

MeetingVis: Visual Narratives to Assist in Recalling Meeting Context and Content

Yang Shi, Chris Bryan, *Student Member, IEEE*, Sridatt Bhamidipati, Ying Zhao, Yaoxue Zhang, and Kwan-Liu Ma, *Fellow, IEEE*

Abstract—In team-based workplaces, reviewing and reflecting on the content from a previously held meeting can lead to better planning and preparation. However, ineffective meeting summaries can impair this process, especially when participants have difficulty remembering what was said and what its context was. To assist with this process, we introduce MeetingVis, a visual narrative-based approach to meeting summarization. MeetingVis is composed of two primary components: (1) a data pipeline that processes the spoken audio from a group discussion, and (2) a visual-based interface that efficiently displays the summarized content. To design MeetingVis, we create a taxonomy of relevant meeting data points, identifying salient elements to promote recall and reflection. These are mapped to an augmented storyline visualization, which combines the display of participant activities, topic evolutions, and task assignments. For evaluation, we conduct a qualitative user study with five groups. Feedback from the study indicates that MeetingVis effectively triggers the recall of subtle details from prior meetings: all study participants were able to remember new details, points, and tasks compared to an unaided, memory-only baseline. This visual-based approaches can also potentially enhance the productivity of both individuals and the whole team.

Index Terms—Design study, information visualization, meeting summarization, natural language processing, visual narrative, voice recognition

1 INTRODUCTION

IN business, collaborative, and team-based environments, meetings are a common occurrence. As the number and duration of meetings is steadily increasing [1], an essential skill for today's workers is being able to recall and reflect the information from previously attended discussions, to better prepare for both upcoming meetings and future work [2].

One way to assist recall is by presenting key points and takeaways in the form of a meeting summary [3]. Automatic, text-based summarization provides a succinct method to describe a meeting's primary focuses [4], but it fails at depicting deeper, multifaceted contexts such as: "What are the contributions of each participant? How important is each topic?" This limits its overall utility.

Meeting browsers provide another way for reviewing meetings, as they combine recorded audio and video in a playback interface [5]. This interface enables a more thorough review, where the complete, fully-recorded meeting

content is available. Nevertheless, using a meeting browser has a drawback, as it can be difficult to perceive the overall discussion at a glance and requires searching across the whole history to find individual conversation points.

In contrast, visual-based meeting summaries present a more efficient and effective approach. Prior work in the area of meeting visualization has used verbal statements to frame a meeting as a set of thematic progressions (e.g., [6], [7]). Key discussion areas can be quickly reviewed to identify the relevant topics of interest as the meeting evolves over time. These efforts primarily focus visualization on a single dimension of meetings by constructing a recognizable picture of *what* is discussed.

However, in considering the narrative structure of an event (e.g., *who? when? what? how?*), we see that individual meetings are composed of several primary elements: specific individuals who interact with each other, a time period during which this interaction occurs, a set of topical points which frame the discussion, etc. There is a lack of systematic study on what meeting elements can be represented using visual cues, how these cues can be visually combined, and what impact they have on memory retrieval.

Motivated by this shortcoming, we have developed *MeetingVis*: a lightweight, intuitive, and informative visual interface to facilitate the review and recall of meetings. MeetingVis is based on the idea that the narrative-based visualization of meeting elements will help a participant to remember both the content and context of a prior-attended meeting. To develop MeetingVis, we first derive a taxonomy of meeting data elements, which are categorized according to their communication channel and analysis degree. From this taxonomy,

- Y. Shi is with the College of Design and Innovation, Tongji University, Shanghai 200092, China. E-mail: yangshi.idx@tongji.edu.cn.
- C. Bryan, S. Bhamidipati, and K.-L. Ma are with the University of California, Davis, CA 95616. E-mail: {cybryan, sbhamidi, klma}@ucdavis.edu.
- Y. Zhao is with the Central South University, Changsha, Hunan 410083, China. E-mail: zhaoying@csu.edu.cn.
- Y. Zhang is with the Central South University, Changsha, Hunan 410083, China, and with the Tsinghua University, Beijing 100084, China. E-mail: zyx@csu.edu.cn.

Manuscript received 22 Sept. 2017; revised 5 Jan. 2018; accepted 2 Feb. 2018.

Date of publication 20 Mar. 2018; date of current version 27 Apr. 2018.

Recommended for acceptance by S. Bruckner, K. Koyamada, and B. Lee.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.

Digital Object Identifier no. 10.1109/TVCG.2018.2816203

we then identify a set of five primary narrative elements to summarize a meeting: *participants*, *topics*, *tasks*, *timeline*, and *interaction*, which map to the major elements of a story. Third, we encode and display the five selected meeting elements by augmenting the base storyline visualization technique [8] with additional visual features. As a result, MeetingVis succinctly summarizes meetings by encoding participant activities, topic evolutions, and task assignments. The backend data pipeline of MeetingVis leverages Automatic Speech Recognition (ASR) and Natural Language Processing (NLP) techniques, including speech recognition, speaker recognition, topic modeling and information extraction.

To evaluate the effectiveness of MeetingVis in assisting recall, we conduct an exploratory user study with five groups of participants. Our study results show that MeetingVis is capable of triggering the memory of both the “specific images” (the content) as well as the “bigger pictures” (the context) from discussions. Visualizing meeting elements in a narrative manner also facilitates reflection by individuals on their participation and overall collaboration. We conclude this paper with a discussion on how visualization approaches can increase productivity for upcoming meetings and provide insight when planning future work.

The contributions of our work include:

- A *visual summary system* that displays the narrative elements to summarize a meeting.
- A *systematic analysis*, based on a taxonomy of meeting elements, that determines a set of primary narrative-based meeting elements for promoting recall.
- A *data pipeline* that captures the audio of a meeting and automatically processes it into the necessary meeting elements for display.
- A *qualitative experiment* that evaluates the effectiveness of the system in supporting recall and reflection, both of meeting content and context.

2 RELATED WORK

MeetingVis builds on prior research for the review and recall of meetings: augmented note-taking techniques, meeting information retrieval methods, and multi-party conversation visualizations.

2.1 Augmented Note-Taking Techniques

Note-taking is the practice of recording information, such as points of discussion during a meeting or oration from a lecturer. Taking notes has been shown to improve recall and learning outcomes [9], [10]. Based on analysis of the note-taking process [11], [12], various augmented note-taking techniques have been employed to help recall meetings. Kalnikaitė et al. [13] designed markup tools where users highlight and/or mark personally relevant notes from transcripts generated by ASR. Nathan et al. [14] integrated annotations created by meeting attendees into recordings and evaluated whether or not they assisted non-attendees in information retrieval. Willett et al. [15] explored how visualization techniques can support reflection, organization, and collaboration around digital notes, such as with a visual dashboard containing notes, text, and metadata.

In general, these works exploit manual note-taking techniques to highlight specific but important points of information.

In contrast, MeetingVis combines visualization and HCI design with a data pipeline architecture to automatically create a meeting summary. It can be used as a complementary, automatic summarization technique, organizing and displaying elements based on their narrative structure and providing a more thematic and contextual overview.

