee-shared annotations for non-attendees of remote meetings. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*. ACM, 339–348.

[106] Rand B Nickerson. 1993. Real-time wireless audience response system. US Patent 5,226,177.

[107] Patrick Olivier and Peter Wright. 2015. Digital Civics: Taking a Local Turn. *interactions* 22, 4 (June 2015), 61–63.

[108] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. *The pagerank citation ranking: Bringing order to the web*. Technical Report. Stanford InfoLab.

[109] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Philadelphia, Pennsylvania, USA, 311–318. https://doi.org/10.3115/1073083.1073135

[110] Annie Piolat, Thierry Olive, and Ronald T Kellogg. 2005. Cognitive effort during note taking. *Applied cognitive psychology* 19, 3 (2005), 291–312.

[111] Robert D Putnam et al. 2000. *Bowling alone: The collapse and revival of American community*. Simon and schuster.

[112] Lawrence R Rabiner, Biing-Hwang Juang, and Janet C Rutledge. 1993. *Fundamentals of speech recognition*. Vol. 14. PTR Prentice Hall Englewood Cliffs.

[113] Wendy M Rahn and John E Transue. 1998. Social trust and value change: The decline of social capital in American youth, 1976–1995. *Political Psychology* 19, 3 (1998), 545–565.

[114] Lisa Marie Rhody. 2016. Why I dig: Feminist approaches to text analysis. *Debates in the digital humanities* (2016).

[115] Evan F Risko, Tom Foulsham, Shane Dawson, and Alan Kingstone. 2013. The collaborative lecture annotation system (CLAS): A new TOOL for distributed learning. *IEEE Transactions on Learning Technologies* 6, 1 (2013), 4–13.

[116] Stephen E Robertson, Steve Walker, Susan Jones, Micheline M Hancock-Beaulieu, Mike Gatford, et al. 1995. Okapi at TREC-3. *Nist Special Publication Sp* 109 (1995), 109.

[117] Gene Rowe and Lynn J Frewer. 2005. A typology of public engagement mechanisms. *Science, Technology, & Human Values* 30, 2 (2005), 251–290.

[118] Mariana Salgado and Michail Galanakis. 2014. "... so what?" limitations of participatory design on decision-making in urban planning. In *Proceedings of the 13th Participatory Design Conference: Short Papers, Industry Cases, Workshop Descriptions, Doctoral Consortium papers, and Keynote abstracts-Volume 2*. 5–8.

[119] Lynn M Sanders. 1997. Against deliberation. *Political theory* 25, 3 (1997), 347–376.

[120] Ethan Seltzer and Dillon Mahmoudi. 2013. Citizen participation, open innovation, and crowdsourcing: Challenges and opportunities for planning. *Journal of Planning Literature* 28, 1 (2013), 3–18.

[121] Guokan Shang, Wensi Ding, Zekun Zhang, Antoine Jean-Pierre Tixier, Polykarpos Meladianos, Michalis Vazirgiannis, and Jean-Pierre Lorré. 2018. Unsupervised Abstractive Meeting Summarization with Multi-Sentence Compression and Budgeted Submodular Maximization. *arXiv preprint arXiv:1805.05271* (2018).

[122] Yang Shi, Yang Wang, and John Chen. 2017. IdeaWall : Improving Creative Collaboration through Combinatorial Visual Stimuli. *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing - CSCW '17* (2017), 594–603.

[123] Huan Sun, Alex Morales, and Xifeng Yan. 2013. Synthetic review spamming and defense. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 1088–1096.

[124] Oguzhan Tas and Farzad Kiyani. 2007. A SURVEY AUTOMATIC TEXT SUMMARIZATION. *PressAcademia Procedia* 5, 1 (2007), 205–213.

[125] Nick Taylor, Justin Marshall, Alicia Blum-Ross, John Mills, Jon Rogers, Paul Egglestone, David M Frohlich, Peter Wright, and Patrick Olivier. 2012. Empowering communities with situated voting devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1361–1370.

