IdeaWall: Improving Creative Collaboration through Combinatorial Visual Stimuli

Yang Shi*
Central South University
shiyang1230@gmail.com

Yang Wang
University of California, Davis
ywang@ucdavis.edu

Ye Qi*
Zhejiang University
ye.charlotte.qi@gmail.com

John Chen
University of California, Davis
jhochen@ucdavis.edu

Xiaoyao Xu*
Zhejiang University
hsu.xiaoyao@gmail.com

Kwan-Liu Ma
University of California, Davis
ma@cs.ucdavis.edu

ABSTRACT

With the recent advances in computer-supported cooperative work systems and increasing popularization of speech-based interfaces, groupware attempting to emulate a knowledgeable participant in a collaborative environment is bound to become a reality in the near future. In this paper, we present IdeaWall, a real-time system that continuously extracts essential information from a verbal discussion and augments that information with web-search materials. IdeaWall provides combinatorial visual stimuli to the participants to facilitate their creative process. We develop three cognitive strategies, from which a prototype application with three display modes was designed, implemented, and evaluated. The results of the user study with twelve groups show that IdeaWall effectively presents visual cues to facilitate verbal creative collaboration for idea generation and sets the stage for future research on intelligent systems that assist collaborative work.

ACM Classification Keywords

H.5.3. Information Interfaces and Presentation (e.g. HCI): Group and Organization Interfaces

Author Keywords

Verbal Collaboration; Brainstorming; Visual Cues; Groupware.

INTRODUCTION

Brainstorming is a form of collaboration by which efforts are made towards problem solving. A list of ideas spontaneously contributed by the participants is gathered over the course [28]. In the phase of ideation, *collecting*, *navigating* and *communicating* information play important roles for creative thinking [31]. Extensive research efforts in computer-supported cooperative work (CSCW) have been made toward designing *groupware* systems that help engage participants in a joint task by increasing interpersonal awareness or coordinating information sharing during collaborative activities (*e.g.*, [12, 21]).

The success of brainstorming can be evaluated through the factors of meeting *process constructs* and meeting *outcomes* [7]. Meeting process refers to communication and teamwork among participants that originate from group dynamics [19]. As an effective augmentation of collaboration, visualization has a proven track record of providing beneficial impacts on discourse in online communications (*e.g.*, [33, 32]), which illustrates individual participation to facilitate social interaction. Although the participant-centered method presents an overview of the individual activities, relevant information exchanged during the conversation is absent. Therefore, we apply a content-centered lens to visualize the semantics of brainstorming content with a purpose of improving the participants' awareness of their progress.

Meeting outcome, on the other hand, refers to the quantifiable results regarding efficiency. One approach shown to effectively improve brainstorming result is presenting conversation-based stimuli (*e.g.*, text [1] or image [34]). This technique enhances the ideation phase with cognitive strategies. To expand existing techniques, our study tests the effect of different combinations of visual cues on brainstorming's success.

Figure 1. A concept art of IdeaWall.

*The work was done while the authors are in residence in UC Davis.

CSCW '17, February 25-March 01, 2017, Portland, OR, USA
© 2017 ACM. ISBN 978-1-4503-4335-0/17/03. . . $15.00
DOI: http://dx.doi.org/10.1145/2998181.2998208

The illustrative scenario, as shown in Figure 1, depicts how a visualization may contribute to communication and productivity improvement. The visualization which conveys a variety of information to brainstorming participants can inspire them to explore new ideas. These inspirations break the ice when the team experiences a creative slump and maintain the idea generation cycle.

In this paper, we present IdeaWall, a real-time system that supports brainstorming by providing visual cues derived from an on-going conversation. We first develop a set of cognitive strategies based on human cognition mechanisms: Semantic Refinement, Pictorial Activation, and Associative Grouping. Semantic Refinement captures keywords from the conversation. By handling information that short-term memory is no longer storing, it minimizes cognitive load and redirects focus towards creative thinking. Pictorial Activation uses cognitive stimulation, attaching information-dense visual media to keywords. The activation of a visual stimulus in a cognitive network can propagate to connected nodes, leading to the generation of new ideas. Associative Grouping structures related concepts into clusters according to associative theory. Through an effective organization of ideas, thoughts can be easily discussed, evaluated, and merged. We then apply the above strategies to the construction of a proof-of-concept application that leverages speech recognition (SR), information extraction (IE) and Natural Language Processing (NLP). A pilot study was first conducted to collect preliminary feedback on visual interface design alternatives. The final system employs three display modes, namely, Keyword Matrix, Caption Matrix, and Dynamic Cells, which correspond to the three cognitive strategies, respectively.

To evaluate the effectiveness of IdeaWall, a user study with 24 participants was conducted. Users became accustomed to the visualization system and quickly integrated it into their idea generation process. Our study results show that instantly presenting visual cues as a collaborative aid leads to a better conversational flow during brainstorming. The qualitative feedback suggests that the majority of users preferred our system which acts as a source of inspiration. We discuss various design implications on developing real-time systems to provide a better means of capturing and displaying collaborative practices.

RELATED WORK

Facilitating Conversation with Visualizations

Conversation visualization highlights salient information and helps participants understand the social structure of the discussion [10]. For example, PeopleGarden [37] reveals the social structure of a conversation by visualizing patterns which exhibit the arrival of new participants and progression of activities. Conversation clock [2] displays individual contribution using participant audio input augmented with interaction information and conversation attributes. While much research has gone into participant-centered visualizations, others have used a content-centered lens. Specifically, Conversation Clusters [3] visually summarizes conversations by showing clusters of words grouped by topics and the evolution of topics over time. Our work focuses on visualizing semantics of a conversation by recording ideas out of brainstorming.

