

Visualizing Sentiments in Business-Customer Relations with Metaphors

Guia Gali

OCAD University
Toronto ON, Canada
ggali@faculty.ocadu.ca

Symon Oliver

OCAD University
Toronto ON, Canada
soliver@faculty.ocadu.ca

Sara Diamond

OCAD University
Toronto ON, Canada
sdiamond@ocadu.ca

Fanny Chevalier

OCAD University
Toronto ON, Canada
fchevalier@ocadu.ca

Abstract

This project explores how visualization of sentiments, extracted from social media posts, can foster transparency and strengthen relations between businesses and its customers. Guided by the nature of the data and an iterative design based on our end users' feedback, we examine a variety of visualization styles and metaphors as possible directions for a common set of tools to benefit both end users.

Keywords

Data visualization, Metaphors, Trust, Social Media

ACM Classification Keywords

H.5.2 User Interfaces: *Graphical user interfaces (GUI)*.

General Terms

Design, Theory

Introduction

With social media activity at its peak, businesses see an opportunity to deploy new strategies in monitoring, engaging with, and identifying customers' concerns [1]. As a result, social monitoring tools that assess customer opinions are of major interest to businesses recognizing the benefits of qualitative online input to inform business decisions. Additionally, social media fosters a collaborative, trust-based forum where potential clients can learn from their peers' past experiences—be it positive or negative—before committing to a product or service.




Figure 1: Money Trees uses the metaphor of growing trees to compare positive, well-discussed banks (big lively green trees) to negative, least discussed banks (small dark trees) and vice versa. Note how sentiment and amount of discussion changes over time.

In this work, we are interested in how we can leverage user-generated online posts as a means to foster transparency and strengthen relations between businesses and their customers, through sentiment analysis and visualization [11]. Based on our industrial partner's needs, we justify why such visualizations should be designed as a common tool for both parties and their differing objectives. We discuss the design methodology we derived to address this goal. To conclude, we illustrate our points with a selection of sketches on bank-related data (e.g. Figure 1).

Background and Motivation

Our work was developed in collaboration with social media users and our industrial partner—*Royal Bank of Canada*—who provided us with bank-related data of over 25,000 online posts pertaining to any of the five major Canadian banks' names. Posts were collected from social network sites and news services, and later pre-processed by a text analytics company¹.

Validity of Data

Due to the private nature of our data mining company, our users faced a black box on how some variables were calculated, thereby causing little interest since it was difficult to understand. Likewise, the complex and nuanced nature of human language was certainly noted in the fair amount of sentiment mislabeling in the dataset. Indeed, the wealth of literature on sentiment analysis—including recent work on analysis of financial blogs [7]—proves the challenging aspect of sentiment analysis. Investigating existing techniques go beyond the scope of this paper as our focus is on the visualization aspect. However, we strive to make our

¹ www.attensity.com

visualization generic so as to accommodate current and emerging sentiment models. As more specific tools are being developed, we believe our visualizations could benefit researchers by making it easier to identify potential flaws in sentiment analysis.

Visualization

Both end users were interested in visualizing opinions and sentiments on banks so as to gauge intent of purchase, faults in services, or even compare against competitors [6]. While ideal commercial solutions for banks (e.g. Sentiment Metrics⁴ or Cymfony⁵) exist for social media monitoring, such tools require strong analytical skills as they are typically designed for strategic marketing, which neglect social media users.

Our end users also see the benefits of this visualization as a potential medium for increasing publicity: be it a bank showcasing good reviews or a customer voicing concerns. Such a tool can help establish a transparent and trustworthy platform in social media marketing [10]. Hence, visualizations should be accessible and engaging for social media users without business training. Existing sentiment visualization tools for such novice users either focus on fields other than finances (e.g. Blews [5]) or are too limited and abstract (e.g. We Feel Fine⁶), and therefore of little interest to our business users.

Thus, our goal is to provide visualizations that are comprehensible and engaging to online users, while being rich and descriptive enough to serve as professional tools for banks.

⁴ <http://www.sentimentmetrics.com>

⁵ <http://www.rcymfony.com>

⁶ <http://www.wefeelfine.org>

Methodology

Based on our users' needs, we believe *casual information visualization* is an ideal candidate for our case. Casual information visualization shares social media's inclusive nature as it lends itself to a wider range of user demographics, from experts to novices and from work-oriented to everyday situations [9]. This visualization style can therefore facilitate a symbiotic relationship where both users and banks can collaborate on betterment of services, creation of well-informed and customized products, and open familiar grounds for discussion [7].

With the aforementioned premise, we derived a general methodology for the design of our visualization, based on the following principles:

- P1 Social transparency:** it is essential to disclose background processes in a familiar and transparent manner for establishing and preserving trust in the visualization. Trust encourages cooperation, which is necessary in facilitating collaborative discussion and co-creation of well-informed products and services between banks and users [11,12].
- P2 Sentiment as judgment indicator:** "Many important trust decisions are made in affect-rich contexts" [2], such as deciding who to trust with one's money. Users with no prior knowledge of particular banks can use visualizations to help them make a decision by comparing banks' sentiments. In this respect, this tool can be an effective outlet for user to rectify (or praise) bank performance. Conversely, banks with positive reviews can also exploit such tools as free advertising generated by bank users [1].
- P3 Monitoring:** Casual information visualization lends itself effectively to an episodic type of usage with

monitoring tools [9]. In our case, quick and clear indications of sentiment changes provide banks the chance to take proactive decisions in remedying complications as they can check (at any given time) the presence of trending topics, abrupt changes in sentiment, or simply to verify assumptions made by their own financial and consumer market analysts.

