

Social Visualization of Health Messages

Brant Chee
 University of Illinois
 Urbana-Champaign
 Institute for Genomic Biology
 1206 W. Gregory Dr.
 Urbana, 61801

chee@uiuc.edu

Karrie G. Karahalios
 University of Illinois
 Urbana-Champaign
 Siebel Center for
 Computer Science
 201 N. Goodwin Ave. 3110
 Urbana, 61801

kkarahal@cs.uiuc.edu

Bruce Schatz
 University of Illinois
 Urbana-Champaign
 Institute for Genomic Biology
 1206 W. Gregory Dr.
 Urbana, 61801

schatz@uiuc.edu

Abstract

Seeking health-related information on the Internet is becoming commonplace. Patients use the Internet to learn about their conditions and medications, and to discuss symptoms and concerns with other patients on chat forums. Physicians also browse these chat sites to understand concerns, symptoms, and beliefs of patients that might not be revealed to the physician due to fear. This paper presents a series of visualization themes applied to patient-generated data to allow for the exploration of these data. The themes we address are clustering, history, and sentiment. We describe the process used to extract and prepare the data for visualization and conclude with a discussion of the visualizations and their use for a specific health community.

1. Introduction

A growing number and variety of health resources are available online, ranging from dedicated commercial health websites such as WebMD to specialized health sites such as CHFpatients.com to personal groups found on Yahoo or IRC. A recent study by the Pew Internet and American Life Project found that 80% of Internet users have searched for health information online, on topics ranging from diseases and medical problems to environmental health hazards to Medicare and health insurance [9].

These topics are reflected in a health site that we have been studying related to congestive heart failure, CHFpatients.com. The topics discussed in the forums on this site range from information-seeking behavior such as finding information on drugs or treatment regimens to opinion gathering on drugs, hospitals, and doctors, to providing support, including bereavement.

Discussions in social forums go beyond pure information seeking. People explore and read them to find a community of peers to share experiences, commiserate with, and ask opinions. Recent studies have shown that a person's relationship with others or social network is relevant to phenomena such as smoking cessation and onset of obesity [6][7]. It has also been shown that support groups are increasingly important in the treatment of chronic illness such as CHF [17].

Facebook, MySpace, and Orkut are examples of burgeoning social networking spaces. However, little to no data are available on the use of social networking sites for healthcare information. There is commercial interest, as evidenced in the startup PatientsLikeMe. While the topic of health as a social networking site theme is nascent, there are numerous other groups, message forums, and other online communities where people may gather. As of May 25, 2008, there were 162,742 health-related groups on Yahoo Health.¹

Current social forums provide a venue for people within a community to discuss and share ideas. However, such communities can be difficult to navigate, making it hard for a new user to find necessary information or people. In particular, sites dealing with personal, sensitive health information can be especially daunting. A newcomer joining such a community might find it difficult to immediately understand the social mores of the new social space, the appropriate conversation topics, and the size and demographic of the audience involved.

Our research aims to provide a "face" for these sites, such that the audience will understand what space they are entering, how many people are involved, what topics are addressed, how people are similar to them, and how the experiences of the community members catalyze interaction. In

¹ Counted from groups at <http://groups.yahoo.com>

particular, we focus on new methods for exploring and investigating *patient derived* health data. We do this by providing a series of visualizations that address three specific themes: clustering, history, and sentiment.

The audience for these visualizations is (1) the community members of CHFpatients.com or similar communities and (2) physicians treating patients in this community. It is commonly understood that patients usually spend a very limited amount of time per year directly interacting with a medical doctor, nurse, or health-care system. A typical face-to-face patient-physician encounter is between 15 and 20 minutes long [19]. Often this is an insufficient amount of time for a physician to go over all aspects of a patient's medical history, leading to patient questions and dissatisfaction when visiting a physician. Twenty-nine percent of the population who use the Internet report that not having time with their doctor is one of the top most frustrating health care experiences [18].

Furthermore, it is well known to physicians that patients with chronic conditions spend time on the internet to learn about their condition through exploration and discussion in chat groups [2]. This patient-generated information is useful to the patient community and their health care. We try to leverage such data with this work.

Physicians are further interested in these sites because they believe patients fear lecturing from clinicians. Therefore, patients may be less than honest about whether they are taking their medication and how much they are taking. This phenomenon is also known as white coat anxiety. Physicians would like to be aware of patient concerns without the patient fearing reprimand from them [2]. Patient-derived data could help physicians better understand the fears and beliefs of the community with respect to their care.

Threading of conversations on CHFpatients.com is deep, and the topics evolve from the discussion of drugs to emotive discussions of feelings of pain, loss, happiness, and sorrow. The evocative nature of the site is not easily seen in current reading interfaces or existing social networking visualizations. We believe that this is a large component of the community that is ever changing and needs to be represented and addressed.