[126] Lars Hasselblad Torres. 2007. Citizen sourcing in the public interest. *Knowledge Management for Development Journal* 3, 1 (2007), 134–145.

[127] Karen Tracy and Margaret Durfy. 2007. Speaking out in public: Citizen participation in contentious school board meetings. *Discourse & Communication* 1, 2 (2007), 223–249.

[128] Gokhan Tur, Andreas Stolcke, Lynn Voss, Stanley Peters, Dilek Hakkani-Tur, John Dowding, Benoit Favre, Raquel Fernández, Matthew Frampton, Mike Frandsen, et al. 2010. The CALO meeting assistant system. *IEEE Transactions on Audio, Speech, and Language Processing* 18, 6 (2010), 1601–1611.

[129] Eleftheria Vaportzis, Maria Giatsi Clausen, and Alan J Gow. 2017. Older adults perceptions of technology and barriers to interacting with tablet computers: a focus group study. *Frontiers in psychology* 8 (2017), 1687.

[130] 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.

[131] Eric von Hippel. 2005. Democratizing innovation: Users take center stage.

[132] Greg Walsh and Eric Wronsky. 2019. AI+ Co-Design: Developing a Novel Computer-supported Approach to Inclusive Design. In *Conference Companion Publication of the 2019 on Computer Supported Cooperative Work and Social Computing*. 408–412.

[133] Tsung-Hsien Wen, Milica Gasic, Nikola Mrksic, Pei-Hao Su, David Vandyke, and Steve Young. 2015. Semantically conditioned lstm-based natural language generation for spoken dialogue systems. *arXiv preprint arXiv:1508.01745* (2015).

[134] Christopher Whitehead and Lydia Ray. 2010. Using the iclicker classroom response system to enhance student involvement and learning. *Journal of Education, Informatics and Cybernetics* 2, 1 (2010), 18–23.

[135] Amy X Zhang and Justin Cranshaw. 2018. Making sense of group chat through collaborative tagging and summarization. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 196.

[136] Amy X Zhang, Bryan Culbertson, and Praveen Paritosh. 2017. Characterizing online discussion using coarse discourse sequences. In *Proceedings of the Eleventh International Conference on Web and Social Media. AAAI Press*.

[137] Ming Zhong, Pengfei Liu, Danqing Wang, Xipeng Qiu, and Xuanjing Huang. 2019. Searching for Effective Neural Extractive Summarization: What Works and What's Next. *arXiv preprint arXiv:1907.03491* (2019).

A APPENDIX

A.1 Summary Evaluation

To summarize augmented transcripts, we wanted a robust unsupervised summarization algorithm. To that end, we improved TextRank [99] to create an attendee-feedback-weighted summarization algorithm. To test the effectiveness of this extractive summarization algorithm, we took three different transcripts obtained from the following meetings[4]:

(1) A meeting discussion about creating a common room for graduate students in the department.
(2) A meeting organized by public school officials discussing challenges to building renovations.
(3) A meeting held on a college campus discussing the rise in racist activities.

We evaluated these transcriptions with 4 different human annotators, who were specifically tasked with identifying sentences that best summarized the discussion transcripts. The annotators were graduate students working in the area of natural language processing. They were familiar with the literature and practices of text summarization tasks. We treated these human-annotated summaries as references and against each of those reference summaries, we evaluated a ROUGE [89] score metric for the auto-generated summary. ROUGE is a widely used metric for evaluating computer-generated summaries. It measures how much of the reference (human-generated) summary is our algorithm "capturing". We ideally wanted to strike a good balance between conciseness (precision) and the amount of information captured (recall) by our auto-generated summary and therefore, we computed the ROUGE F1-scores which in turn measure both, precision and recall. We evaluated ROUGE-1 F1 and ROUGE-2 F1 scores where ROUGE-N denotes N-gram comparisons of auto-generated and reference summaries. For example, ROUGE-1 refers to the overlap of unigrams between our auto-generated summary and the reference summary. Similarly, ROUGE-2 would denote the bigram comparisons. Upon evaluating our auto-generated summaries for the three selected transcripts against their respective reference summaries, we noticed nearly-consistency ROUGE scores. Our full results are summarized in Table 2. These results suggest that our algorithm produced summaries of consistent quality across different meeting transcripts.