Generating Better Ideas using Content-Driven Media

When fostering idea generation, a stage of problem solving and decision making [27], multiple input sources can help elicit profound thought. With appropriate technology assistance, external stimuli may be used to enhance performance. InspirationWall [1] displays concepts which relate to the conversation, enriching ongoing idea generation. Similarly, Nguyen *et al.* [26] introduces a system that uses ranking recommendations to generate personalized topic suggestions during a conversation. While some have used textual cues, others have emphasized on presenting pictorial stimuli. For example, IdeaExpander [34] improves upon existing group brainstorming paradigms by adding a picture channel to a chat window. After the conversation's semantic content is analyzed, it is shown on the visual channel along with related images. Lewis *et al.* [20] apply digitally embedded image stimuli which affect users. Their results suggest that positive emotional priming improves the quality of ideas. IdeaWall augments idea generation with related pictorial stimuli and effective concept organizations to improve creative collaboration.

Maintaining Group Awareness in the Workspace

A well-designed shared workspace is essential to collaboration. Presenting information on a display maintains group awareness of the current task and widens the breadth of discussion with additional perspectives while reducing communication ambiguities and information distortion [18, 9]. Dynamo [15] supports meetings in locations such as outdoor venues by exchanging and distributing live information among devices and screens. A multi-surface collaboration space, WeSpace [36], allows simultaneous visual exploration coalesced from multiple data sources. Otmar *et al.* [14] design an electronic brainstorming system which uses two different interactive surfaces; a tabletop surface that facilitates writing and group cohesion and a wall display that provides an overview to maintain context awareness. IdeaWall uses a touch screen to help users interactively explore ideas in terms of words or images.

DESIGN

Our goal is to design a real-time visualization system that stimulates collaborative creative meetings. Grounded in human cognition theories, three cognitive strategies were developed to guide the design process. Corresponding layouts were made to evaluate the effect of the strategies on how visual stimuli affect users' responses.

Cognitive Strategies

After reflecting on information assimilation and extensive examination of background literature, we come to three cognitive strategies that can be used to facilitate idea generation: Semantic Refinement, Pictorial Activation, and Associative Grouping.

Semantic Refinement - Capturing Keywords for Review

The Semantic Refinement strategy captures and displays keywords extracted from the conversations during a meeting. As reported by the 3M Study [24] as well as Mosvick and Nelson

[25], up to fifty percent of the time spent in a meeting was wasted due to information being lost or distorted, suboptimal decision making, and meeting mismanagement. The reason is that humans have limited capability for short-term memorization of information chunks [23]. A typical technique for minimizing cognitive load is externalization [16]. Semantic Refinement utilizes the participants' spatial memory and records the flow of keywords on a medium accessible to the entire group. Each keyword can be assigned with a degree of significance, correspond to its meaning, frequency, and surrounding context. Semantic Refinement reduces the cognitive load on users' short-term memory, and then redirects the focus towards creative tasks such as association and reflection.

Pictorial Activation - Stimulating Knowledge Using Pictures

The Pictorial Activation strategy provides content-based image stimuli to the problem-solvers. According to Nijstad, Stroebe [27], and Wang [35], the concept of cognitive stimulation suggests that using pictures as stimuli constitutes a strategy for presenting multiple concepts at once, through the simultaneous delivery of visual components. A single picture contains multiple sources of stimuli, including color, shape, context, and more [35]. Each stimulus acts as a new entry node into the cognitive network [30]. A node in the cognitive network represents a perceptual concept. The activation of an entry node will propagate to other connected nodes in the cognitive network, and lead to the generation of new ideas. By representing a conversation using pictures, Pictorial Activation allows the provision of additional information-dense stimuli, then leads thinkers to new and alternative explorations of their knowledge.

Associative Grouping - Structural Organization of Concepts

The Associative Grouping strategy structures related concepts into clusters. It is derived from the concept of categorization, through which ideas and objects are recognized, differentiated, and understood [5]. According to Mednick's [22] associative theory, creating new connections between previously unrelated concepts and bringing associative elements into ideational contiguity increases the speed and chance of arriving at a creative solution. For example, in brainstormings, one common method of developing ideas is to aggregate similar concepts such as circling related ideas or putting ideas written on post-its next to each other. Thoughts can then be easily discussed, evaluated, and merged. Associative Grouping adds meaningful structure to otherwise scattered information, producing a more systematic and organic construction of ideas.

Design Approach

To demonstrate the feasibility of cognitive strategies, we equipped the real-time meeting visualization system with three display modes, namely: Keywords Matrix, Captioned Matrix, and Dynamic Cells. The layout design follows perceptual principles drawn from a class of theories known as the Gestalt laws of grouping. By utilizing the grouping methods while maintaining a straightforward aesthetic, the design improves readability and minimizes distraction. The design principles are as follows:

DP1 Each display mode addresses features of one cognitive strategy.

DP2 Each display mode incrementally builds up from previous ones; consistency is maintained across all modes.

DP3 Each display mode contains intuitive and minimalistic visual components.

Keywords Matrix - A Grid of Significant Phrases

The Keywords Matrix mode addresses the Semantic Refinement strategy by presenting keywords extracted from conversational content, as shown in Figure 2. Keywords are arranged in a 3×7 matrix. The layout is symmetric about the x-axis and y-axis. When a new keyword is captured, it is placed in the central position. Replaced words are moved to one of the eight adjacent positions in clockwise order, and then pushed toward outermost columns. If no space available in that column, the keyword with the lowest significance is discarded and replaced. The font size of a keyword indicates its significance. As a result, the more important keywords are visually striking and persist for a longer duration. These design decisions are made based on Semantic Refinement strategy that the displayed keywords help users maintain their train of thought and externalize cognitive load.