To address this issue, we mine the chat sessions on CHFpatients.com for common topics. Specifically, we began by automatically extracting drug names and their corresponding outcomes in the form of side effects and attitude towards them. These features are then used to drive the visualizations.

With these visualizations, the patient and medical community will be able to:

- Explore patient-derived data about drugs, symptoms, family, and community concerns
- Uncover sentiment relationships over time in the community
- Apply statistical methods on these data in a simple manner to gain insights about the data
- Use the data as a training set to validate future groupings of drug regimens and symptoms

This paper is organized in the following manner. We begin by discussing work related to our visualization approach. Then, we describe CHF and the CHF community we are studying. The three visualization themes and the methodology used to organize the visualization data are discussed next. We conclude with a discussion of the pros and cons of the visualization approaches and directions for future work.

2. Background and Related Work

Information visualization is “the computer-assisted use of visual processing to gain understanding” [5]. Some of the goals of information visualization are to (i) give users a deeper understanding of data, (ii) encourage the discovery of details and relations that would be difficult to notice otherwise, and (iii) support the recognition of relevant patterns by exploiting the visual recognition capability of users [5]. Social visualization is a subset of information visualization that specifically deals with *information about people, for people* [14].

The majority of medical visualizations, to date, deal with problems such as image acquisition and processing (for example Magnetic Resonance Imaging scans or Computer Aided Tomography images) [5]. Until recently, historical patient data have been located at various medical facilities and have been archived inconsistently. Currently, sites like Google Health² and MicrosoftHealth³ allow individuals to create an archive of their own history. At the moment, these sites are for patient's personal use. There is no visualization or relationship to community. They are useful for patients to find personal health-related patterns; however, they have little connection to outside sources.

² <https://www.google.com/health>

³ <http://www.healthvault.com/>

In a similar vein, the LifeLines project visualizes a patient's personal health history. This work is seminal in the area of timeline health visualization [16]. It visualizes a patient's medical history as a timeline, depicting relevant events broken into various categories such as problems, allergies, and lab pathology. This differs from our visualizations in several ways. The data used to generate the LifeLines visualization are gathered from a patient's medical record, not from publicly-accessible message board postings. Second, instead of a visualization for a single patient, we anonymize patient data and aggregate it over many people.

Although they do not present health information, the visualization approach used by Viégas and Smith for Authorlines presented the same information from two perspectives [20]. Authorlines visualized Usenet groups using two visualizations: One focused on individual author posting information; the other focused on characteristics of the group posting as a whole. Seeing the same data from two perspectives can highlight patterns that one might not observe from only one vantage point.

Similarly, the CodeSaw visualization [11] showed author timeline information in an open source community. The CodeSaw visualization gleaned its data from that group's open source CVS archive and associated public mailing list. This led members of the community to reflect on timing patterns and allowed for them to message each other and further discuss their progress, bugs, and deadlines.

The approach we are taking contrasts with other approaches such as PatientsLikeMe, in that we do not ask patients to fill out a questionnaire about drugs, side effects, and mood. Completing forms can often be time consuming; continually updating forms requires the discipline of routine visits and additions to the form. For our visualizations, we derive the data from existing conversation archives. We would like the extraction of such information to be automatic and invisible to the user. Our approach may also alleviate skewed data due to the Hawthorne effect. It remains to be seen if an extraction-based approach is as accurate as an explicit form-based one such as PatientsLikeMe.

3. Community

The CHFpatients site discusses topics relating to Congestive Heart Failure. CHF is the leading cause of mortality in the United States [4]. CHF occurs when the heart does not pump enough oxygenated blood through the body. It is a chronic illness, an ongoing condition in which a patient may live with the illness for the rest of their lives. 96% of CHF

patients have one or more co-morbidities (illnesses) at the same time [1].

The chronic nature of CHF lends itself to support groups and communities where people develop long-term relationships. While there are numerous resources including newsgroups on the topic of CHF, the CHFpatients community and website is unique in that it is one of the larger forums and has an ongoing established community.

The CHFpatients community consists of over 10,884 messages and 1,100 people, many who have been involved since 1999. The messages are on a variety of topics, ranging from advice about medicine to bereavement over a loved one. This group is somewhat unique in that the moderator spell checks and edits the messages for clarity. Sometimes the moderator will also include comments at the end of posts and append links to disambiguate drug names. The following is a sample message that was anonymized from the message archives. <Link></Link> indicates that the moderator inserted a link into the message:

Hi All, I need your input. I'm having about 27,000 extra pre-ventricular beats in a 24 hour period, per a Holter monitor test. My electrophysiologist and cardiologist agree that I should go on <Link>sotalol</Link>/Betapace. They are putting me in the hospital on February 26 to titrate me up on it. I've refused the drug in the past because it is such a dangerous drug.

