

Visualizing the Affective Structure of a Text Document

Hugo Liu, Ted Selker, Henry Lieberman

MIT Media Laboratory

20 Ames St., Bldg E15

Cambridge, MA 02139 USA

{hugo, selker, lieber}@media.mit.edu

ABSTRACT

This paper introduces an approach for graphically visualizing the affective structure of a text document. A document is first affectively analyzed using a unique textual affect sensing engine, which leverages commonsense knowledge to classify text more reliably and comprehensively than can be achieved with keyword spotting methods alone. Using this engine, sentences are annotated using six basic Ekman emotions. Colors used to represent each of these emotions are sequenced into a color bar, which represents the progression of affect through a text document. Smoothing techniques allow the user to vary the granularity of the affective structure being displayed on the color bar. The bar is hyperlinked in a way such that it can be used to easily navigate the document. A user evaluation demonstrates that the proposed method for visualizing and navigating a document's affective structure facilitates a user's within-document information foraging activity.

Keywords

Affective UI, Document Analysis, Visualization, Color Psychology, Information Foraging, Hypermedia

INTRODUCTION

Visualizing a document's structure is proven to be useful in aiding information retrieval. The classic example is the table of contents, often found at the beginning of a book or lengthy document. In a table of contents, entries are thematic units such as chapters, sections, and subsections. Glancing through the table, a person can quickly glean the overall thematic structure of the document, allowing him/her to quickly access a section of interest.

Thematic structure, however, is only one of the ways in which people can naturally conceptualize a document. Another way by which people mentally index a document is in terms of its affective structure. In talking about stories, people often make references such as "the sad part" or "the climax." These are references to a story's *affective* structure. This paper introduces an approach for visualizing a document's affective structure represented in a color-based navigation bar.

First, we describe how the affective structure of text is pro-

duced through analysis. Second, we explain the choice of color as the mode of representation. Third, we describe how the color bar is used to navigate a document. Fourth, we discuss a user evaluation. We conclude with a summary of contribution and directions for future work.

AFFECTIVE STRUCTURE OF A DOCUMENT

Deducing the affective structure of a document takes place in two stages. First, each sentence in the document is classified into one of the six basic emotion categories of happy, sad, angry, fearful, disgusted, surprised, plus neutral. Second, a heuristic smoothing algorithm accretes sentence-level affect into incrementally larger units, thus rendering the affective "structure" of the document.

Sentence-level analysis. Sentences are classified into one of the six basic emotion categories proposed by Ekman (see [1]), plus the neutral category. Classification is done using four linguistic models of text built from Open Mind Commonsense, a large corpus (400,000 facts) of commonsense knowledge about the affective quality of things, actions, and events [1]. This novel method of classification goes beyond mere keyword spotting by evaluating the *underlying semantic content* of text. Consider the text, "My wife left me; she took the kids *and* the dog." While shallow spotting of *mood* keywords would fail, our approach is able to classify this as a sad sentence based on commonsense knowledge about divorces. The import of such a comprehensive approach is that we are now able to classify text reliably enough to be able to render the affective structure of large text documents at arbitrary levels of granularity.

Rendering affective structure. Affectively annotated sentences are combined into successively larger units using a heuristic smoothing algorithm. The algorithm integrates information from a Bayesian model, layout structure, and discourse cues to determine how the affect of text can be generalized to successively larger structures.

To summarize the affect of text over multi-sentence regions, a naïve approach would be a simple *winner-take-all scheme* in which the most frequently occurring emotion wins. However, this assumes complete independence of the emotion categories, which is not true in our case. Along different affect dimensions such as anxiety or contentment, some Ekman emotions are naturally co-dependent, such as *sad*, *angry*, and *fearful* along the anxiety dimension. To account for co-dependencies, we use a simple Bayesian network to compute, for each sentence, the probability of being in each of the seven categories given




Figure 1. The affect color bar in the top frame, document in the bottom frame.

the annotation of the current sentence. The network was learned from an annotated training corpus of 510 sentences across 15 documents, over various sub-genres of story text. Summing the probability of being in each category over all the sentences in a text region, we were able to achieve a better smoothing model than the naïve voting method.

