Augmenting Visual Representation of Affectively Charged Information using Sound Graphs

Nadya A. Calderon* Bernhard E. Riecke † Brian Fisher‡

School of Interactive Arts and Technology
Simon Fraser University

ABSTRACT

Within the Visual Analytics research agenda there is an interest on studying the applicability of multimodal information representation and interaction techniques for the analytical reasoning process. The present study summarizes a pilot experiment conducted to understand the effects of augmenting visualizations of affectively-charged information using auditory graphs. We designed an audio-visual representation of social comments made to different news posted on a popular website, and their affective dimension using a sentiment analysis tool for short texts. Participants of the study were asked to create an assessment of the affective valence trend (positive or negative) of the news articles using for it, the visualizations and sonifications. The conditions were tested looking for speed/accuracy trade off comparing the visual representation with an audiovisual one. We discuss our preliminary findings regarding the design of augmented information-representation.

Index Terms: H.5.2 [Information Interfaces and Presentation]: User Interfaces(D.2.2, H.1.2, I.3.6)—; K.4.1 [Computers and Society]: Public Policy Issues—Human Safety;

1 INTRODUCTION

From Visual Analytics (VA) perspective, people engage in analytical processes in order to transform information into knowledge to later test it, and share it. The display of information should be done in such way that facilitates the detection of patterns in datasets and sense-making processes [8]. Within opinion mining or sentiment analysis, the objective is to extract sentiment-related information from unstructured text and, similar to VA scenarios, the process involves a data set to explore and a cognitive task that needs to be enhanced. Although this analysis has been used mostly for commercial tasks, there is an increasing interest in the analysis of affective dimension of the social web content [7]. For sentiment analysis there are different mining algorithms that analyze informal text and derive affective polarities, however there are not that many tools to support the analysis of the information extracted from the mining process [3].

VA provides guidelines on how to use visualization strategies to represent information with specific characteristics, however and regardless of our awareness of sound as an important channel for universally accessible interface design, designers still lack of guidelines on how to use sound in their work [5]. The sonification field by means of non-verbal sound, has been an ongoing topic of interest for the last decades and the most prominent results up to date are those on the use of sound as alarms, auditory icons and earcons as status indicators [5]. However, the use of sound as means for data exploration is still a field in experimental phase with no concrete usable guidelines. Although several research contributions validate data mappings into sound [9],[1], there is little consensus on what should be considered standard and more important, how to apply it regularly within the interaction design process.

We are interested in sound as an augmented information representation to support analytic tasks. Correspondingly, it is important to find and document mappings of attributes of information into multimodal representation that combine sound with other sensorial stimuli. Our assumption is that given its evocative nature, sound is a natural metaphor to represent the affective attribute of information when present. Thus, in this study we test if adding data sonifications to a data visualization improve participant's response latency and accuracy when judging the affective valence of an article.

2 AFFECTIVE VISUALIZATIONS AND AUDITORY GRAPHS

Public comments, or replies to social media posts on the Internet, are an example of information charged with an affective dimension. People actively participate and leave a trace of opinions online in different social media instances. We used data consisting of social exchanges and debates between people on the web around news articles and topics in the *Digg* website (`http://digg.com/`). The comments were annotated with a sentiment charge estimated by the *SentiStrength* opinion mining system [6]. The SentiStrength algorithm assigns both a positive and negative valence to every comment, each valence ranging from 1 to 5 with 1 being not negative/positive and 5 being very negative/positive. After having annotated each comment, the affective valence trend was defined adding all the signed differences between positive and negative valences of the individual comments.

2.1 Affective Visualizations

We used two complementary representations in order to visualize the affective content of the comments for a single article (see Figure 1). Given the specific characteristics of the affective values generated with the mining algorithm (having always both values, positive and negative for each comment), we followed Gregory et al.'s advice of representing valence axes in a way they can be viewed individually as well as in pairs[3].

The bar chart represents the frequency distribution of positive and negative values independently. The upper set of bars presents the distribution of positive valence while the lower set of bars represents the negative distribution. Figure 1 shows an article for which the positive valence of its comments is mostly high (mainly between 3 and 5) while the negative valence is for the most part, a low value (1). This pattern is consistent for an article with an overall positive trend. The second graph, to the right of the diagram, corresponds to a heat matrix of the individual comments bi-valence. The matrix uses the same color code as the bars, and the square's area represents the number of comments with that specific combination of positive and negative valence. Following our example, the heatmap shows a majority of comments falling into the [high-positive, low-negative] area of the matrix.

*e-mail: ncaldero@sfu.ca
†e-mail:ber1@sfu.ca
‡e-mail:bfisher@sfu.ca

IEEE Conference on Visual Analytics Science and Technology 2012
October 14 - 19, Seattle, WA, USA
978-1-4673-4753-2/12/$31.00 ©2012 IEEE

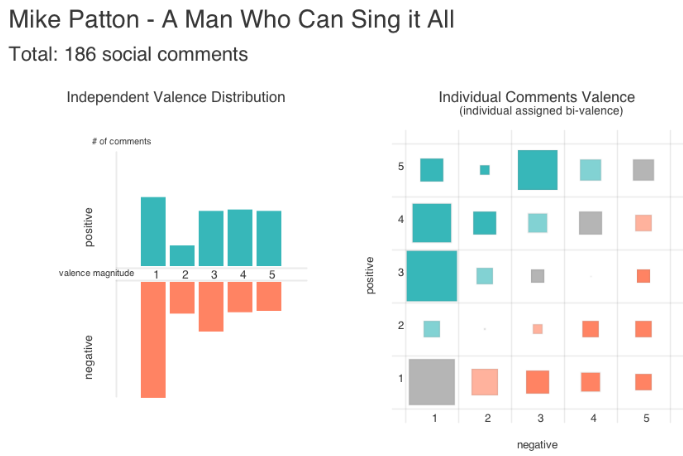

Figure 1: Affective valence visualization for a single article

2.2 Affective Auditory Graphs

We guided our simple melodic design by Lemmens [4] and Giard [2] experiments. They use major and minor chord mappings, and simple high and low frequency tones respectively. High pitch sounds are associated with positive valence and low pitch sounds associated with negative ones. The design of sound for this study followed that simple strategy and evolved from testing continuous sinusoidal wave tones to use discrete musical notes.