Captioned Matrix - Pictures Attached to Keywords

The Captioned Matrix mode supplements Keywords Matrix mode with Pictorial Activation through the attachment of related pictures to conversationally-retrieved keywords, as shown in Figure 3. We keep the 3×7 matrix layout and use

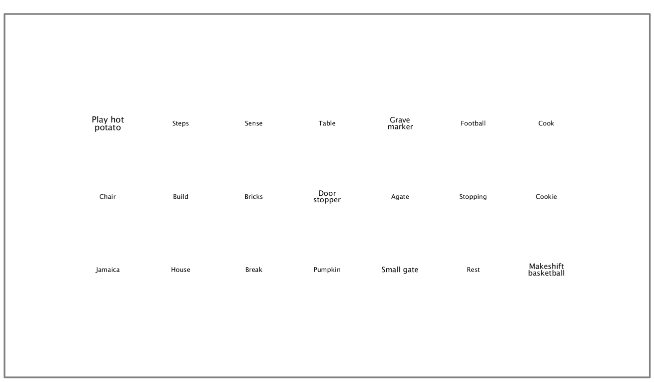

Figure 2. A Snapshot of Keywords Matrix Mode

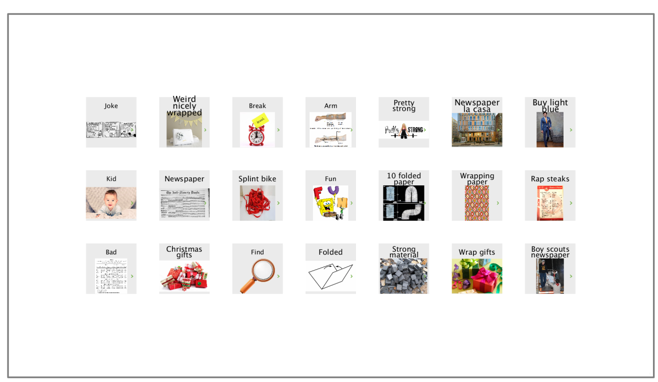

Figure 3. A Snapshot of Captioned Matrix Mode

a rectangular cell for each unit in the matrix. Same as the previous mode, the most recently added cell is displayed in the center, and other cells are updated following the same rule as the Keywords Matrix mode. All cells are of identical size with the keywords as captions and associated pictures as visual contents. The font size of the caption is used to indicate the significance of the keywords. Each cell has up to ten picture candidates with a green arrow indicator showing if there are additional pictures to explore. Users could interact with the visualization to cycle through pictures for given keywords and are encouraged to do so to benefit from the Pictorial Activation strategy. Through the introduction of pictorial stimuli, this mode allows users to investigate additional resources beyond what was said during conversations.

Dynamic Cells - Clustering Structure Based on Association

The Dynamic Cells mode combines the three strategies of Semantic Refinement, Pictorial Activation, and Associative Grouping by displaying keywords with related pictures and representing relations amongst them through clustering, as shown in Figure 4. We apply a force-directed layout to utilize its grouping features. Clusters of cells are organized according to their semantic similarity. Large clusters or those with a high degree of similarity to other clusters gravitate toward the center. If an incoming cell is associated with an existing one, it would be placed within the same cluster using the Ulam spiral method. Otherwise, it forms a new cluster on its own. Up to 21 cells are presented on display for maintaining consistency. When the number of cells exceeds the maximum amount, the oldest cluster fades out of view. We measure each group's age by recording the time of its most recent addition. By grouping keywords based on similarity, this mode helps organize ideas within existing clusters and prompts new categories of thoughts.

Implementation

The architecture of IdeaWall contains four backend components: Speech Recognition, Keyphrase Extraction, Image Retrieval, and Similarity Calculation (see Figure 5).

Speech Recognition The speech recognition component of our prototype system records participants' speech, generates a vocal input stream, and returns with a text output stream.

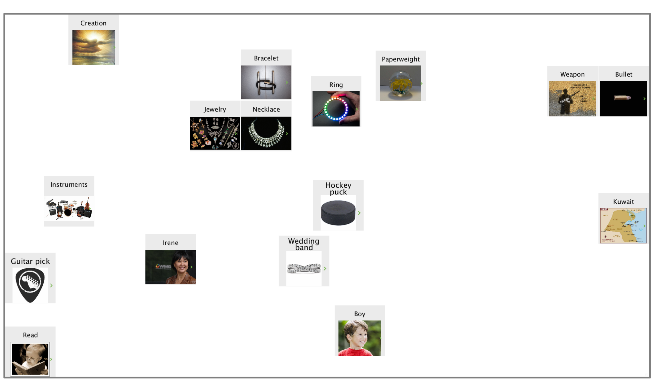

Figure 4. A Snapshot of Dynamic Cells Mode

We conducted tests on the state-of-the-art speech recognition services: Bing Speech API[1], CMU Sphinx[2], and Google Web Speech API[3]. We chose Google Web Speech API because it demonstrates good performance in the following facets: fast recognition speed with low delays where latency ranges between 50ms to 100ms; efficient message communication with the keyphrase extraction component; and high recognition accuracy when used with the dominant local dialect (*i.e.*, in the optimal environment with no background noise, the recognition rate for users with the Standard American accent were above 93.5%).

Keyphrase Extraction Keyphrase extraction performed by IdeaWall is conducted using the RAKE algorithm [29]. After removing stop words and punctuation, RAKE chunks the remaining tokens in the sentence according to Part-of-Speech tags, leaving only nouns as keyphrase candidates. Then, it builds a word co-occurrence graph, scores its significance determined by the quotient of its word degree and word frequency, with emphasis on words predominantly appearing in longer phrases.

RAKE stands out against other complex models (*e.g.*, SVM classification in [34]) in terms of speed, which causes nearly zero delays while maintaining decent accuracy, especially for lengthy input sentences. Besides, RAKE requires few or no training sets. This feature enables generating acceptable results even at the beginning of a test case when the accumulated textual input is not sufficient.

Image Retrieval To implement the back-end image retrieval service, which is initiated after search terms have been determined, we adopted the Custom Search API from Google[4]. This API provides comprehensive and precise results within an acceptable period (*i.e.*, around 300ms for any given request). We also customized the search filter to optimize the applicability of each image to the topic (*e.g.*, set emphasized sites and restrict results from certain fields).