In order to determine boundaries for regions over which affect is summarized, we consider layout structure and discourse cues. Layout structure prefers that affective structure boundaries correspond with layout breaks, such as paragraph, scene, or chapter breaks. Discourse cues are keywords and phrases which saliently denote a shift of discourse, such as sentences beginning with “all of a sudden,” or “surprisingly.” These cues provide more fine-grained clues about affective structure boundaries, and are used in conjunction with layout structure to partition text.

VISUALIZING AFFECT USING COLORS

Our motivation to represent affect using color is two-fold. First, colors have an inherent affectively evocative nature. Using [2] as a guideline, we chose one particular affect-to-color correspondence scheme. Since no color encoding can claim to be completely intuitive, there is a *learn-it* component to this scheme. Also, since the emotional connotation of color has a cultural component (for example, *white* represents death in some asian cultures), it may be desirable in the future to develop culturally sensitive schemes.

A second motivation for representing affective structure using color is that colors are visually distinct, can codify a large number of categories, and are identifiable even in small slices. Consider the alternatives. A graph plot visualization would require a different contour curve to represent each emotion. Regions annotated with labels or patterns would not be identifiable if occurring in small slices.

NAVIGATING A DOCUMENT USING THE COLOR BAR

The affect color bar (Figure 1) displays in a frame above the text document being visualized. The x-axis position of the bar corresponds to the y-axis position of the document. The granularity of affect analysis can be varied using the zoom controls to the right of the bar. Mousing over any part of the bar will pop-up a small tooltip whose background is of the same color. An emoticon corresponding to the emotion depicted by that color is included to help reinforce the meaning of that color. A text caption indicates the chapter name / section title of the story region being moused over. The bar is hyperlinked such that clicking an x-position on the bar scrolls the document in the bottom

frame to the corresponding y-position. Optionally, the document text can be colored to match the colors in the bar.

EVALUATION

We presented four users with four story documents (e.g. news article, novel) each, two that each user characterized as *familiar*, and two unfamiliar. For each document, users were asked to perform two timed information access tasks (e.g. *navigate to where the wolf eats the grandmother*). In all cases, users were asked to access an important event in the story. The first task was performed with the affect bar, and the second task, with a layout identical to the affect bar except that the color is uniformly yellow and no emoticon is shown in its tooltip (*control*). The results of this initial study showed that the affect bar improved the speed of the task by an average of 27% for familiar story documents and 36% for unfamiliar story documents. Improvement was not as dramatic for familiar stories, where users more frequently relied on their knowledge of the *order of events* in navigation. Also, users sometimes reverted to browsing by thematic events when there was confusion about the meaning of the colors. A longer-term study in which participants have more fully acquainted themselves with the color system might yield better results.

CONCLUSION AND FUTURE WORK

This paper outlines an approach for using a color-based affect bar to visualize and navigate a text document. An evaluation shows that such a visualization tool improves a user’s ability to quickly navigate to points of interest in a document, suggesting that affective document visualization is a promising information navigation tool.

In future work, we hope to more fully evaluate the approach, especially with respect to genre-generic texts such as web pages. We would like to evaluate the effectiveness of our color-encoding scheme versus other representation styles, such as composite plots of affect along single dimensions (e.g. anxiety). Finally, we are working on visually overlaying thematic and affective structure and studying whether an integrated visualization where affective structure is contextualized by thematic units might be more intuitive for users.

REFERENCES

1. Liu, H., Lieberman, H., Selker, T. A Model of Textual Affect Sensing using Real-World Knowledge. To Appear in *Proceedings of IUI 2003*. Miami, Florida.
2. Valdez, P., & Mehrabian, A. Effects of color on emotions. *Journal of Experimental Psychology: General*, 123, 394-409. (1994).