2.2 Meeting Information Retrieval Methods

Various information and multimedia channels can be collected from a meeting. The simplest approach is filming the discussion or recording the spoken audio, which can be loaded into meeting browsers for review and information retrieval [5]. Kern et al. [16] suggested wearable sensor systems for meeting participants. These generate personal, contextual data points, such as speaker segments, body movements, and the changing postures of participants. Jaimes et al. [17] improved meeting video retrieval by graphically illustrating memory cues such as room layout, participant names, and sitting positions. Yu et al. [18] developed MeetingAssistant, featuring context-aware browsing in which meeting information is selected and shown based on users’ situational contexts. Girgensohn et al. [19] created HyperMeeting, which synchronizes and indexes videos from videoconferencing attendees by speaker and topic. Hunter et al. [20] developed MemTable, which supports media capture through channels such as drawing on physical paper, text input, semantic searching, and temporal browsing.

Most of these approaches focus on augmenting direct media representations (e.g., videos, audios) with additional context-enabling cues, rather than abstractly encoding data points with visualization. However, our work is similar in that MeetingVis uses cues to help trigger recall. Both approaches are based in cognitive theory to help participants reconstruct the meeting’s flow, overall experience, and specific items that were discussed.

2.3 Multi-Party Conversation Visualizations

Conversation-based visualizations present discussions between two or more participants. Usually, the focus includes showing the topical aspects of the meeting, allowing for thematic review and analysis of the discussed content. ChAT [6] uses a matrix view to track meeting characteristics, including who talked to whom, when it happened, and the entities that were discussed. Meeting Adjourned [21] presents topical discussions as a series on a linear timeline with a corresponding list of action items. SUVI [7] visualizes meetings in two ways: a storyboard depicts the chronological sequence and a newspaper layout illustrates topics in relation to their relevance. ConToVi [22] uses a visual sedimentation metaphor to display participants’ behavior patterns and conversational flows in a radial manner. IdeaWall [23] dynamically groups the content of conversations based on the thematic structure. NEREx [24] presents the relations of named-entity pairs in conversation transcripts to help the exploratory analysis of multi-party conversations.

Similarly, MeetingVis visualizes the discussed topical content of meetings. However, by taking a narrative-based approach, MeetingVis includes significant context information, such as how participant communication, interaction, and involvement changes over time.

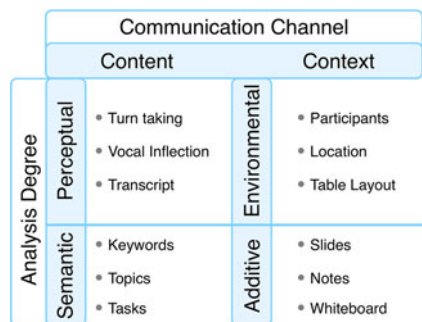


Fig. 1. Taxonomy of meeting data types and example data elements. Each element is placed in one of four classes based on two classification dimensions: its *communication channel* and *analysis degree*.

3 CONSTRUCTION OF MEETING-STORY MAPPING

Our goal is to design a narrative-based visual summary for the purpose of improving meeting recall. To this end, we analyze salient meeting elements, identify elements based on their narrative structures and potential to invoke recall, and map these elements onto augmented storyline visualizations. The resulting meeting-to-story mapping is used to drive the design of a visual summary system.

3.1 Identifying and Categorizing Meeting Elements

We begin by identifying meeting elements that can be used to depict or review a meeting. To do this, we reviewed the data and metadata points discussed in prior research (e.g., [18], [25]) and by existing applications for meeting review (e.g., Minute [26], MeetingKing [27]).

These elements are organized into a taxonomy of four classes shown in Fig. 1. There are two main dimensions for determining the classification of a meeting element: its *communication channel* and *analysis degree*. Each dimension has two options that an element can fall into.

Communication channel divides elements based on whether they are *content*- or *context*-based. Content data is recordable and/or extractable from the meeting itself. This includes words spoken by attendees and the thematic topics that are discussed. Alternatively, contextual meeting elements act as cues by being present within the meeting environment itself. This includes the participants present and additional resources that are brought to the meeting, such as presentation slides or meeting notes.

The second classification dimension is analysis degree. It refers to how meeting elements may be cognitively divided. *Perceptual* and *environmental* describes low-level data schemas, such as speaker vocal inflections and the location of the meeting. *Semantic* and *additive* describes data from higher-level syntheses such as keyword extraction and the analysis of additional, contextualizing meeting resources.

3.2 Selecting Elements to Show Narrative Structure

Showing all possible meeting elements at once would be inefficient and is likely unnecessary for triggering recall. In our case, we attempt to identify a subset of salient elements that will simultaneously provoke specific memories while providing an efficient summary of the meeting for review purposes.

To help determine the set of meeting elements to use, we first conducted a sample survey with ten professionals who

TABLE 1
Five Meeting Elements Chosen to Invoke Recall

Meeting Element	Narrative Structure Descriptor	Description
Participants	Who	The people taking part in the discussion
Topics	What	The discrete subjects discussed
Tasks	What	The action items assigned for future work
Timeline	When	The temporal context of other meeting elements
Interaction	How	The level and type of discussion between participants

regularly attend meetings (five scientific researchers, two professors, two software engineers, and a business product manager). Each was asked questions about note-taking including: (1) What is the meeting information that participants commonly record for later recall? (2) What is the meeting information that participants want but cannot acquire due to constraints (e.g., time limits, analysis requirements)? Based on their feedback, the three most popular meeting elements were defined as: (1) *participants*, the people taking part in the meeting, (2) *topics*, the discrete, thematic subjects discussed during the meeting, and (3) *tasks*, the action items assigned for future work.

We next considered recall of meeting elements from a cognitive perspective. At this level, recall can be defined as information retrieval from episodic memory [3]. Retrieving such memories requires that the narrative structure (e.g., the *who*, *when*, *what*, *how* knowledge of the occurrence), which compose a concrete, underlying topology of episodic memory, be specified explicitly [28], [29], [30]. The three above-defined meeting elements already answer the *who* and *what* of the meeting. To fulfill the *when* and *how*, we specify two additional meeting elements: (4) meeting data should be arrayed along a *timeline* to provide temporal context, and (5) the *interaction* between participants addresses the context as the discussion evolves. These five elements, noted in Table 1, are chosen to summarize a meeting and invoke recall.

3.3 Meeting-Story Mapping

The next step is to structure the five defined meeting elements to visually summarize a meeting. Based on their definitions, we map the five meeting elements to a set of narrative components that exist within all stories: *character*, *setting*, *plot*, *conflict*, *theme* and *point of view* [31] (see Table 2). These story elements and their connections to each other are shown in Fig. 2a, (note that *point of view* is not included because it is largely

TABLE 2
Five Essential Elements of Stories

Story Element	Narrative Structure Descriptor	Description
Characters	Who	The people or other being that the story is about
Plot	What	The sequence of events in a story
Theme	What	An idea or point that is central to a story
Setting	When, Where	The time and location describing the surroundings of the story
Conflict	How	The challenge main characters need to solve to achieve their goals




Fig. 2. Mapping between the essential elements of (a) a story and (b) a meeting.

determined by writing convention), and the equivalent structuring of meeting elements is shown in Fig. 2b.

In a story, the setting, characters, and plot are intricately connected. Characters interact with each other through conflict. These elements are in support of an overall theme, the central idea of the story.

The meeting elements are similarly arranged. Here, timeline, participants, and topics are connected. Interactions between participants are demonstrated as communication. The overall resolutions of the meeting are akin to tasks, which are decided upon as future agendas or action items that participants agree to perform for future work.