- P4 Metaphor:** using common contextual metaphors will help ease interpretation; hence establishing familiar diction where banks and users can easily communicate. While a less intuitive visualization could guard against one of the dangers of user-driven design, simply that metaphors only reiterate existing knowledge [3], tools should always be accessible in a direct and natural way, "*first to gain acceptance from the users, and second to promote unimpeded investigation of their data without having to think about the visualization methods.*" [8]

Design Sketches

Our design approach was largely charrette-based (i.e. series of sketches, workshops, brainstorm and critique sessions) [10] over the course of three months. Featured in this paper are three sketches (of many) that use a variety of metaphors and methods.

Each sketch focused on clear sentiment visualization of online posts regarding the five banks. Sketches are based on hypothetical situations informed by the dataset, which will be further tested in the prototype development stage as part of our extended work. Banks and social media users were present throughout the concept generation phase for feedback to identify the most promising key points and approaches that best address their needs.




Figure 2: SentiWheel shows sentiments around a wheel that is divided among the five banks. Lines are drawn between sentiments and banks based on shared keywords. The size of wheel also changes as days pass, depending on the total number of posts.

Dataset

For demonstration purposes, our initial dataset was shortened from 39 to 14 days (May 19 to June 1, 2011) and further reduced through deletion of every other entry to retain proportion. Three iterations later, just over 1000 entries were left in the dataset (12.5% of the 7000 reduced sample), a manageable size for manually producing our sketches based on real data. Extracted sentiment values encompass positive and negative (liking or disliking a product, respectively), neutral (neither positive nor negative) and mixed (equally positive and negative).

SentiWheel

Inspired by pie charts, SentiWheel presents a holistic view of the financial market as banks share the same large pool of users. This timepiece visualization takes the user through daily snapshots of online posts. The wheel is divided among five banks, with each bank's share further subdivided by its volume of sentiments.

SentiWheel balances a casual information visualization approach with pragmatic aspects of traditional visualizations by combining qualitative dataset with quantitative visual conventions. In Figure 2, we can compare banks and sentiments through their portion sizes in the wheel (P3). If users desire, details are accessible through interaction where numerical breakdowns of the data are provided (e.g. Bank 1 shares 32%, with 2100 posts on May 19).

The wheel's size dynamically adjusts as data is crawled online, reflecting the most current volume of posts in real time. The option of positioning several wheels side-by-side allows users to assess activity by comparing size and color patterns over days. Lines connect shared

keywords between banks, thus facilitating comparison of sentiments on similar topics. For instance, 'service hours' could rate as positive in Bank 1 but neutral in Bank 2. These connections allow users to determine what makes certain banks work over others.

While SentiWheel appears to balance both ease and pragmatic sensibilities, many users found the pie chart metaphor as weak and too analytical, taking more time than anticipated to make sense of relations and proportions between banks. Hence, a stronger visual metaphor that easily conveys and equally resonates with banks and its customers, may present the data in a more unobstructed and engaging manner (P4).

Money Trees

Ironically alluding to the idiom 'money doesn't grow on trees,' this tongue-in-cheek visualization embodies each bank as a tree with leaves from lively green (positive) to brown murky colors (negative), as shown on Figure 1. Similar to SentiWheel, it takes a daily snapshot of online sentiments. When running in real time, the trees appear to be alive as leaves appear and disappear as the dataset is updated.

While the colloquial reference to corporate money is apparent, using trees show a humanistic portrayal of the hierarchical and symbiotic relations between banks and customers. We can observe how banks cultivate their trees (i.e. customers) with the goal of nurturing big green leafy trees. This strong metaphor downplays the seriousness of financial affairs by positioning a system where banks treat customers like a live organism, in need of careful attention. This presents banks in a positive and proactive light where they persistently strive in maintaining customer patronage.




Figure 3: Emotional Tapestry uses layers of patterns to represent complexities and variance of sentiments online. As particular sentiments increase, its respective pattern becomes more prominent.

Money Trees should function well as an ambient monitoring tool (P3) as it is accommodating of episodic use. It can act as a complement to market analysis sourced by banks, which could illuminate underlying problems invisible in statistical dataset. In Figure 1, the trees begin with a few neutral posts (dark green color). As time passes, more negative and mixed sentiments arise. Days later, negative and mixed sentiments become more visible. A monitoring tool like this can aid in damage control by observing 'dying' trees and determining its causes through topic detection.