When compared to using a self-built local dataset as in [34], our real-time retrieval component can respond to a variety of requests. By instantly collecting and presenting information, it allows exploration of a much larger database and increases of user engagement.

[1]https://www.microsoft.com/cognitive-services/en-us/speech-api
[2]http://cmusphinx.sourceforge.net
[3]https://dvcs.w3.org/hg/speech-api/raw-file/tip/speechapi.html
[4]https://developers.google.com/custom-search/

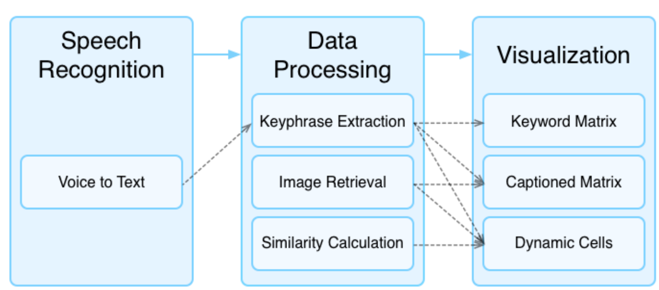

Figure 5. System Architecture

Similarity Calculation IdeaWall calculates the similarity between two key phrases using the normalized cosine distance between word2vec[5] word vector representations trained from the English Wikipedia corpora in advance. Once a new key phrase is detected, its most similar counterpart among existing key phrases is computed, and an edge is defined based on the similarity of this pair.

Word2vec is a neural network model that can capture most of the regularities in the training dataset. It maintains better computational efficiency in comparison with other distributional representations (*e.g.*, ESA model in [3]).

EVALUATION

The evaluation process began with a small, limited pilot study designed to provide early-stage feedback. It was then followed by a formal user study which determines the effectiveness of the presented visualizations.

Pilot Study

In the pilot study, IdeaWall's functionality was tested to collect preliminary results in preparation for the main user study. A total of 10 participants aged 20 to 28 took part in five pairs. We explored a few design alternatives in the pilot study. For example, the proper number of items to be shown in display modes were tested (*e.g.*, 1×11 matrix, 3×7 matrix, 5×5 matrix). We decided upon the layout based on user requirements to ensure they feel neither overwhelmed by the amount of text nor wanting for more information. In addition, participants were presented with a word cloud in early design phase. Through our observations, participants indicated a preference for additional pictorial information, which is more eye-catching and visually rich. Based on the design principles and participants' feedback, we reduced the added complexity between display modes to conduct a fair comparative study.

User Study

After modifying the visual design, we conducted a laboratory study to investigate how IdeaWall facilitates collaborative creative meetings. We posed three research questions that we aimed to answer:

[5]https://radimrehurek.com/gensim/models/word2vec

Figure 6. Verbal Collaboration Assistant in Action

Table 1. Factors used in the study

Type	*ID*	*Description*
Condition	M0	No visual cues
	M1	Display mode: Keywords Matrix
	M2	Display mode: Captioned Matrix
	M3	Display mode: Dynamic Cells
Task	B	List unique uses for Bricks
	N	List unique uses for Newspapers
	C	List unique uses for Coffee mugs
	Q	List unique uses for Quarters

RQ1 How does the system influence brainstorming?

RQ2 How well do the three display modes of the visualization perform?

RQ3 How do participants integrate the system into their discussion?

Participants

We recruited 24 participants (8 females) aged 18 to 31 (mean 24). We paired them into 12 groups (participants sometimes knew each other before the user study). We used pairs in order to optimize team synergy, duration of speaking time, and user engagement. All participants are fluent American English speakers with experience conducting brainstorming meetings. Participants included college students, researchers, and professionals from various backgrounds such as computer science, psychology, electrical engineering, biology, neurology, and music.

Apparatus

Our user study was conducted using an iMac with 2560×1440 display connected to a 27-inch touch screen. The iMac was used as the backend server for speech recognition, data processing, and monitoring logs generated during the experiment. The participants were able to interact with the touchscreen and explore image candidates when desired. The brainstorming processes were recorded using a camera and on-screen interactions were recorded using a screen capture tool for later references. A studio condenser microphone was used to record the conversation (see Figure 6).

Methodology and Tasks

We conducted a within-subject study and employed three display modes of IdeaWall (*i.e.*, Keywords Matrix, Captioned Matrix, and Dynamic Cells) as experimental conditions. The three experimental conditions were counterbalanced using a Latin Square while a control condition with no visualization was always presented first. In each of the four conditions, participants were instructed to conduct a five-minute Guilford's Alternative Use Task [13]. Guilford's task has been used by previous brainstorming studies (*e.g.*, [20]) as an activity generalizable to other creative meetings. The order of the tasks with different topics was also counterbalanced to minimize learning effect (see Table 1).

Table 2. Measures used in the study

	Type	*ID*	*Description*
Performance	Group Dynamics	ΔLull	Total duration of lulls in conversation
		#Lull	Quantity of lulls in conversation
	System Engagement	ΔStare	Total duration spent looking at screen
		#Click	Quantity of clicks on screen
	Creative Output	#Idea	Quantity of ideas
		∼Idea	Similarity of ideas
Experience	Usefulness	Q1	The keywords capture the essence of the ideas.
		Q2	The images are stimulating.
		Q3	The organization of the keywords are helpful.
	Enjoyability	Q4	Overall, the visual cues are helpful.
		Q5	Generally, we had a productive discussion.

Procedure

The participants were brought to the laboratory and given a brief introduction to the purpose of IdeaWall. They were instructed about the task and provided with the four conventional brainstorming rules [28] before the study: 1) the more ideas, the better; 2) the wilder ideas, the better; 3) improving or combining ideas are better; and 4) do not be critical. Each group received instructions (adapted from [13]) at the beginning of the session to complete a modified Guilford's Alternative Use Tasks, as follows:

In this task, your goal is to think of as many unique and unusual uses for a common object. For example, using a paperclip as an earring is an unusual and unique use. However, using a paperclip to bind papers is not unique or unusual. Try to think of as many unique and unusual uses as possible.