Given this narrative-based framing, a chosen visualization metaphor should emphasize these types of connections. We choose storyline visualization [8] for our base technique, as it is effective for representing interacting and grouped entities over time. Meeting participants map to lines and cluster together as they participate in the discussion over the duration of the meeting. We describe the visual mapping in detail the next section.

4 MEETINGVIS DESIGN

We now describe the visual interface and data pipeline implementation details for MeetingVis. The design requirements based on the sample survey are as follows:

- DR1 *Provide an intuitive representation of meeting elements.* Recognizable, explorable, simplified visuals make it easy for participants to understand meeting data by lowering the required cognitive threshold.
- DR2 *Organize meeting elements to trigger recall of specific memories.* Visually summarizing a meeting requires connecting elements as correlated components instead of independent events. Highlighting key interaction points (e.g., *who* talks about *what* at *when* during the meeting) reveals the underlying narrative structure when supporting memory recall.
- DR3 *Enable user refinement of automatically generated meeting results.* Meeting attendees might modify their notes afterward, adding to their memory banks. Human-in-the-loop (HITL) is an effective strategy for optimizing

or amending the summary results. Users should be provided with the options to edit, add/remove topics, and tweak other stylistic settings.

We note here that, in this design, our focus is creating a lightweight design that effectively summarizes and visualizes a meeting while triggering recall. Therefore, we intentionally omit standard data points that are commonly found in meeting browsers— notably access to the raw recorded audio, video, and speech transcript. For a full-featured system, these metadata elements can easily be integrated. See Section 7.1 for more discussion.

4.1 Augmented Storyline Visualization (DR1 & DR2)

MeetingVis’s visual interface shows the five meeting elements around three components: participant activities, task assignments, and topic evolutions. Three data points—*interactions* between *participants* along the meeting’s *timeline*—map to the base storyline technique. To show *topics* and *tasks*, we augment the view with additional features.

4.1.1 Participant Activities and Task Assignments

Each participant is shown as a color-coded line that tracks from left-to-right, as shown in Fig. 3 B-(i). Interaction is demonstrated by horizontal temporal positioning; when participants discuss the same topic, their lines cluster together for that portion of the timeline. We measure each person’s conversational contribution by counting the amount of words spoken over the course of the meeting [13], which is encoded in their line thickness (a thicker line means a participant is speaking more at that point in time). Participant activities can be analyzed by reviewing how their lines converge/diverge and how their thicknesses compare in relation to each other.

To denote assigned tasks, we append a triangle glyph to participant lines. To see a task description, hovering on the glyph pops up a tooltip with detail information, as shown in Fig. 3 B-(ii). As opposed to a linked list of action items, integrating tasks directly onto the presentation of participants immediately reveals which participants have been assigned tasks, how many tasks are assigned, and what each task entails in a single view.

4.1.2 Topic Evolutions

To show extracted topics that are discussed during the meeting, we append “topic bubbles” behind the storyline, as shown in Fig. 3 B-(iii). These are color-coded according to the number of participants that join the topic. A topic is colored “warm” when it contains more than half of meeting participants, “cool” when it has less. Each bubble has an attached word cloud, as shown in Fig. 3 B-(iv). This cloud contains a set of keywords which are extracted from the discussion pertaining to the topic. Keyword size is based on its importance. Word clouds are hidden by default and shown by hovering over a topic bubble with the mouse.

4.2 User Interface and Interaction (DR3)

The augmented storyline visualization makes up the main view in our tool (called the Storyline View); other components include the Title View, Edit Panel, and Control Panel, as shown in Fig. 3 A-D. The purpose of these are:




Fig. 3. Screenshot of the MeetingVis interface. (A) Title View shows header information about the meeting; (B) Storyline View displays augmented storyline visualization of the meeting. Users can click to review the details of a topic bubble and a task triangle; (C) Edit Panel enables topic adding, editing, or removing; (D) Control Panel is used to brush and zoom along the main Storyline View. The storyline summarizes the meeting of one discussion group from our user study with encoded participants' names.

A Title View: Shows header information about the meeting, its name, date, and a color key for visual elements.

B Storyline View: Displays the augmented storyline visual interface.

C Edit Panels: Used to adjust stylistic and data elements, such as editing, adding/removing a topic or a keyword.

D Control Panel: Used to brush and zoom along the main Storyline View.

4.3 Data Pipeline Implementation

MeetingVis's data pipeline (Fig. 4) utilizes NLP techniques to automatically capture and process meeting data. The architecture contains four backend components: speech recognition, speaker recognition, topic modeling, and information extraction.

4.3.1 Speech Recognition

Given a full audio recording of a meeting, the speech recognition component outputs the speech of participants as a word file. For this, we use Google Speech Recognition wrapped in the SpeechRecognition python module [32]. Google Speech API has a consistently higher accuracy in determining the words spoken (above 90 percent with the Standard American accent) when compared to other tools.




Fig. 4. The MeetingVis data pipeline for capturing, processing, and viewing meeting data.

4.3.2 Speaker Recognition

The speaker recognition component uses Microsoft Cognitive Services API [33]. We first train the voices of each participant with three 20-second samples. Next, we segment the meeting into 5 second intervals and request the service for the corresponding speaker of each segment. We match the speaker to speech recognition results of that particular segment. To ensure the accuracy and consistency of both speech and speaker recognition, we use the same microphone for voice training and meeting recordings.

4.3.3 Topic Modeling

We use the Topics over Time [34] model to find relevant topics by time segment throughout the meeting. The algorithm is based on word co-occurrences and time-localization. We utilize this topic model because it works well with timestamp metadata and time segmentation, as the meeting transcript is broken into time fragments.

4.3.4 Information Extraction

Natural Language Toolkit [35] is used to perform information extraction, analyzing the transcript to assign tasks to certain participants based on the content of the meeting. We go line by line through the transcript, breaking each sentence into tokens and looking for particular verb sequences that indicate being assigned a task, such as "to do", "to handle", or "to take on". Once a sequence is found without a negative in front, we backtrack from that point to find a pronoun or proper noun in order to map the action to a particular individual. We then analyze the line or lines that follow to check that the individual confirmed the task. Finally, we add this task to an object representing that individual. This process eventually builds up a table that stores the tasks participants have assigned to one another.

5 EVALUATION

To evaluate the effectiveness of MeetingVis, we conducted a qualitative experiment with five discussion groups (termed G1-G5). To help direct analysis we posed three research questions to answer:

- RQ1** How do participants perceive the visualization?
- RQ2** How does the system support memory retrieval?
- RQ3** What refinement strategies are applied to the automatically generated result?

5.1 Method

The study was composed of three phases: (1) the group's *meeting* phase, (2) a post-meeting *recall* phase to view the visual summary, and (3) the subsequent *explore* phase where participants could refine the autogenerated results.

Meeting Phase. In the meeting phase, we recorded each group's discussion. For G1, G2, and G3, we attended a regularly scheduled meeting for their ongoing projects, which is the most prevalent meeting purpose [36]. For G4 and G5, each group was told that they were a committee tasked with organizing a charity event [13]; they had to make decisions such as the charity's cause, the food and refreshments, and assign responsibilities.

A camera and microphone recorded each meeting; the experimenters were not present to minimize any obtrusive effect. Participants were encouraged to maintain their usual meeting behavior which might boost memory such as taking notes, and were told a follow-up session would be conducted the next week. After each meeting, the audio feed was processed according to the data pipeline steps in Fig. 4. To ensure accuracy, two external, independent auditors compared these results by checking them against the meeting's video recording.