Moreover, Money Trees was well received by users as it quickly communicates through culturally embedded metaphors. Interactivity will allow users to expand leaf clusters for details, track topics and movement of particular posts (as it is reblogged), and observe how negative/positive sentiments proliferate in the online community. For instance, if a brown leaf is followed by more of the same, is it due to direct influence?

Emotional Tapestry

The success of a humanistic visualization approach led us to Emotional Tapestry that reveals the complexity of human emotions by layering sentiments on top of one another (see Figure 3). A grid symbolizing a full month—one square for each day—shows an overview of the wealth and variety of sentiments in online posts. Each sentiment is embodied by a specific pattern: positive values have rounded, pleasing shapes and colors while negative values are loud and aggressive. These patterns become more prominent as online posts of similar sentiment increases.

As the tapestry fills itself, users can spot problems when particular squares or grids begin to look unsightly

or disturbing. Motion can be added to patterns when the state of a square becomes alarming, emphasizing urgency in the situation. Similar to Money Trees, customers can influence the tapestry by posting critical comments as a means of protesting or raising attention to unfavorable services or products. This provides a potent outlet for criticism as it significantly impacts the overall visual feel of the tapestry.

Using a tapestry shows a wider historical overview of sentiments as opposed to temporal snapshots like in Money Trees. With this tool, banks can display the efficacy of their actions as they listen to customers (i.e. another means of enhancing market analysis). Adding interactivity will aid users in categorizing by topics, banks, etc., thus helping banks and customers see into details as they track changes over time.




Figure 4: Overlapping of different grids to generate aggregated views of opinions over the course of months.

Additionally, it is possible to overlap tapestries to build aggregated views of overall opinions (see Figure 4). In our example, we present a calendar-based view. However, the tapestry is versatile enough that each square can use data based on other criteria. Indeed, users welcomed the flexibility it provides as they see the huge potential for personalization and exploration.

Discussion and Conclusion

Overall, the visualization sketches were well received by end users. Though the pie chart sensibilities of SentiWheel seemed like an understandable metaphor, we have found upon user feedback that humanistic metaphors had the most trustworthy and profound impact on users. Money Trees resonated culturally with both banks and customers while the personalization potential, flexibility, and ability to display complexities of emotions through layering make Emotional Tapestry equally strong and versatile. We have discovered strong candidates for facilitating understanding and establishing trust at a metaphorical level, able to captivate users and establish diction with ease.

Effective sentiment visualizations, founded on fostering transparency and strengthening relations, could have profound consequences on how banks view or treat customers, and vice versa. By using humanistic metaphors, we can establish trust in the visualization, thus encouraging cooperation and collaboration through social media applications. Use of these metaphors through convenient, coherent and customizable tools further facilitate understanding and exploration as users compare information and track changes online.

Future Work

Our team will begin developing functional, interactive prototypes for user testing. This will aid in validating the design concepts and metaphors and further refine the visualizations through iterative participatory design. In addition, the prototypes will be used to test the dataset against different emotion models and sentiment engines. Information from this research can possibly aid in improving precision of sentiment labeling and development of more specific tools.

References

- [1] E. Cambria and A. Hussain. Sentic Computing for social media marketing. In *Multimedia Tools and Applications*, pp 1-21, 2011.
- [2] J. Dunn, M. Schweitzer. Feeling and Believing: The Influence of Emotion on Trust. In *Journal of Personality and Social Psychology*, pp 736–748, 2005.
- [3] S. Diamond, Artists & Designers: An Experiment in Data Visualization. In *Proc. ACM Creativity & Cognition*, pp 195-196, 2011.
- [4] B. Fogg and H. Tseng. The Elements of Computer Credibility. In *Proc. of ACM CHI '99*, pp 80–87, 1999.
- [5] M. Gamon, S. Basu, D. Belenko, D. Fisher, M. Hurst and A.C. König, BLEWS: Using Blogs to Provide Context for News Articles. In *ICWSM 2008*, 2008.
- [6] A. Kaplan and M. Haenlein. Users of the world, unite! The challenges and opportunities of Social Media. In *Business Horizons*, 53, pp 59-68, 2010.
- [7] N. O'Hare, M. Davy, A. Birmingham, P. Ferguson, P. Sheridan, C. Gurrin and A.F. Smeaton. Topic-dependent sentiment analysis of financial blogs. In *Proc. of CIKM workshop on Topic-sentiment analysis for mass opinion (TSA '09)*, pp 9-16, 2009.
- [8] A. Pang and M. Clifton. Metaphors for Visualization. In *Proc. of Visualization in Scientific Computing*, pp 1-9, 1995.
- [9] Z. Pousman, J. Stasko and M. Mateas. Casual Information Visualization: Depictions of Data in Everyday Life. In *IEEE TVCG*, 13(6), pp 1145-1152, 2007.
- [10] H. Sanoff. *Community Participation Methods in Design and Planning*. Wiley Ed., 2000.
- [11] M. Sawhney and E. Prandelli. Communities of Creation: Managing Distributed Innovation in Turbulent Markets. In *California Management Review*, 42(4), pp 24–54, 2000.
- [12] B. Shneiderman. Designing Trust into Online Experiences. In *Comm. of the ACM*, 43(12), 2000.