InspirationWall: Supporting Idea Generation Through Automatic Information Exploration

Salvatore Andolina[1], **Khalil Klouche**[1,3], **Diogo Cabral**[1], **Tuukka Ruotsalo**[2], **and Giulio Jacucci**[1,2]
[1]Helsinki Institute for Information Technology HIIT,
Department of Computer Science, University of Helsinki, Finland
[2]Helsinki Institute for Information Technology HIIT, Aalto University, Finland
[3]Helsinki Media Lab, Aalto University, Finland

ABSTRACT
Collaborative idea generation leverages social interactions and knowledge sharing to spark diverse associations and produce creative ideas. Information exploration systems expand the current context by suggesting novel but related concepts. In this paper we introduce InspirationWall, an unobtrusive display that leverages speech recognition and information exploration to enhance an ongoing idea generation session with automatically retrieved concepts that relate to the conversation. We evaluated the system in six idea generation sessions of 20 minutes with small groups of two people. Preliminary results suggest that InspirationWall contrasts the decay of idea productivity over time and can thus represent an effective way to enhance idea generation activities.

Author Keywords
Idea generation; Information Exploration; Automatic Speech Recognition.

ACM Classification Keywords
H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

INTRODUCTION
Collecting and navigating through information is an important phase in creative processes [13], which fosters associative and inspirational learning [2]. Previous work that sought to support for example brainstorming referred to the semantic network structure of human memory, where concepts feature as nodes with associative links [15]. In brainstorming, one cognitive operation to generate ideas is to retrieve concepts from associative memory. Expanding the current context of topics has been investigated through topic suggestion algorithms designed to generate candidate topics that are novel but related to the current context [9]. As brainstorming is often a collaborative practice, recent creativity systems support groups. Groups generally perform better than individuals in a variety of tasks [6]. Group brainstorming can be effective in generating creative ideas as suggested by cognitive approaches [3], and technology may help minimizing the effect of negative social processes [5]. A beneficial feature in group brainstorming is the ability to detect the context and content of the brainstorming through utterances of participants. IdeaExpander [14] is a tool to support group brainstorming by intelligently selecting pictorial stimuli based on the group's conversation on a chat. The pictures generally enhanced performance as measured by both originality and diversity of ideas [15]. Less investigated are face to face systems in group sessions that suggest keywords instead of pictures. Systems suggesting keywords and topics have recently been applied successfully to improve exploratory search processes [1, 10]. Such systems predict the current intent model of the user in the exploration and suggest possible explorations. These approaches have also been found useful in avoiding keyword input by selecting and manipulating suggested keywords by touch [8]. The present work investigate further alternative input modalities such as speech to text that permit the system to run in the background without interrupting the creative process but providing a continuous resource.

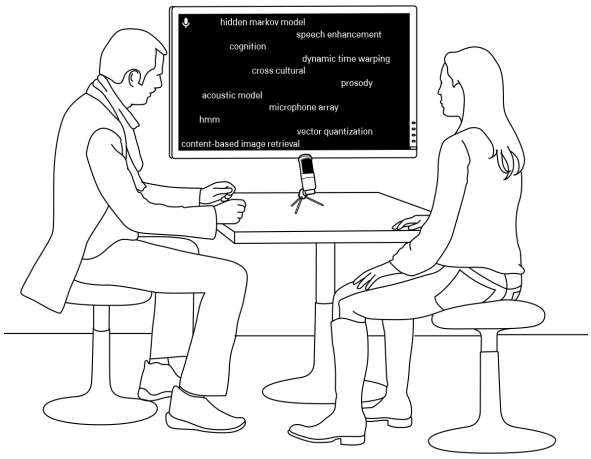

Figure 1. InspirationWall is an unobtrusive display supporting idea generation by leveraging speech recognition and automatic information exploration. It monitors users' discussion and automatically suggests keywords to support their idea generation.

C&C '15, June 22–25, 2015, Glasgow, United Kingdom

DOI: http://dx.doi.org/10.1145/2757226.2757252

Figure 2. Left: Two participants in a brainstorming session. Right: A screen capture from the InspirationWall interface.

SYSTEM DESIGN

InspirationWall is a non-intrusive source of diverse ideas (Figure 1). It was designed as a low-key visual aid, as to not interfere with the user's idea generation process. InspirationWall continuously monitors the discussions through a conference microphone and the input to the system is recorded from users natural interaction via speech recognition. Speech recognition is performed in real-time using Googles implementation of the HTML5 Web Speech API [16].