Each group first tested the control condition M0 to warm up and became accustomed to the task, and then completed the remaining three experimental conditions, M1, M2, and M3. Before each task, participants were given a short introduction of the visual cues that will be shown on the screen and the interaction types that were supported in this display mode. Participants were free to inspect or interact with the visualization when they wished. To allow comparison between different conditions, a questionnaire was administered when users completed four conditions at the end of the study.

Measures

We evaluated our prototype based on measurements from two perspectives: *performance* and *experience*.

Performance We analyzed two quantitative aspects: meeting process attributes (*e.g.*, teamwork) and meeting outcomes (*e.g.*, efficiency) [7]. Two external coders were asked to review video recordings of the participants' conversations and system interactions. First, the coders gauged the meeting process between users by recording the verbal facet of *group dynamics* and the physical facet of *system engagement*. In particular, we define group dynamics as the degree of participation, in which silent lulls is used as the main measure. While lulls give group members time to formulate ideas, they also mark a lack of motivation and energy [17]. In our user study, we determined that lulls often coincide with low creative energy in a group. Extended silence from both members indicates that they are having difficulty collaborating. Afterward, the two coders measured the overall quality of the users' results in *creative output*. (see Performance section in Table 2)

Experience To qualitatively evaluate user experience, we designed a questionnaire consisting of five questions. Responses to the questionnaire were provided on a 7-likert scale (1 = strongly disagree; 7 = strongly agree). Questions Q1 to Q3 were intended to examine the *usefulness* of each display mode, whereas Q4 and Q5 focused on the users' overall *enjoyability* (see Experience section in Table 2).

RESULTS AND ANALYSIS

We next report the results of our study and interpret them to answer RQ1-3 listed earlier.

RQ1: How does the system influence brainstorming?

Improvements in Conversational Engagement We assessed the total duration of conversational lulls as ΔLull, and their numbers with #Lull in the conversation across all conditions. This is used as an indicator of group dynamics. One-way ANOVA tests suggest that there exist significant differences in ΔLull ($F_{3,44}$=6.165, $p<0.005$) and #Lull ($F_{3,44}$=11.104, $p<0.0001$). Post-hoc Tukey test further reveals that M1 (ΔLull: $\mu = 25.36$, $\sigma = 31.02$, #Lull: $\mu = 2.64$, $\sigma = 2.46$), M2 (ΔLull: $\mu = 16.82$, $\sigma = 39.75$, #Lull: $\mu = 1.09$, $\sigma = 2.02$), and M3 (ΔLull: $\mu = 7.09$, $\sigma = 14.04$, #Lull: $\mu = 1.00$, $\sigma = 1.48$) significantly outperform M0 (ΔLull: $\mu = 60.91$, $\sigma = 34.33$, #Lull: $\mu = 5.91$, $\sigma = 2.91$). These findings show that IdeaWall is capable of reducing the duration and amount of lulls during a meeting. The use of visual cues result in a marked improvement in collaboration dynamics.

Benefits of System Interaction To further investigate how the effect of system interaction on brainstorming dynamics, we analyzed the total duration spent staring at screen ΔStare and quantity of clicks #Click on screen. We found a significant negative correlation between ΔStare and #Lull in display modes which contain images (M2: $R = -0.679$, $p<0.05$, M3: $R = -0.671$, $p<0.05$), suggesting that interacting with IdeaWall can facilitate group dynamics. When conveying information, pictorial stimuli act as an effective complement to the text, improving cognitive performance.

We looked at quantities of ideas #Idea and their similarity ∼Idea to analyze creative performance (∼Idea between each pair of ideas was computed using the method described in Similarity Calculation section). We observed no significant enhancement within the creative dimension, implying that our system may not be able to alter users' inherent creative abilities significantly. The reason may be that by tasking the users with developing unusual uses for familiar items, we are not fully exploring the users' creativity. As one of the participants suggested, "*I'd like to try more difficult and challenging topics*" (P13). Thus, we suggest that an alternative scenario of using IdeaWall can be improvising uses for unfamiliar objects.

RQ2: How well do the three display modes perform?

Keywords Tracking Aids Review and Recall A majority of the participants found keywords helpful (see Figure 7). "*Not all the [captured] keywords made sense, but a lot of them were helpful... You can look back at words when you are stuck. It helps when relating and connecting ideas*", stated by P6. "*I didn't have to stop to recall what we came up with, they were already there*" (P23). Some of the participants suggested adding interesting features, "*it might work better if you had clusters and showed connections among those [keywords]*" (P19).

Images Inspire Thought and Reduce Ambiguity Many participants commented that the retrieved images were "*inspiring*".

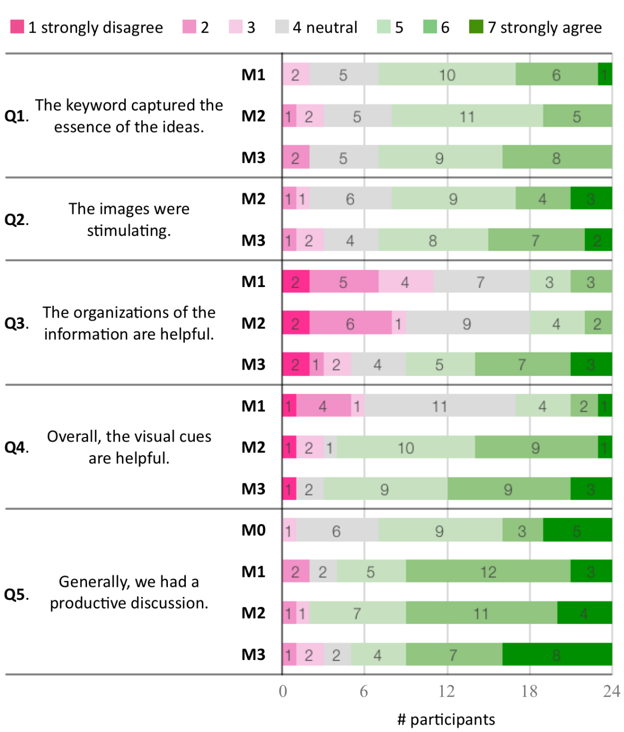

Figure 7. Responses to five evaluation questions illustrated in a stacked bar chart.