The data selected to be represented with the auditory graph was the same as the bar chart data displayed in Figure 1. Each magnitude value was mapped to a note. Two different octaves were selected to map the independent valences this way, each auditory graph was composed by 5 tones using a higher octave (C4-C5) to represent the positive chart and a lower one (C2-C3) for the negative values. The two melodic graphs were played in sequence having the positive followed by the negative one. At the end, the mapping was designed in accordance to what a traditional music training would suggest as an ascendent and a descendent easily perceptible sound. Examples of the representations for different trends can be seen at: `http://www.youtube.com/watch?v=DbuwdLQU7Cg`

3 Evaluation and Results

We conducted a controlled within-subjects experiment having response time and accuracy measures as dependent variables. 8 participants were asked to assess the affective valence of the articles, answering if they observed a positive trend, a negative one or simply no evident trend displayed. Each participant completed 24 trials each for the visual-only display and the audiovisual representation in balanced order. Additionally, participants also filled out a questionnaire evaluating the difficulty of both representations and their preference.

The analysis over the pilot data collected does not show a significant advantage of the audiovisual stimuli over the visual stimuli. Although mean response times in the visual condition were overall lower (M=5746 ms, SD=3605 ms) than in the audiovisual condition (M=7189 ms, SD=3038ms), this trend did not reach significance, $F(1,7)=1.10$, $p=0.33$. We were looking for an speed/accuracy trade off, expecting that if there was a larger consumption of time then the accuracy would improve, however it was not possible to find out the effect only throughout the pilot study conducted. We marked up some guidelines to evolve our design in a next stage as well as observe some implications for musically trained people as we discuss next.

4 Discussion

We learned that if a response time effect is expected, using static visualizations to be compared with audiovisual representations, constitutes a disadvantage for the multimodal version. We need to consider that having sound is adding a time dimension that is more comparable with an animated version of a visualization, including sound is by default animated. This is relevant regarding the design of augmented or multimodal information-representation systems. We want to research if perception of time-varying information is better with sound than with visuals. Since our intention is having augmented representations, time is likely to be a dimension that needs to be present in each mode. Thus, a future comparison should consider having animated visualizations although we understand that having animation is helpful dependent on the mapping of information and needs to be carefully designed.

The sound design, was a major issue for this project from its conception. Including knowledge from ecological sound design is the next exploration for our project. The pairing of evocative-sound and affective attributes of information may still be a valid connection but our ideas about simple mapping of data to notes need to be revised.

Finally, the design of the auditory graph was done in a way that having traditional music training would conveyed easily the information enclosed. A simple glance to the effect (detailed graphics depicted in the poster), although not statistically significant suggested that melodic auditory graphs tailor a specific set of skilled people while could be deviating the attention to not trained people. We need take this into account in order to determine how to select proper sound mappings that could enhanced inner trained abilities without having the opposite effects on people lacking them.

Acknowledgements

The authors would like to thank the *SentiStrength* group for providing us access to their tool and dataset: `http://sentistrength.wlv.ac.uk/`

References

[1] L. Brown, S. Brewster, S. Ramloll, R. Burton, and B. Riedel. Design guidelines for audio presentation of graphs and tables. In E. Brazil and B. Shinn-Cunningham, editors, *9th International Conference on Auditory Display (ICAD)*, pages 284–287. International Conference on Auditory Display, 2003.

[2] M. H. Giard and F. Peronnet. Auditory-visual integration during multimodal object recognition in humans: A behavioral and electrophysiological study. *Journal of Cognitive Neuroscience*, 11(5):473–490, 1999.

[3] M. L. Gregory, N. Chinchor, P. Whitney, R. Carter, E. Hetzler, and A. Turner. User-directed sentiment analysis: visualizing the affective content of documents. In *Proceedings of the Workshop on Sentiment and Subjectivity in Text*, SST '06, page 2330, Stroudsburg, PA, USA, 2006. Association for Computational Linguistics.

[4] P. M. C. Lemmens, A. de Haan, G. P. van Galen, and R. G. J. Meulenbroek. Emotionally charged earcons reveal affective congruency effects. *Ergonomics*, 50(12):2017–2025, 2007.

[5] M. Nees and B. Walker. Auditory interfaces and sonication. In *The Universal Access Handbook*, pages 507–522. New York: CRC Press, inpress.

[6] G. Paltoglou, M. Thelwall, and K. Buckley. Online textual communications annotated with grades of emotion strength. Technical report, 2009.

[7] M. Thelwall, K. Buckley, and G. Paltoglou. Sentiment strength detection for the social web. *Journal of the American Society for Information Science and Technology*, 63(1):163–173, Jan. 2012.

[8] J. J. Thomas and K. A. Cook, editors. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. National Visualization and Analytics Ctr, 2005.

[9] B. N. Walker and G. Kramer. Sonification design and metaphors. *ACM Transactions on Applied Perception*, 2(4):413–417, Oct. 2005.