Recognized expressions are processed by an entity-based keyword suggestion system that returns related keywords by discovering associated and novel information related to the input [11]. Returned keywords are then displayed as slowly crossing the screen from top to bottom as to allow a progressive refreshing of displayed keywords. The graphical interface of Inspiration Wall is minimal: it runs fullscreen with a black background. Keywords are displayed in white. Every two seconds providing the keyword buffer is not empty a new keyword appears at the top of the screen at a random horizontal position and falls slowly towards the bottom of the screen.

Keyword Suggestion System

As the set of potential keywords matching any part of users' discussion is likely to be much higher than what can be presented for the user, and the discussion can contain misleading cues due to the natural dialogues that the system listens, we use a centrality-based ranking of the keywords in a large knowledge-graph.

Intuitively, this approach allows the system to suggest central keywords that are related to the user input via the knowledge-graph rather than only suggesting keywords directly matching to the input. This can help discovering keywords that are highly relevant for the input, but at the same time central to the overall discussion [12]. The knowledge-graph G is undirected and labeled and consists of a disjoint union set of keywords and documents (called nodes $n \in G$) and the set of edges between the documents and the keywords. Each keyword in the graph is connected to a document it describes. For example, an article about "relevance feedback for information retrieval" could be described with a set of keywords, such as "information retrieval", "relevance feedback", "implicit feedback", "web search", and so on. In addition, we index the text of the articles that is used to retrieve an initially relevant set of documents from which the knowledge-graph is constructed.

The user's query consists of one or more words detected via the speech recognition system. A set of keywords detected are called preference keywords $q \in G$ in the graph, where $|q| = 1$ and q_j denotes the preference for keyword j. In our case no weighting is conducted for the keywords so the preference is uniformly distributed over for the given keywords in q.

We use the personalized PageRank method [7] to compute the ranking of the nodes given the q. It can be then formalized as follows. Let an individual node be denoted as n, and by $I(n)$ and $O(n)$ denote the set of in-neighbors and out-neighbors of n in G respectively. Let A be the matrix corresponding to the graph G, where

$$A_{ij} = \frac{1}{|O_{ij} \cup I_{ij}|}$$

if the node i links to the node j or vice versa, and $A_{ij} = 0$ otherwise. For a given q, the personalized PageRank equation can be written as

$$v = (1 - c)Av + cq,$$

where $c = 0.15$ is the teleportation rate. The solution v is a steady-state distribution of random surfers, where a surfer teleports at each step to a node n with probability $c \cdot q(n)$, or moves to a random neighbor otherwise. We compute the steady distribution by using the power iteration method with 100 iterations.

The weights of the v are directly used in ranking the keywords. As the size of our knowledge-graph is hundreds of millions of nodes, the computation is not possible on-line for the complete graph. To make the PageRank computation feasible with an acceptable latency, we approximate the set of nodes to be included in the initial graph by using a language

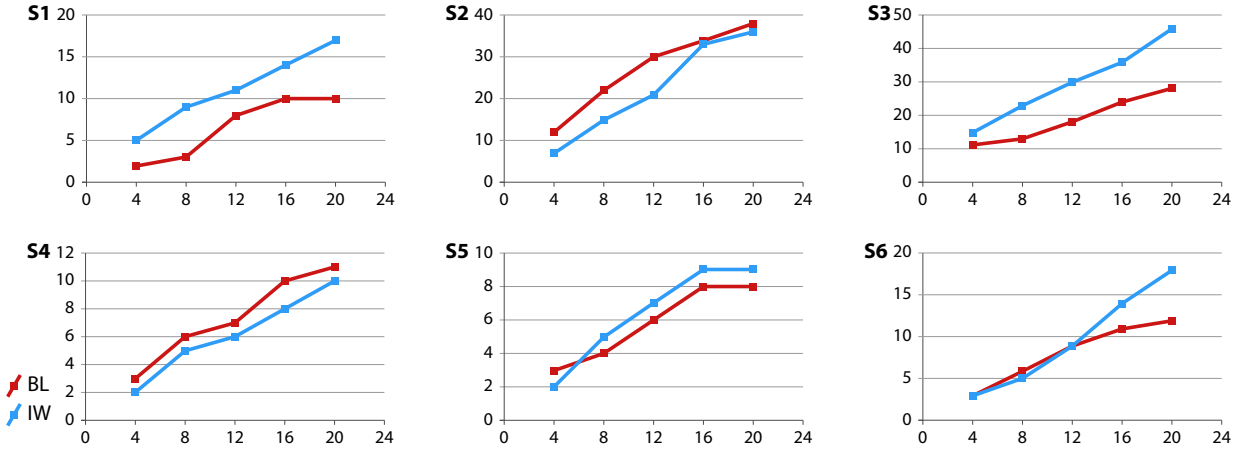

Figure 3. Accumulation of ideas per condition (BL = Baseline; IW = InspirationWall) in the different sessions S1,...,S6. On the Y-axis is the cumulative number of ideas, and on the X-axis is the time from the beginning of the session (minutes).

model approach of information retrieval [17] and select 3000 documents and the corresponding entities to be cumulatively added in the knowledge-graph at each iteration.