Recall and Explore Phases. Meeting participants were asked to come to our lab one week after the recorded meeting (before their next regular meeting) to conduct the recall and explore phases. Both phases were done individually to prevent group synergy.

For the recall phase, we asked each participant to verbally walk through their prior meeting using only memory. This established a recall baseline from which to estimate the add-on benefits that our tool could provide. We then demonstrated a mock, summarized meeting to teach the visual layout. After this, we loaded the participant's own meeting in the viewer and asked them to review and investigate it. While doing this, they used think-aloud protocol to state any additional remembered details. Once satisfied in reviewing the meeting, we administered a short survey and questionnaire where participants rated how the visualization affected their memory experience (adapted from [37]) using a 5-point Likert scale. Specifically, the questionnaire asked to rate how MeetingVis affected the following:

- *Helpfulness.* This visualization augmented your memory of the meeting.
- *Visual encodings.* This visualization offered a good representation of the meeting.
- *New details.* You were reminded of aspects of topics from the meeting you had not previously thought of.

TABLE 3
Demographics and Meeting Data from the Discussion Groups

Group ID	# of Participants (females)	Ave Age	Discussion Time	Discussion Topic
G1	9 (4)	19.7	0:14:25	Hackathon Schedule
G2	5 (2)	33.4	1:08:58	Writing Club
G3	4 (1)	21.2	0:31:51	Management Coursework
G4	5 (0)	27.8	0:23:29	Charity Event Planning
G5	5 (1)	25.8	0:26:52	Charity Event Planning

- *Usefulness.* The new details you recall are useful for the next meeting.
- *Informativeness.* The important points of the meeting are represented in the summary.
- *Readability.* It is generally easy to tell who or what is being referred to in the summary.
- *Vividness.* Your memories are very vivid and detailed.
- *Time perspective.* Your memories for the hour when the event took place are clear.
- *Accessibility.* Your memories are easy for you to recall.
- *Coherence.* Your memories come back to you as a logical, coherent story, not in bits and pieces.
- *Sensory details.* Your memories for this event involve a lot of physical reactions and sensory information (sounds, etc.).
- *Valence.* The impact on your memory is positive.

For the explore phase, we introduced the tool's edit functionality. Each participant was allowed to refine and stylize the automated summary according to their personal preferences. When they reached a satisfactory result, a semi-structured interview about the usability of the system was conducted.

For both phases, participants sat with a proctor at a computer station (an iMac with 2560 × 1440 display). A microphone recorded audio, and the proctor observed the actions and strategies used by participants. The program logs of user actions were recorded for later references.

5.2 Participants

Discussion groups were recruited via college mailing lists. Our recruitment material specified that we were looking for pre-existing project groups of 4-10 members who have regular weekly meetings. The five groups had a total of 28 participants (7 females) with an average age of 24.9 (SD = 5.9), and included college students, researchers, and professionals from genetics, computer science, bioinformatics, neurobiology, and managerial economics backgrounds (see Table 3). Participants completed a brief demographic survey before starting the experiment and we recorded a voice sample of each for the speaker recognition step in the data processing pipeline.

All participants took part in the meeting phase, and 19 participants (4 females) with an average age of 27.4 (SD = 5.4) completed the recall and explore phases. A majority of participants (79 percent) noted that they have the habit of recalling a meeting to prepare for the next meeting using unaided memory (73.7 percent) or personal notes (42.1 percent).




Fig. 5. Diverging stacked bar chart of usability ratings with respect to overall experience, summary evaluation, and memory evaluation of MeetingVis. Each participant answered on a 5-point likert scale with the questionnaire.

6 RESULTS AND ANALYSIS

In the recall phase, users first walked through the meeting using memory only. In assessing memory baselines, most users could either list off the main set of topics that were discussed or go through the meeting points, but struggled to remember many of the details.

Compared to the baseline, all 19 users were able to recall new items from their meetings by using MeetingVis, such as more detailed descriptions of discussions, critical points, and the order of topics discussed. When reviewing a meeting summary, most users (12/19) first explored the topic bubbles and word clouds, moving through each in chronological order, before reading the tasks assigned to each person. Several users (4/19) reviewed topics based on size, starting with the largest.

Quantitatively, our tool rated highly based on the questionnaire results, including aspects regarding user experience, summary evaluation, and memory evaluation (see Fig. 5). There was variation however in both *how much* it helped and *which* of the visual elements were considered as relevant. To better explain this, we now discuss observations and user feedback for the three posed research questions (RQ1-RQ3).

6.1 RQ1: How Users Perceived the Visualization

6.1.1 Topic Bubbles to Bring up More Details

All users found the topic bubbles and keyword clouds helpful; these were usually the first visual feature noted. Users thought this visualization reminded them of parts and aspects of the meeting that they had not previously thought of (see New Details rating in Fig. 5, $\mu = 4.33$, $\sigma = 0.69$). “As we go over this, like a bigger view, we could be like, ‘oh, I remember at

that time at the meeting we’re talking about...’ And we just go to this block specifically” (G1P5). “I completely did not recall this ‘Airbnb’ conversation, that’s an important part in our meeting, so it definitely helps me remember that” (G3P1). Most users were quickly able to scan the weighted words for each topic and pick out personally relevant ones. “The word cloud helps me to remember the context that we were talking them in. So when I think of a keyword, I think, ‘okay, why did we say that keyword. Oh, cause we were talking about these at that time’ ” (G4P1).

Five users mentioned that the topic bubbles imply discrete breaks in conversation; they considered this confusing. They suggested that the tool should show the transition between topics better. Another noted that some topics seemed to bleed into each other due to the nature of their discussion and that the topic segmentation could be improved. A majority of users recommended that we should include automatic topic labeling as a feature.

6.1.2 Participant Lines for Tracking Individual Activities and Interactions

More than half of the users analyzed the participant lines and their groupings, it was easy for them to map their meeting experience onto participant line presentation. “It’s very clear on who talked, and how much they talked. This line did help in terms of if I’m still there or not” (G1P4). One method of identifying individual contribution was to track the line thickness over time. “I obviously talked a little too much” (G4P1). “Thickness is a very straightforward representation. [G4P1] was driving this meeting, it’s easy to see that! [G4P4] also talked part of it, and the rest of us talked equal parts” (G4P2).

Several users asked about the vertical temporal placement of lines, and noted that it should be used to encode individual verbal contribution. Another user suggested that we use line breaks to denote inactive meeting participants as opposed to having lines leave and re-enter topic bubbles.

6.1.3 Timelines for Cognitive Orientation

Timelines organize events in a chronological sequence, from which users were able to reconstruct the experience (see Time Perspective rating in Fig. 5, $\mu = 4.00$, $\sigma = 0.84$). “Going through the timeline puts me in a time perspective and adds a lot of visual content to my recollection of the memory. So I can replay it from beginning to end with greater accuracy” (G3P1). Participants used the time scale to evaluate each topic of the discussion. “I feel it’s really useful to have a timeline for me to say, ‘hey this conversation took twenty minutes or that took ten seconds’ by looking at the horizontal axis.” (G1P3).

Most users brushed timeline in the control panel to switch between global perspective and detail perspective. One user suggested showing meeting transcript as one moves mouse to a particular point on the timeline, as it helps navigate through a more detailed summary and find information quickly.