"*Watching them [pop up] while brainstorming is fun*" (P2, P14, P17). "*I like colors, so it's helpful to me. Color works well. I think it's easy to associate an object with the color after you see this*" (P17). In addition, participants occasionally refocused their conversation by using the visual cues. "*Oh, I thought you were referring to... Yes, that would make sense"* (P7). "*Yeah, that's what I was talking about, look... "* (P20).

Organized Thoughts Improve Cue's Effectiveness Clustering structure is described as an "*interesting*" feature to assist participants' idea generation process. "*Small clumps in level three offered ideas*" (P2). "*When I was trying to come up with ideas, I moved one cluster to another"* (P4). "*When two keywords are combined, it helps trigger new ideas*" (P10). Participants thought the layout reduces mental strain. "*It takes time if I had to sort [the information] by myself, this [automatic sorting] allows me to think about new ideas*" (P9). "*Sometimes [the captioned pictures] moved around too much, that could be distracting. But it helps by shuffling ideas, and it's interesting to the eye*" (P5).

When compared to other modes, Dynamic Cell mode received more feedback regarding potential changes. "*Mode three highlights related images and ideas. There is room for more features like you could make the cluster selectable, make the keywords indicate high-level ideas*" (P3). "*I wish I could delete useless clusters or create a new one by myself*" (P2). "*You may consider using the existing clusters to generate new blocks. That could be very interesting. It gives ideas even though I didn't say a word about it*" (P14).

RQ3: How do participants integrate the system into their discussion?

Images Augment Text Cues Most groups expressed a preference for pictures over text. As remarked by P3, P14, P19, and P22: "*Pictures convey more*". When asking about the order of preference when comparing three display modes, a majority of the participants indicated that they prefer the modes with images. "*I definitely like level two and three the best, they're more eye-catching, though level one could be helpful somehow*" (P1). Participants also suggested various design alternatives, such as "*bigger images*" (P3, P7, P16), "*more picture candidates*" (P13), and "*self-cycling pictures*" (P5, P9).

Mismatched Cues may Give Unexpected Benefits Several participants suggested that even mismatched keywords sometimes provide benefits. "*It's interesting that the computer picked up words we didn't say... Keywords that are incorrect give new ideas*" (P5, P11). Participants noted the benefits of the mismatched pictures, as such pictures could deliver unanticipated but helpful information and then prompt alternative discussions. "*It's actually more helpful when images didn't match [the keyword]*" (P1, P6, P18). P16 added: "*I found unrelated words add confusion but photos don't... I'd prefer more variance in photos, like for 'cat', show me some cat toys instead of three photos of ordinary cats*".

Various Uses Based on Collaborative Style Groups used the visualization system differently based on their communication and collaboration styles. For example, P6 from group 3 remarked that "*we were so good at this that I didn't really look at the screen*", while P12 from group 6 stated, "*my partner didn't say much. It was a lot harder to come up with ideas on my own. I had to look at the screen for inspiration now and then. It's extremely helpful when you're running out of ideas.*" Participants tended to have very different ratings and opinions about their user experience based on the groups' performance. How different group styles affected their collaboration is discussed later on.

DISCUSSION

In this section, we discuss design implications to help set the stage for future development of real-time collaborative systems.

Mismatched Queries and Idea Generation

We observed mismatched queries in both keywords and images due to inaccuracy inherent within the speech recognition and information extraction components of IdeaWall. These inaccurate queries interfere with users' thought processes and result in distraction. The accuracy of queries could be more critical for tasks that involve organizational logic and pattern recognition. For example, if participants are told to "*determine the most effective use of bricks*" instead of "*think of as many uses of bricks as possible*", the system should maintain high accuracy to keep track of the items that influence the decision-making process. This way, participants can propose related ideas more quickly, as the distraction is minimized by providing more focused cues.

Although inaccuracy poses undesirable factors, we found that users are able to draw inspiration from some mismatched cues and use them to achieve improvements. For example, during a user study session where the task was brainstorming unique uses for bricks, the system captured "*break*" when one of the participants said "*brick*". This inaccurate result caught the participants' attention, and they immediately came up with the idea to "*break bricks to make sculptures*". In this case, discrepancies between the spoken and captured words result in the activation of the train of thought. Similarly, when a picture differs from the spoken word, it inspires users to explore new concepts along that track.

Mismatched results can be leveraged to generate a wider variety of cues by promoting lateral thinking [8]. Instead of continuing to move in familiar directions, lateral thinking takes off laterally to a new and innovative direction using random inputs (*e.g.*, word, picture, and sound). We suggest considering lateral stimulus, which can range from similar words in pronunciation to synonyms or antonyms. For example, when the keyword "*cat*" is captured, rather than presenting information on mundane cats, a lateral approach would present results derived from varying classifications, such as "*cartoon cat*", "*Catwoman*", or even "*dog*". These loosely related concepts reap additional benefits by delivering unanticipated but helpful information from a unique perspective.

Use Context to Mark Areas of Interest

To further improve the idea generation process, multiple contextual dimensions can be attached to each of the ideas generated in brainstorming. Contextual data of the raw audio input can be used to depict interesting conversational features, such as emotional states and volumes (*e.g.*, [6]). For example, a participant might be excited when expressing his ideas in a loud voice. The context suggests that such ideas of interest are valuable to explore and review in post-meeting session.