EVALUATION

We designed an experiment to evaluate the effect of the system on the idea generation process. The goal of the experiment was (1.) to understand if and how InspirationWall helped small groups generating more ideas, and (2.) to assess the overall effectiveness of the system as a creativity support tool using standard metrics.

Participants

The evaluation was conducted in groups of two persons (Figure 2). We recruited twelve participants (six pairs) with experience in idea generation activities from the computer science departments of two universities. Three of the participants were females and the mean age was $M = 28.33$, $SD = 3.98$. To simulate more natural discussions and brainstorming activity, we ensured that participants in the same pair knew each other. Participants were non-native English speakers from different countries and cultures (Iran, Canada, Spain, Nepal, Italy, Turkey, Sri Lanka, Rwanda, Kenya) with a similar level of proficiency in oral English. Their levels of education were: 25% PhD, 67% Master, 8% Bachelor. Participants received two movie tickets as a compensation for their participation.

Tasks

We used a within-group design, where groups were asked to perform two tasks: one with the support of InspirationWall and one without external support. We counterbalanced by changing the order in which the two tasks were performed and the order in which the groups were subjected to each condition.

The task was created to support an idea generation scenario and formulated as follows: *Imagine you have to come up with novel student projects on topic X. Please generate as many ideas as possible for new technologies, interaction techniques, methodologies, application scenarios, and so on, that might be used as more specific topics of the projects on topic X.* Two topics were used in the evaluation sessions: (1.) Robotics, and (2.) Wearable computing.

Metrics and Results

Quantity of Ideas

Since we were interested to check whether our application influenced the number of ideas generated, we have looked to the cumulative number of ideas considering time and session (Table 1). In total, the six groups have produced 107 ($M = 30.57$, $SD = 12.24$) ideas without external support and 136 ($M = 38.86$, $SD = 14.99$) using InspirationWall. In Figure 3, it is shown the accumulation of ideas per condition in the sessions, considering intervals of 4 minutes. In addition, video recordings obtained from the camera placed between the participants and the InspirationWall display allowed us to count the occurrences of participants looking at the screen (results shown in Figure 4). It is interesting to observe that the three groups (S1, S3 and S6) that have looked at the display the most, improved their performance with respect to the baseline condition, presenting a higher number of generated ideas and a more constant productivity.

Creativity Support Index

To measure the performance of our system in terms of creativity support, we involved participants in the assessment of the Creativity Support Index (CSI) [4].

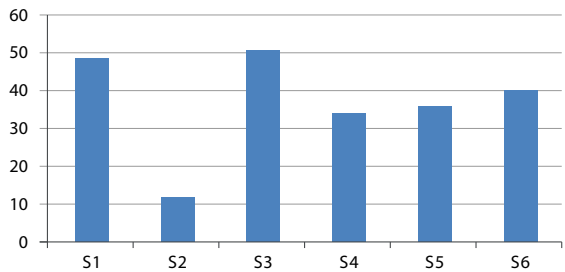

Figure 4. Number of occurrences of participants looking at the screen in sessions S1,...,S6.

Table 1. Accumulation of ideas per condition (BL = Baseline; IW = InspirationWall) at time $T = 4, 8, 12, 16, 20$ in sessions S1,...,S6. For each point in time, p-values from paired t-tests are also shown.

T		S1	S2	S3	S4	S5	S6	M (SD)	p
4	**BL**	2	12	11	3	3	3	5.67 (4.55)	1
	IW	5	7	15	2	2	3	5.67 (4.97)	
8	**BL**	3	22	13	6	4	6	9.00 (7.27)	0.61
	IW	9	15	23	5	5	5	10.33 (7.34)	
12	**BL**	8	30	18	7	6	9	13.00 (9.38)	0.73
	IW	11	21	30	6	7	9	14.00 (9.51)	
16	**BL**	10	34	24	10	8	11	16.17 (10.48)	0.23
	IW	14	33	36	8	9	14	19.00 (12.30)	
20	**BL**	10	38	28	11	8	12	17.83 (12.24)	0.17
	IW	17	36	46	10	9	18	22.67 (15.00)	

This index is computed from two sets of six questions and each question related with a factor. The six factors that compose the CSI are: *Collaboration*, *Enjoyment*, *Exploration*, *Expression*, *Expressiveness*, *Immersion* and *Results Worth Effort*. Each pair of questions are weighted based in pair wise comparisons of the factors made by each participant. The result of the CSI was $M = 53.36$, $SD = 13.35$. The most important factors for the participants were *Expressiveness* ($M = 3.58$, $SD = 1.24$) and *Exploration* ($M = 3.83$, $SD = 0.94$).