6.1.4 Task Icons Work as Reminders

Regarding the task icons, feedback was uniformly positive; users thought these straightforwardly represent group contract of decisions and commitments. “I think the triangles are really helpful. Because in the end, we were like, ‘okay, this is what

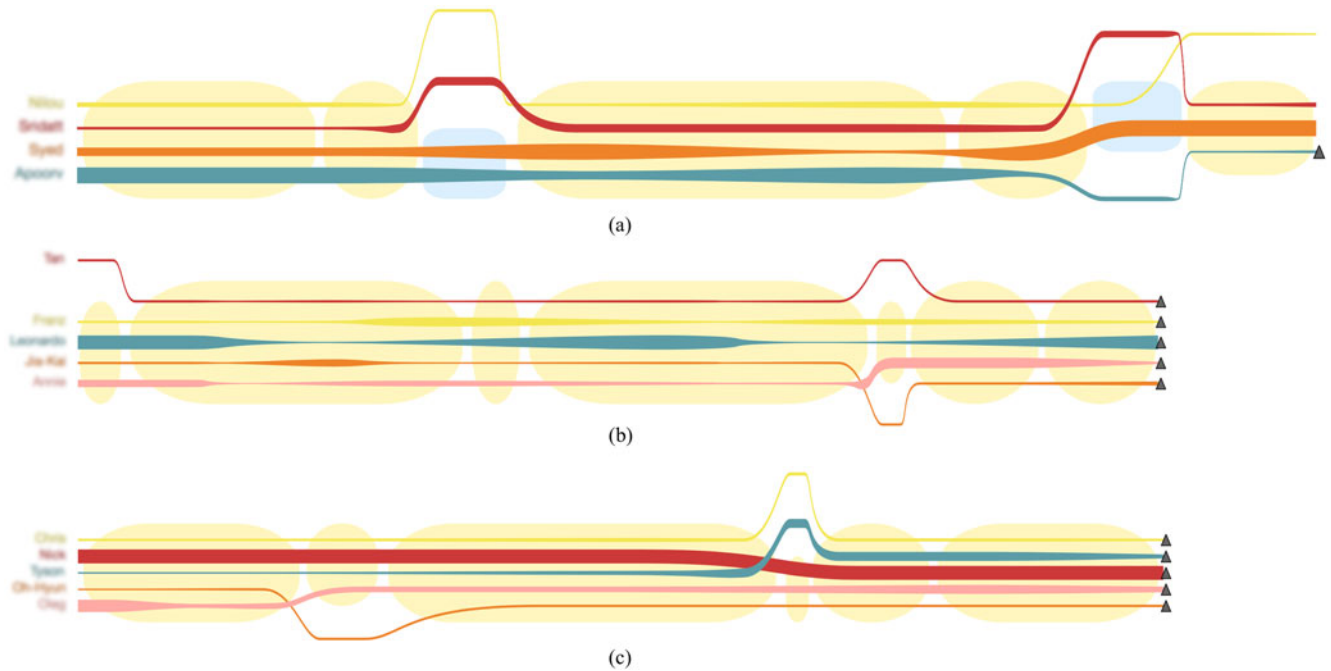


Fig. 6. Visualization of meetings by discussion groups (a) G3, (b) G4, and (c) G5. Note that participant names have been blurred out, and the meeting in (a) has a longer duration than (b) and (c).

we need to get done.’ This is a really good reminder of what has been assigned” (G1P5). Users usually remembered their own task(s), but rarely recalled those of others. “The visualization reminded me what tasks were assigned to who, which is very important” (G4P3).

Task tooltips were reviewed by all users. One user suggested that task reminders should be sent to each participant’s smart phone, which could automatically notify them of work assignment before their next regular meeting.

6.2 RQ2: How the System Supported Memory Recall

6.2.1 It’s a Self-Contained Meeting Summary

Users mostly thought MeetingVis provides a succinct and efficient way to review the major points and flows of their meetings. “You can see who’s there at first glance. As you progress mouse to mouse, you can see how the conversation progresses by the frequency of the words. Good simplicity, good intuition for big picture scale” (G2P4). “Even if it’s just the lines, I can really tell how much we talked, and who talked. That helped me remember about the topics (in the meeting)” (G5P3). “Audio recordings of meetings are the most monotonous things to listen to. This tool sort gives a transcript without actually transcribing it. You can go through the timeline and realize what happened without having to read the whole thing or listen to the whole thing” (G3P1).

Through the exploration of the visual retrieval cues such as topic bubbles and task icons, users recalled more specific details, ideas, and sub-discussions. “It made me remember specific images or specific scenes of what happened, I remember [G1P6] talking specifically” (G1P4). “Even though I have a good memory on this [meeting]. This is definitely helpful. So the big points I definitely remember. But sometimes I don’t remember how I got to that topic in the first place, or what happened next necessarily. By looking at the smaller ones, it helped me remember smaller details that I wouldn’t remember otherwise” (G1P5).

Some users noted that while our system supports the recall of the “big picture” or “overview” of a meeting, it didn’t help much in remembering sensory information of the discussion (see Sensory Details rating in Fig. 5, $\mu = 2.89$, $\sigma = 1.32$). “It’s trying to give you an at-a-glance view of the event. It doesn’t leave as much room for details. The only really big [draw-back] I see is the loss of detail, for example, sound” (G4P4). “The aspect of tone is not really covered in this. I don’t know if someone is necessarily angry. I wish there’s a way to analyze emotion, like how their tone of voice is.” (G1P3).

6.2.2 It Shows Different Discussion Structures

The visualization results show different discussion structures from group discussions. Participants were able to recall their participation and evaluate the productivity accordingly. For example, Fig. 6a shows the meeting result of G3, which presents two “cold” word clouds, which are colored light blue. These suggest that several participants were not involved in these sub-discussions and the discussion during this time period might not be an effective, cohesive one. Several participants commented on this chart, “it’s a lot of information, so the coherence perhaps is not fully there at least for me to interpret it” (G3P1). “I think overall its helpful, but maybe missing some keywords [from the word clouds]” (G5P1).

We found different participation patterns when comparing the visualizations of charity event planning meetings of group G4 (see Fig. 6b) and G5 (Fig. 6c). For G4, the thickness of participant lines shows that different phases in the meeting were led by specific users. “Having things split up into different phases, that’s very nice. I had in my mind, we went through phases, and then being able to go back and see some of the main things we talked about. The different phases help as well” (G4P1). For G5, the thickness of participant lines suggests that the red participant was leading the whole meeting session while other participants had less opportunities to discuss and express their thoughts. “It provides you with a structure to

the meeting. Maybe I wanna improve and be a better participant next time, I can see what makes me less engaged in the meeting.” (G5P2).

6.2.3 It Acts as a Note-Taking Assistant

Many users felt MeetingVis provides a way to replace or at least supplement manual note-taking. Although an auto-generated meeting summary might produce mismatched results or biased key points, it reduces cognitive load and provides an adaptable base. “A lot of times with meetings, you’ll have somebody take notes and you’ll send out meeting minutes. Probably nobody gonna read through that entire email, it has to be summarized. And rather than somebody writing up and customizing summaries, and instead what you do is somebody would say ‘here’s the automatic layout, let’s rearrange this a little bit’” (G4P1).

For meetings that are regularly scheduled, two users mentioned this tool could help better plan for future discussions. “Normally, I might have just written down my task for the next meeting without recalling much of the thought process behind the tasks we came up with. With the visualization tool, I was able to remember more of the ideas we explored, which might help us in the next meeting not to waste time going over ideas we already thought about, or remembering why we came to certain decisions” (G5P2).