Log data that shows participants' interaction with the system can also be collected to track areas of interest (*e.g.*, [4]). Users may exhibit certain tendencies regarding different stimuli. In our case, some participants actively cycled through pictures for clues while others tried to re-organize idea cluster. These engagement patterns can be found through the analysis of mouse click and eye movement data. Visualization of log data helps determine the components (*e.g.*, text, image) that attract more attention and achieve greater effects, thus providing new insights into the effectiveness of IdeaWall.

Identify Individual Contribution

The current design of IdeaWall focuses on analyzing the semantics of brainstorming content, its functionality could be expanded by using various participant-related information to organize visual cues. For example, methods revealing individual participation such as speaker recognition and opinion extraction can be used to better understand the efficacy of meetings. Differentiation between speakers in a conversation could alleviate distraction by separating clusters based on individuals. This speaker-based clustering organizes thoughts in a way that augments depiction of individual characteristics and supports reflection on group productivity.

Other mentioned venues for enhancement include adding argumentative dimension. For example, within a topic discussion, there may be a conflict variation, as some people disagree with others. It would be beneficial to capture this argumentative divide in ideas, that is, whether all participants agreed on the same decision, whether there were two sides discussing over a divide, or whether the topic was a free-for-all of opinions. Further studies could be conducted to compare different clustering methods based on group criteria. By highlighting thought process in each camp, IdeaWall acts as a research tool to assess how different participants contribute to certain ideas or design decisions.

Make Passive Displays Participate

The design of interactive systems should consider the manifestation of user behavior. In our deployment of IdeaWall, collaborations between participants pose different forms dependent upon group chemistry and dynamics. Among the different styles of teams, groups with long idle times between ideas relied heavily on the visual support interface to receive inspiration. On the other hand, highly productive pairs with better rapport preferred to focus on the ongoing conversation. They used the visualization if they got stuck or ran out of ideas. For both collaboration styles, the system proved most effective when conversations reach a lull. In addition, participants' roles play a large part in determining how they interact with the system. Participants who were listeners were able to complete additional tasks. They may retrieve clues with the picture cycling feature of the visualization while directing the speaker's attention toward valuable leads. The active speakers were less likely to use the interface until they had completed their current thought.

An advanced real-time intelligent system could engage in a conversation as a knowledgeable participant; it gives different responses based on specific situations, whether active, reactive, or passive. For example, it suggests a new direction to explore when a collaborative discussion reaches a deadlock. It gives feedback and suggestions accordingly when other participants are presenting their thoughts. Through the use of situation-based strategies, further development of such systems could be designed to present selected cues based on the perceived participants' intentions (*e.g.*, by mining a broad knowledge base of human behavior [11]).

Limitations and Future Work

There were several limitations discovered in our study that leave room for future improvement. By adopting group brainstorming solutions for Guilford's Alternative Uses Tasks, we chose to task users with objectives that could be fulfilled immediately. Meanwhile, the brainstorming was conducted using pairs in order to optimize team synergy and user engagement. However, evaluating Guilford's task on pairs with limited sample size makes it difficult to generalize the findings to an open brainstorm of a group. We suggest evaluating the tool in a more realistic context, such as that of the major academic discussions or design agency meetings. Developing an ideal scenario would allow us to draw stronger conclusions regarding the effectiveness of our system for larger, more complex meetings.

Further, the design of clustering in Dynamic Cell uses a force-directed layout. According to qualitative feedback, this dynamic organization is novel but could also be distracting due to its progressive movement. It is not clear if other designs might achieve better or worse results. For example, making use of Treemapping, Sunburst, and Sankey diagrams could visualize the information in more static ways. We suggest future studies focus on conducting a comparison between the different clustering layouts to obtain a deeper understanding of the visual forms most helpful in supporting collaborative creative activities.

CONCLUSION

This paper presents IdeaWall, a groupware system which comprehends human conversation and provides helpful visual aids. The state-of-the-art speech recognition and information extraction technologies were used to develop a proof-of-concept application. Using design strategies grounded in mechanisms of human cognition, the content-centered visualization with novel layouts act as an effective tool for improving verbal group discussions.

Through a laboratory study, our work was found to bestow several benefits: a majority of users preferred our system with instant visual feedback; pictorial cues were considered to be far more helpful than purely textual cues; and structuring information organically positively affected creative performance. Our study results indicate that IdeaWall provides a better means of capturing and displaying collaborative practices using visualization techniques. Our work also suggests areas for more studies on converting multi-party conversation data into effective visual aids to help inspire idea generation, identify individual contribution, and reflect on group productivity.

ACKNOWLEDGMENTS

This research is supported in part by the U.S. National Science Foundation via grants NSF DRL-1323214, NSF IIS-1528203 and NSF IIS-1320229, the U.S. Department of Energy through grant DE-FC02-12ER26072, and National Natural Science Foundation of China under grant No.61402540, No.61672538. We thank Yaoxue Zhang, Fangfang Zhou, Ying Zhao from Central South University and Yi-Ling Chen from University of California, Davis for their assistance. We also thank Justin Ma for creating the concept art.

REFERENCES

1. Salvatore Andolina, Khalil Klouche, Diogo Cabral, Tuukka Ruotsalo, and Giulio Jacucci. 2015. InspirationWall: Supporting Idea Generation Through Automatic Information Exploration. In *Proceedings of the ACM SIGCHI Conference on Creativity and Cognition*. 103–106.
2. Tony Bergstrom and Karrie Karahalios. 2007. Conversation Clock: Visualizing audio patterns in co-located groups. In *Proceedings of the 40th Annual Hawaii International Conference on System Sciences*. IEEE, 78–78.
3. Tony Bergstrom and Karrie Karahalios. 2009. Conversation clusters: grouping conversation topics through human-computer dialog. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*. 2349–2352.
4. Tanja Blascheck, Markus John, Kuno Kurzhals, Steffen Koch, and Thomas Ertl. 2016. VA2: A Visual Analytics Approach for Evaluating Visual Analytics Applications. *IEEE transactions on visualization and computer graphics* 22, 1 (2016), 61–70.
5. Henri Cohen and Claire Lefebvre. 2005. *Handbook of categorization in cognitive science*. Elsevier.
6. Andrew J Cowell, Michelle L Gregory, Joe Bruce, Jereme Haack, Doug Love, Stuart Rose, and Adrienne H Andrew. 2006. Understanding the dynamics of collaborative multi-party discourse. *Information Visualization* 5, 4 (2006), 250–259.
7. Robert Davison. 1997. An instrument for measuring meeting success. *Information & Management* 32, 4 (1997), 163–176.
8. Edward De Bono. 2010. *Lateral thinking: a textbook of creativity*. Penguin UK.
9. Joan Morris DiMicco, Anna Pandolfo, and Walter Bender. 2004. Influencing group participation with a shared display. In *Proceedings of the ACM conference on Computer supported cooperative work*. 614–623.