DISCUSSION AND CONCLUSIONS

Creative ideas are often triggered by unexpected associations. InspirationWall offers a quiet additional source of information to fuel the activity of collaborative idea generation. This paper presents the implementation and a preliminary evaluation of such a system. Participants were asked to generate ideas but not explicitly to use or interact with the system which was simply provided as is. Our study shows that participants that used InspirationWall more – as indicated by the count and duration of gazing occurrences obtained through video analysis – tended to generate more ideas in total and over time. Those results suggest that InspirationWall contrasts the decay of idea productivity over time typical of traditional idea generation sessions. Although the CSI does not show a high value, it is still above the median value of the scale, with the most important factors for the participants being *Expressiveness* and *Exploration*. Such results confirms the effectiveness of automatic information exploration and keyword suggestion on idea generation, opening a variety of directions for future work, including for example application to other datasets, and allowing richer interactions with the system through touch. The novel approach on idea generation support described in this paper, the simple design of our prototype and the positive results of this preliminary study are the contributions of our work to the future of digital tools for creativity support.

ACKNOWLEDGEMENTS

The data used in the experiments is derived from the Web of Science prepared by THOMSON REUTERS, Inc., and the digital libraries of ACM, IEEE, and Springer. This work has been partially supported by TEKES (Re:Know) and the Academy of Finland (278090).

REFERENCES

1. Andolina, S., et al. Intentstreams: smart parallel search streams for branching exploratory search. In *Proc. IUI*, ACM (2015).
2. Binder, T., Ehn, P., De Michelis, G., Jacucci, G., Linde, G., and Wagner, I. *Design things*. MIT Press Cambridge, MA, 2011.
3. Brown, V. R., and Paulus, P. B. Making Group Brainstorming More Effective: Recommendations From an Associative Memory Perspective. *Current Directions in Psychological Science 11*, 6 (Dec. 2002), 208–212.
4. Cherry, E., and Latulipe, C. Quantifying the creativity support of digital tools through the creativity support index. *ACM Trans. Comput.-Hum. Interact. 21*, 4 (June 2014), 21:1–21:25.
5. Diehl, M., and Stroebe, W. Productivity loss in brainstorming groups: Toward the solution of a riddle. *J Pers Soc Psycho* (1987), 497–509.
6. Hill, G. W. Group versus individual performance: Are n+1 heads better than one? *Psychological bulletin 91*, 3 (1982), 517–539.
7. Jeh, G., and Widom, J. Scaling personalized web search. In *Proc. WWW*, WWW '03, ACM (2003), 271–279.
8. Klouche, K., Ruotsalo, T., Cabral, D., Andolina, S., Bellucci, A., and Jacucci, G. Designing for exploratory search on touch devices. In *Proc. CHI*, ACM (2015).
9. Maguitman, A., Leake, D., and Reichherzer, T. Suggesting novel but related topics: towards context-based support for knowledge model extension. In *Proc. IUI*, ACM (2005), 207–214.
10. Ruotsalo, T., et al. Directing exploratory search with interactive intent modeling. In *Proc. CIKM*, ACM (2013), 1759–1764.
11. Ruotsalo, T., et al. Interactive intent modeling: Information discovery beyond search. *Commun. ACM 58*, 1 (Jan. 2015), 86–92.
12. Ruotsalo, T., and Hyvönen, E. Exploiting semantic annotations for domain-specific entity search. In *Proc. ECIR*, Springer (2015).
13. Shneiderman, B. Creativity support tools. *Commun. ACM 45*, 10 (Oct. 2002), 116–120.
14. Wang, H.-C., Cosley, D., and Fussell, S. R. Idea expander: Supporting group brainstorming with conversationally triggered visual thinking stimuli. In *Proc. CSCW*, ACM (2010), 103–106.
15. Wang, H.-C., Fussell, S. R., and Cosley, D. From diversity to creativity: Stimulating group brainstorming with cultural differences and conversationally-retrieved pictures. In *Proc. CSCW*, ACM (2011), 265–274.
16. Web Speech API. `https://dvcs.w3.org/hg/speech-api/raw-file/tip/speechapi.html`.
17. Zhai, C., and Lafferty, J. A study of smoothing methods for language models applied to ad hoc information retrieval. In *Proc. SIGIR*, ACM (2001), 334–342.