6.3 RQ3: How Users Refined the Automated Visual Summary

6.3.1 Less-Important Topics were Deleted, not Ignored

All users performed a combination of adding labels to topics, adjusting bubble colors, and playing with topic start/stop times. Seven users deleted topics that they felt were unimportant (smaller topics). “I think I would delete the topics. There were several times we did sidetrack on particular topic (G4P5). “I found these ‘tweener’ topics contains a mix of keywords from the larger surrounding (topic) bubbles (G4P2). Three users actually added new topics, “having the option to create a new topic is helpful. Maybe in between here, we could have smaller ones that didn’t get caught but we know where they were” (G1P5). One user remarked that our visual design and interface was a good way to make the summary presentation-worthy: “That’s really helpful if I wanted to edit the meeting and show it to my boss” (G5P2).

6.3.2 Less-Important Keywords Were Ignored, not Deleted

Interestingly, only a few users (3/19) actually inserted or removed keywords into the topic word clouds. Instead of deleting keywords that were not seen as relevant to the topic, users tended to adjust the weights of them within the word cloud (that is, make them larger or smaller). Fourteen users performed this action on keywords for at least one topic. One reason could be to its perceived context. For example, when de-weighting a word, users explained that even if it was frequently spoken and thus highly weighted by our system, contextually it might be an aside to the main discussion at hand: “Some of the representation might be a little bit misleading. ‘Bricks’ is actually a little larger than some other stuff. ‘Nugget’ was definitely more important than ‘bricks’,

‘bricks’ we’re mostly joking about. We actually did spend a fair amount of time talking about it even though it’s less important for a meeting” (G4P1).

7 DISCUSSION

We discuss design implications from our user study, limitations in our current work, and plausible future research directions.

7.1 Narrative Structure in Meeting Summary

MeetingVis performs a straightforward mapping of the five selected meeting elements to visual encodings. One concern raised by participants was that these retrieval cues are limited in showing some of the more nuanced aspects of their discussions, such as the fine-grained context and topic sentiments. For example, adding sensory details such as how speaker tones and emotional valence change would be instructive, as the underlying meaning of words can be interpreted different ways depending on a speaker’s inflection.

Different architectural models of narrative schemas can be explored in allowing for more varieties and complexities of meetings. León [38] proposed that in storing episodic content, the memory units are linked by narrative relations (i.e., *time, location, causality, agency, abstraction* and *composition*). In applying this theory to MeetingVis, topics discussed can further be organized based on their *causal relations* with regard to preconditions and effect. A causality network can be constructed to reveal a meeting’s path towards its conclusion or the train of thoughts of the participants (e.g., [39]).

Alternatively, different levels of *abstraction* can be applied to the topics, allowing hierarchical investigation of discussion details. For example, conventional data points commonly found in meeting browsers can be considered. Several subjects requested the ability to play back audio and video snippets from conversation, or be able to access specific speech transcript sections as a “details on demand” component, which supports expanding a specific timeframe and contextualizing a key point.

7.2 Visualizations of Multi-Party Discourse

The current design of MeetingVis adapts the storyline visualization idiom, which is typically reserved for movie plots, to a new and practical context; it shows who is talking about what topics at what time in a meeting. One of the strengths of storylines is the diverging and merging lines along the vertical axis as characters move between several groups or categories over time. In contrast, MeetingVis has no parallel representation as all participants are assumed to only be discussing one topic at a time.

We suggest further exploring the design space of narrative visualization to convey stories using meeting data [40] (e.g., annotated charts [41]). For example, the synthetic, charity meetings generally had more back-and-forth conflict between participants. Instead of using discrete topic bubbles, ThemeRiver-based visualization [42] could help users identify if participants reach consensus during the meeting. The overall argumentative variation such as idea merging, branching, and interleaving, can be represented as a flow graph (e.g., [43], [44]) behind the storyline.

We evaluated our system with discussion groups of small size. As the number of participants increase (e.g., research meeting with international partners usually reach more than 20 participants), storylines can have perceptual issues when too many lines are simultaneously plotted. To ensure scalability for large group discussions, level-of-detail (LOD) techniques are worth exploring [45]). For example, for participants who share the same opinion, their lines can be bundled together to represent an opinion group.

7.3 Visual Encodings of Meeting Elements

In our design, the word clouds present the extracted keywords, the lines visualize participants, and the triangle glyph shows tasks. According to user feedback, these visual encodings can be improved.

In future MeetingVis versions, we are considering to replace word clouds with multi-word phrases or sampled quotes to provide a better recall support for a topic. In addition to font size, other visual channels, such as font color intensity, the use of additional marks, and spatial layout [46], can be explored to show the importance of keywords.

One participant suggested that showing lines leaving and re-entering topic bubbles is not a straightforward representation of inactive participation. Linear diagrams [47] can be used as an alternative to spatially group participants; line breaks will then denote inactive meeting participants.

MeetingVis can be generalized to other forms of multi-party discourse by introducing new glyph for characteristic interactive events. For example, we tested MeetingVis in a different context by collecting and visualizing data from our lab's messaging application: Slack. In analyzing its group discussions, we noted that sharing digital links and resources is common. Augmenting MeetingVis with "resource glyphs" that attach to the storyline is one way to handle this.

7.4 Collaborative Refinement for a Full Meeting Picture

A comprehensive meeting summary can provide new insights for decision-making and retrospection. In the explore phase, we observed that our MeetingVis users had a tendency to refine cues that were more immediately relevant to their own participatory experience, such as the sub-discussions, keywords, and tasks they were involved in. To develop a full picture, it is important to gather the scattered refinements from different participants and reconstruct a full, holistic meeting experience.

To support synchronous refinement activities, a collaborative workspace that enables organization, sharing, and coordination of modifications can be integrated (e.g., [48]). This would allow for iteratively merging and connecting separate edits collected from multiple users, adding them to the auto-generated meeting summary. In such scenarios, it will be important to provide views showing edit histories to see what has been changed from the default summary. When detecting matching or conflicted edits regarding a specific meeting point, a matched graph can be built to evaluate and analyze these retrieved details from different sources (e.g., [49]). Meeting notes or additional resources could also be included to validate the effectiveness of these changes.

7.5 Comparative Analysis of Meeting Sequence

Effective analysis of previous meetings can have a positive influence on the planning of future work. Our current user study solicited users' first impressions with MeetingVis. While the preliminary results indicate we are on the right path, it would be more practically useful to conduct a longer-term evaluation. For example, group meetings held according to a regular schedule over a long-term period can be processed and pieced together to form a linked set of visual summaries. Comparative analysis (e.g., [50]) of multiple meeting sequences can help users reflect on participation, contribution, and collaboration at both individual and group levels.

Various visual cues in our design are beneficial in reminding participants of individual responsibilities, commitments, and involvements. As shown in Fig. 3, if Charizard who mainly contributed to only one topic wants to be more active, he can review the visual summary and reflect on what makes him less engaged in the meeting. If multiple meeting summaries are available, he can trace his performance along a sequence of summaries to see how he makes progressions over time.

At the group level, we believe our tool is adaptable for large-scale project development. Comparative analysis of visual meeting summaries could facilitate review of a project's history, evolution and development. Through the observation of differences in conversational patterns and the exploration of strategies adopted by the participants, a group can gain insight on whose contributions led to certain design decisions, how collaboration improves synergy effects, and how previous projects can be used to make better decisions in the future.