10. Judith Donath, Karrie Karahalios, and Fernanda Viegas. 1999. Visualizing conversation. *Journal of Computer-Mediated Communication* 4, 4 (1999).
11. Ethan Fast, William McGrath, Pranav Rajpurkar, and Michael S Bernstein. 2016. Augur: Mining Human Behaviors from Fiction to Power Interactive Systems. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*. 237–247.
12. Saul Greenberg and Michael Rounding. 2001. The notification collage: posting information to public and personal displays. In *Proceedings of the ACM SIGCHI Conference on Human factors in computing systems*. 514–521.
13. Joy Paul Guilford. 1967. *The nature of human intelligence.* McGraw-Hill.
14. Otmar Hilliges, Lucia Terrenghi, Sebastian Boring, David Kim, Hendrik Richter, and Andreas Butz. 2007. Designing for collaborative creative problem solving. In *Proceedings of the ACM SIGCHI Conference on Creativity & Cognition*. 137–146.
15. Shahram Izadi, Harry Brignull, Tom Rodden, Yvonne Rogers, and Mia Underwood. 2003. Dynamo: a public interactive surface supporting the cooperative sharing and exchange of media. In *Proceedings of the ACM symposium on User interface software and technology*. 159–168.
16. David Kirsh. 1995. The intelligent use of space. *Artificial intelligence* 73, 1 (1995), 31–68.
17. Jeffrey A Kottler and Matt Englar-Carlson. 2009. *Learning group leadership: An experiential approach*. Sage, Chapter Understanding Group Dynamics and Systems, 76–79.
18. Robert E Kraut, Darren Gergle, and Susan R Fussell. 2002. The use of visual information in shared visual spaces: Informing the development of virtual co-presence. In *Proceedings of the ACM Conference on Computer supported cooperative work*. 31–40.
19. Daniel Levi. 2016. *Group dynamics for teams*. Sage Publications.
20. Sheena Lewis, Mira Dontcheva, and Elizabeth Gerber. 2011. Affective computational priming and creativity. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*. 735–744.
21. Joseph F McCarthy, Ben Congleton, and F Maxwell Harper. 2008. The context, content & community collage: sharing personal digital media in the physical workplace. In *Proceedings of the ACM Conference on Computer supported cooperative work*. 97–106.
22. Sarnoff Mednick. 1962. The associative basis of the creative process. *Psychological review* 69, 3 (1962), 220.
23. George A Miller. 1956. The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological review* 63, 2 (1956), 81.
24. Peter R Monge, Charles McSween, and JoAnne Wyer. 1989. *A profile of meetings in corporate America: Results of the 3M meeting effectiveness study*. Annenberg School of Communications, University of Southern California.
25. Roger K Mosvick and Robert B Nelson. 1987. *We've got to start meeting like this: a guide to successful business meeting management*. Scott Foresman Trade.
26. Tien T Nguyen, Duyen T Nguyen, Shamsi T Iqbal, and Eyal Ofek. 2015. The Known Stranger: Supporting Conversations between Strangers with Personalized Topic Suggestions. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*. 555–564.
27. Bernard A Nijstad and Wolfgang Stroebe. 2006. How the group affects the mind: A cognitive model of idea generation in groups. *Personality and social psychology review* 10, 3 (2006), 186–213.
28. Alex F Osborn. 1953. *Applied imagination.* Scribner's.
29. Stuart Rose, Dave Engel, Nick Cramer, and Wendy Cowley. 2010. Automatic keyword extraction from individual documents. *Text Mining* (2010), 1–20.
30. Eric L Santanen, Robert O Briggs, and Gert-Jan De Vreede. 2000. The cognitive network model of creativity: a new causal model of creativity and a new brainstorming technique. In *Proceedings of the Hawaii International Conference on System Sciences*. IEEE.
31. Ben Shneiderman. 2007. Creativity support tools: Accelerating discovery and innovation. *Commun. ACM* 50, 12 (2007), 20–32.
32. Annie Tat and M Sheelagh T Carpendale. 2002. Visualising human dialog. In *Proceedings of International Conference on Information Visualisation*. IEEE, 16–21.
33. Fernanda B Viégas and Judith S Donath. 1999. Chat circles. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*. 9–16.
34. Hao-Chuan Wang, Dan Cosley, and Susan R Fussell. 2010. Idea Expander: Supporting group brainstorming with conversationally triggered visual thinking stimuli. In *Proceedings of the ACM conference on Computer supported cooperative work*. 103–106.
35. Hao-Chuan Wang, Susan R Fussell, and Dan Cosley. 2011. From diversity to creativity: Stimulating group brainstorming with cultural differences and conversationally-retrieved pictures. In *Proceedings of the ACM conference on Computer supported cooperative work*. 265–274.
36. Daniel Wigdor, Hao Jiang, Clifton Forlines, Michelle Borkin, and Chia Shen. 2009. WeSpace: the design development and deployment of a walk-up and share multi-surface visual collaboration system. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*. 1237–1246.
37. Rebecca Xiong and Judith Donath. 1999. PeopleGarden: creating data portraits for users. In *Proceedings of the ACM symposium on User interface software and technology*. 37–44.