Table 2. `ROUGE-1` and `ROUGE-2` F1 scores of generated summaries (T_n) against reference summaries produced by human annotators (A_n).

Transcriptions	Rouge-1 F1				Rouge-2 F1			
	A_1	A_2	A_3	A_4	A_1	A_2	A_3	A_4
T_1 (Common room discussion)	51.29	45.83	48.12	42.41	34.56	33.23	24.64	30.98
T_2 (School building renovations)	48.23	12.34	44.43	49.19	28.21	31.39	29.79	23.98
T_3 (Activities on college campus)	68.21	70.61	52.28	55.09	33.94	35.90	27.29	21.93

Our summarization method allowed us to generate short unsupervised summaries of meeting transcripts by modeling those transcripts as a graph with sentences being the nodes. A similarity function is used to build edges between those nodes. BM25 [116] is a ranking utility function widely used to measure similarity in state-of-the-art methods for information extraction tasks. To measure the effectiveness of using BM25 as a similarity measure in our case and to further assess the robustness of our summarization system on a widely used standardized dataset, we evaluated the effect of changing the similarity metric used in TextRank through some ablation studies. We carried out these experiments on the standard AMI Meeting Corpus[5], which is a widely

---

[4]All of which are presently analyzed in the CommunityClick interface.
[5]http://groups.inf.ed.ac.uk/ami/corpus/

used meeting-summarization benchmark. We used the traditional test set of 20 meetings, each associated with a human-annotated summary of approximately 290 tokens. We compared the vanilla similarity metric [99] against the BM25 similarity metric and evaluated the ROUGE scores for the AMI test-set reference summaries. Results for these experiments are summarized in Table 3. We observe that our approach resulted in a 3-point gain in recall (amount of information captured) and a 2-point gain in precision over the vanilla similarity metric used in the TextRank algorithm on the AMI-meeting corpus.

Table 3. Macro averaged results for different similarity functions within TextRank.

Similarity Function (f)	AMI ROUGE-1		
	P	**R**	**F-1**
TextRank (with BM25)	**42.98**	**35.75**	**37.39**
TextRank (with cosine-similarity)	41.21	32.66	36.54
TextRank (vanilla)	40.11	33.48	36.27

As we evaluated the performance on a domain-specific dataset (AMI meeting corpus), we argue that these gains are consistent across our data as well. Furthermore, we found that our results were in-line with recent state-of-the-art methods [121, 137] on unsupervised extractive summarization.

A.2 Automatic Speech Recognition Evaluation

To assess how accurately our Automated Speech Recognition captures the meeting audio to generate meeting transcripts, we carried out a tests to evaluate the Bilingual Evaluation Understudy Score (BLEU) [109] on the same transcripts described in the previous section. A BLEU score is a metric for evaluating a generated sentence to a reference sentence and is often used in machine translation tasks. We adopted it to evaluate our transcripts by counting matching n-grams in the generated text to matching n-grams in the reference text. A score 1.0 indicates a perfect match and a lowest score of 0.0 would indicate a complete mismatch. Our reference transcripts were generated by a single human annotator who manually transcribed the meeting audio. We consistently observed high BLEU scores, indicating that tokens appearing in the reference transcript are also present in the generated transcript, thereby implying that our ASR system was indeed reliable. Our full results are summarized in Table 4.

Table 4. Evaluation of generated transcripts against reference transcripts that were manually generated by humans. BLEU-n scores here effectively indicate n-gram overlaps between sentences from reference transcript and generated transcript.

	T_1 (16 mins)	T_2 (55 mins)	T_3 (75 mins)
BLEU - 1	0.9028	0.9024	0.9046
BLEU - 2	0.8811	0.8753	0.8733
BLEU - 3	0.8152	0.8232	0.8381
BLEU - 4	0.7812	0.7910	0.7604