7.6 Limitations

Our study reveals a few limitations with the initial version of MeetingVis. First, we use ASR to transcribe participant speech. To ensure valid results, we asked two external auditors to compare these results by checking them against the meeting's recording. In real scenarios, accuracy can decrease due to, for example, heterodox speaking styles. Regional variation, speech tempo, and slang use all act as confounding factors against accurately capturing keywords. Second, to ensure the accuracy of speaker recognition, participants were instructed to speak one at a time. However, in lively discussions, participants certainly tend to talk simultaneously. The use of more advanced audio capturing hardware and emerging NLP methods with better multi-user handling could provide some improvement.

8 CONCLUSION

As real-world problem solving is often done with a collective effort through multiple rounds of deliberation on design ideas and test results, a technology like MeetingVis that can facilitate more productive group discussion is timely and desirable. We have demonstrated the use of a visual, narrative approach to organize and present essential meeting elements in a lightweight, intuitive, and informative manner can improve meeting recall and prompt attendees to reflect on their prior participation, contributions, and overall collaboration within the discussion. Our work has

provided a good proof of concept and our findings can help direct further research, which may include employing multimodal sensing and advanced natural language processing techniques as well as designing multiple representations of different aspects of the captured knowledge and collaborative activities.


ACKNOWLEDGMENTS

This research was sponsored in part by the Fundamental Research Funds for the Central Universities No. 22120180012, the U.S. National Science Foundation through grants IIS-1528203 and IIS-1320229, and the National Natural Science Foundation of China under grant No.61402540, No.61672538. This presented work was done while Yang Shi was with VIDI Labs, UC Davis.

REFERENCES

- [1] N. C. Romano Jr and J. F. Nunamaker Jr, "Meeting analysis: Findings from research and practice," in *Proc. IEEE Hawaii Int. Conf. Syst. Sci.*, 2001, Art. no. 1072.
- [2] S. Whittaker, R. Laban, and S. Tucker, "Analysing meeting records: An ethnographic study and technological implications," in *Proc. Mach. Learn. Multimodal Interaction*, 2005, pp. 101–113.
- [3] E. Tulving, "Episodic and semantic memory," in *Organization of Memory*. London, U.K.: Academic, 1972.
- [4] M. Gambhir and V. Gupta, "Recent automatic text summarization techniques: A survey," *Artif. Intell. Rev.*, vol. 47, no. 1, pp. 1–66, 2017.
- [5] S. Whittaker, S. Tucker, and D. Lalanne, "Meeting browsers and meeting assistants," in *Multimodal Signal Processing: Human Interactions in Meetings*. Cambridge, U.K.: Cambridge Univ. Press, 2011, ch. 12, Art. no. 204.
- [6] A. J. Cowell, M. L. Gregory, J. Bruce, J. Haack, D. Love, S. Rose, and A. H. Andrew, "Understanding the dynamics of collaborative multi-party discourse," *Inf. Vis.*, vol. 5, no. 4, pp. 250–259, 2006.
- [7] S. Castronovo, J. Frey, and P. Poller, "A generic layout-tool for summaries of meetings in a constraint-based approach," in *Machine Learning for Multimodal Interaction*. Berlin, Germany: Springer, 2008, pp. 248–259.
- [8] Y. Tanahashi and K.-L. Ma, "Design considerations for optimizing storyline visualizations," *IEEE Trans. Vis. Comput. Graph.*, vol. 18, no. 12, pp. 2679–2688, Dec. 2012.
- [9] T. Makany, J. Kemp, and I. E. Dror, "Optimising the use of note-taking as an external cognitive aid for increasing learning," *British J. Edu. Technol.*, vol. 40, no. 4, pp. 619–635, 2009.
- [10] J. L. Fisher and M. B. Harris, "Effect of note taking and review on recall," *J. Edu. Psychol.*, vol. 65, no. 3, 1973, Art. no. 321.
- [11] F. Khan, *A Survey of Note-Taking Practices*. Palo Alto, CA, USA: Hewlett-Packard Laboratories, 1993.
- [12] G. Oleksik, N. Milic-Frayling, and R. Jones, "Study of electronic lab notebook design and practices that emerged in a collaborative scientific environment," in *Proc. ACM Conf. Comput. Supported Cooperative Work Social Comput.*, 2014, pp. 120–133.
- [13] V. Kalnikaitė, P. Ehlen, and S. Whittaker, "Markup as you talk: Establishing effective memory cues while still contributing to a meeting," in *Proc. ACM Conf. Comput. Supported Cooperative Work*, 2012, pp. 349–358.
- [14] M. Nathan, M. Topkara, J. Lai, S. Pan, S. Wood, J. Boston, and L. Terveen, "In case you missed it: Benefits of attendee-shared annotations for non-attendees of remote meetings," in *Proc. ACM Conf. Comput. Supported Cooperative Work*, 2012, pp. 339–348.
- [15] W. Willett, P. Goffin, and P. Isenberg, "Understanding digital note-taking practice for visualization," *IEEE Comput. Graph. Appl.*, vol. 35, no. 4, pp. 38–51, Jul./Aug. 2015.
- [16] N. Kern, B. Schiele, H. Junker, P. Lukowicz, and G. Tröster, "Wearable sensing to annotate meeting recordings," *Personal Ubiquitous Comput.*, vol. 7, no. 5, pp. 263–274, 2003.
- [17] A. Jaimes, K. Omura, T. Nagamine, and K. Hirata, "Memory cues for meeting video retrieval," in *Proc. ACM Workshop Continuous Archival Retrieval Personal Experiences*, 2004, pp. 74–85.
- [18] Z. Yu, M. Ozeki, Y. Fujii, and Y. Nakamura, "Towards smart meeting: Enabling technologies and a real-world application," in *Proc. ACM Int. Conf. Multimodal Interfaces*, 2007, pp. 86–93.
- [19] A. Girgensohn, J. Marlow, F. Shipman, and L. Wilcox, "HyperMeeting: Supporting asynchronous meetings with hyper-video," in *Proc. 23rd ACM Int. Conf. Multimedia*, 2015, pp. 611–620.
- [20] S. Hunter, P. Maes, S. Scott, and H. Kaufman, "MemTable: An integrated system for capture and recall of shared histories in group workspaces," in *Proc. SIGCHI Conf. Human Factors Comput. Syst.*, 2011, pp. 3305–3314.
- [21] P. Ehlen, M. Purver, J. Niekrasz, K. Lee, and S. Peters, "Meeting adjourned: Off-line learning interfaces for automatic meeting understanding," in *Proc. ACM Int. Conf. Intell. User Interfaces*, 2008, pp. 276–284.
- [22] M. El-Assady, V. Gold, C. Acevedo, C. Collins, and D. Keim, "ConToVi: Multi-party conversation exploration using topic-space views," in *Proc. Eurographics Conf. Vis. Comput. Graph. Forum*, 2016, pp. 431–440.
- [23] Y. Shi, Y. Wang, Y. Qi, J. Chen, X. Xu, and K.-L. Ma, "IdeaWall: Improving creative collaboration through combinatorial visual stimuli," in *Proc. ACM Conf. Comput. Supported Cooperative Work Social Comput.*, 2017, pp. 594–603.
- [24] M. El-Assady, R. Sevastjanova, B. Gipp, D. Keim, and C. Collins, "NEREX: Named-entity relationship exploration in multi-party conversations," in *Proc. Eurographics Conf. Vis. Comput. Graph. Forum*, 2017, pp. 213–225.
- [25] S. Tucker and S. Whittaker, "Accessing multimodal meeting data: Systems, problems and possibilities," in *Proc. Int. Workshop Mach. Learn. Multimodal Interaction*, 2004, pp. 1–11.
- [26] Minute. [Online]. Available: <https://www.getminute.com/>
- [27] Meetingking. [Online]. Available: <http://meetingking.com/>
- [28] P. Scripts, *Goals, and Understanding: An Inquiry into Human Knowledge Structures*. Hillsdale, NJ, USA: Erlbaum, 1977.
- [29] J. Bruner, "The narrative construction of reality," *Critical Inquiry*, vol. 18, no. 1, pp. 1–21, 1991.
- [30] D. C. Rubin and D. L. Greenberg, "The role of narrative in recollection: A view from cognitive psychology and neuropsychology," in *Narrative and Consciousness*, New York, NY, USA: Oxford Univ. Press, 2003, pp. 53–85.
- [31] R. Bran, "Message in a bottle telling stories in a digital world," *Procedia-Social Behavioral Sci.*, vol. 2, no. 2, pp. 1790–1793, 2010.
- [32] Speechrecognition Python Module. [Online]. Available: <https://pypi.python.org/pypi/SpeechRecognition/3.4.3>
- [33] Microsoft Cognitive Services API. [Online]. Available: <https://www.microsoft.com/cognitive-services/en-us/apis>
- [34] X. Wang and A. McCallum, "Topics over time: A non-Markov continuous-time model of topical trends," in *Proc. ACM Int. Conf. Knowl. Discovery Data Mining*, 2006, pp. 424–433.
- [35] Natural language toolkit. [Online]. Available: <http://www.nltk.org/>
- [36] J. A. Allen, T. Beck, C. W. Scott, and S. G. Rogelberg, "Understanding workplace meetings: A qualitative taxonomy of meeting purposes," *Manage. Res. Rev.*, vol. 37, no. 9, pp. 791–814, 2014.
- [37] M. Luchetti and A. R. Sutin, "Measuring the phenomenology of autobiographical memory: A short form of the memory experiences questionnaire," *Memory*, vol. 24, pp. 592–602, 2016.
- [38] C. León, "An architecture of narrative memory," *Biologically Inspired Cogn. Archit.*, vol. 16, pp. 19–33, 2016.
- [39] D. Shahaf, C. Guestrin, and E. Horvitz, "Trains of thought: Generating information maps," in *Proc. ACM Int. Conf. World Wide Web*, 2012, pp. 899–908.
- [40] E. Segel and J. Heer, "Narrative visualization: Telling stories with data," *IEEE Trans. Vis. Comput. Graph.*, vol. 16, no. 6, pp. 1139–1148, Nov./Dec. 2010.
- [41] C. Bryan, K.-L. Ma, and J. Woodring, "Temporal summary images: An approach to narrative visualization via interactive annotation generation and placement," *IEEE Trans. Vis. Comput. Graph.*, vol. 23, no. 1, pp. 511–520, Jan. 2017.
- [42] S. Havre, E. Hetzler, P. Whitney, and L. Nowell, "ThemeRiver: Visualizing thematic changes in large document collections," *IEEE Trans. Vis. Comput. Graph.*, vol. 8, no. 1, pp. 9–20, Jan.–Mar. 2002.
- [43] W. Cui, S. Liu, Z. Wu, and H. Wei, "How hierarchical topics evolve in large text corpora," *IEEE Trans. Vis. Comput. Graph.*, vol. 20, no. 12, pp. 2281–2290, Dec. 2014.

- [44] X. Wang, S. Liu, Y. Chen, T.-Q. Peng, J. Su, J. Yang, and B. Guo, "How ideas flow across multiple social groups," in *Proc. IEEE Conf. Vis. Anal. Sci. Technol.*, 2016, pp. 51–60.
- [45] S. Liu, Y. Wu, E. Wei, M. Liu, and Y. Liu, "StoryFlow: Tracking the evolution of stories," *IEEE Trans. Vis. Comput. Graph.*, vol. 19, no. 12, pp. 2436–2445, Dec. 2013.
- [46] C. Felix, S. Franconeri, and E. Bertini, "Taking word clouds apart: An empirical investigation of the design space for keyword summaries," *IEEE Trans. Vis. Comput. Graph.*, vol. 24, no. 1, pp. 657–666, Jan. 2018.
- [47] P. Rodgers, G. Stapleton, and P. Chapman, "Visualizing sets with linear diagrams," *ACM Trans. Comput.-Human Int.*, vol. 22, no. 6, 2015, Art. no. 27.
- [48] N. Mahyar and M. Tory, "Supporting communication and coordination in collaborative sensemaking," *IEEE Trans. Vis. Comput. Graph.*, vol. 20, no. 12, pp. 1633–1642, Dec. 2014.
- [49] X. Wang, S. Liu, J. Liu, J. Chen, J. Zhu, and B. Guo, "TopicPanorama: A full picture of relevant topics," *IEEE Trans. Vis. Comput. Graph.*, vol. 22, no. 12, pp. 2508–2521, Dec. 2016.
- [50] J. Zhao, Z. Liu, M. Dontcheva, A. Hertzmann, and A. Wilson, "MatrixWave: Visual comparison of event sequence data," in *Proc. ACM SIGCHI Conf. Human Factors Comput. Syst.*, 2015, pp. 259–268.



Yang Shi received the PhD degree in computer science from Central South University, China, in 2017. She is currently an assistant researcher with the College of Design and Innovation, Tongji University, Shanghai, China. Her current research interests include data visualization and human computer interaction.



Chris Bryan received the BS degree in computer science from the University of Arkansas, Fayetteville. He is working toward the PhD degree at the University of California, Davis, studying computer science and information visualization under Kwan-Liu Ma as a part of ViDi research group. His research interests include data storytelling, large-scale data analysis and visualization, and machine learning. He is a student member of the IEEE.



Sridatt Bhamidipati received the graduate degree in computer science & engineering and a minor in technology management from the UC Davis, in 2017. He is a software engineer with the LeanData, a startup in the Bay area. During his time at Davis, he worked in the ViDi lab as an undergraduate researcher and co-founded an event called HackDavis that aims to cultivate a hacker community at Davis.



Ying Zhao received the PhD degree in computer science from Central South University, in 2015. He is currently an associate professor in the School of Information Science and Engineering, Central South University. His research interests include information visualization and visual analytics.



Yaoxue Zhang received the BSc degree from the Northwest Institute of Telecommunication Engineering, China, in 1982 and the PhD degree in computer networking from Tohoku University, Japan, in 1989. Currently, he is a professor with the School of Information Science and Engineering, Central South University, China, and also a professor with the Department of Computer Science and Technology, Tsinghua University, China. His research interests include computer networking, operating systems, ubiquitous/pervasive computing, transparent computing, and big data. He has published more than 300 technical papers in international journals and conferences, as well as nine monographs and textbooks. He is a fellow of the Chinese Academy of Engineering and the editor-in-chief of the *Chinese Journal of Electronics*.



Kwan-Liu Ma received the PhD degree in computer science from the University of Utah, in 1993. He is a professor of computer science and the chair of the Graduate Group in Computer Science (GGCS), University of California-Davis, where he directs ViDi Labs and UC Davis Center of Excellence for Visualization. His research spans the fields of visualization, computer graphics, high-performance computing, and user interface design. During 1993-1999, he was with ICASE/NASA Langley Research Center as a research scientist. He joined UC Davis in 1999. He is a fellow of the IEEE.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.