Is there anyone out there who could give me an idea of how you've done on this drug? I'd sure appreciate hearing about your experiences. Thanks so much.

The moderator's additions suggest a dedication to, and personal stake in, this community.

Many members post bios with demographic information including location, age, gender, and a description of their condition. For example, a user posted the following about herself: "CHF, hypothyroidism, hypertension, type 2 diabetes, breast cancer, sleep apnea, emphysema." Members in the community are eager to share support and information. These sentiments are echoed in the results of a Pew study of Internet users with chronic conditions [8].

4. Visualizations

The visualizations we present in this paper fall into three themes: clustering, history, and sentiment. We present them in the following subsections.

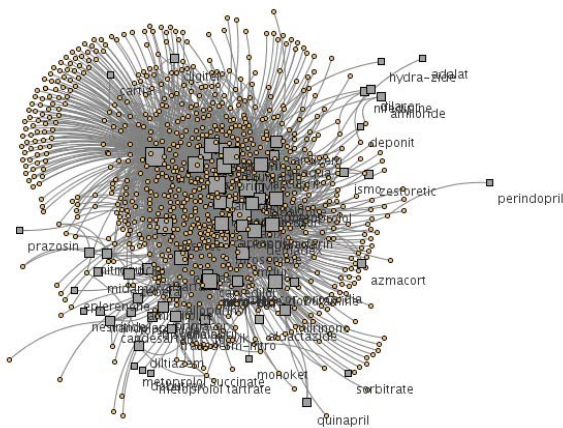


Figure 1. Overview of CHFpatients forum. Circular nodes represent users in the forum. Square nodes represent drugs. The size of the square nodes is proportional to the number of people discussing that drug.

4.1 Clustering

The goal of this visualization theme is to cluster similar information together. That is, we want to cluster people together with similar medications. This is the people-centric view. Users can also cluster families of drugs, including generic substitutions as well as common drug combinations.

Figure 1 depicts one such visualization. Specific drugs are represented as square nodes. The most discussed drugs appear in the center. Drugs that are infrequently talked about or not talked about in conjunction with other drugs appear near the edge of the graph. Drugs that are highly talked about and/or talked about in conjunction with other drugs appear near the center of the graph. Circular nodes connected to the drugs represent people discussing that drug in the forum. From this overview, we can see at a glance the general size of the forum, the number of drugs discussed, and the most talked about medications. Users can further zoom in and explore details of the visualization.

It is not assumed that anyone who is on a drug discusses it in a post. Conversely, it is also not assumed that when a person mentions a drug, they are taking it. However, we have noted that people who post about particular drugs are often taking them, especially if they frequently mention the drug or complain about its side effects. Many messages are anecdotes of experiences on particular drugs.

Figure 2 depicts a zoomed-in view of a selected drug, in this case, Nitro-time. Nitro-time is a time-

release version of nitroglycerin, a vasodilator for relaxing blood vessels to reduce blood pressure. This is a commonly-used drug for CHF patients. When a drug is selected in the visualization, the nodes of people discussing the drug are highlighted in red. In this manner, the user can see what segment of the population is interested in or knowledgeable about that drug.

This visualization uses a Fruchterman-Reingold (FR) layout algorithm that “attempts to produce aesthetically-pleasing, two-dimensional pictures of graphs by doing simplified simulations of physical systems. We are concerned with drawing undirected graphs according to some generally accepted aesthetic criteria: 1. Distribute the vertices evenly in the frame. 2. Minimize edge crossings. 3. Make edge lengths uniform. 4. Reflect inherent symmetry. 5. Conform to the frame” [10].

This form of visualization is suited to large overviews for general information. It is meant to be a starting point for further exploration.

4.2 History

The previous visualization showed all the information for the entire time span of the forum. If a user wants to explore time-based patterns, a time-based visualization is more appropriate. Figure 3 is a rendering of people’s drug notations over reconfigurable durations of time.

The vertical axis lists member names⁴ in




Figure 2. Zoomed view of Figure 1. When a drug is selected, in this case, Nitro-time, the user can see all the members who have discussed this drug in the forum highlighted in red.

⁴ Names have been altered for this publication.




Figure 4 Visualization of aggregate community sentiment over time. The green rectangular bars represent drug mentions. The blue rectangular bars represent symptoms. The dark blue waveform is a measure of sentiment. The bottom two bar graphs appear when a section from the top graph is selected. Drug mentions within the time frame appear on the left graph and side effect mentions appear on the right.

The right bar graph represents all of the side effects within the selected time period. The area corresponding to a side effect represents the percentage of mentions about that side effect. When a user moves their mouse over a portion of the graph, the drug name or side effect name appear to the left or right of the graph, respectively. This can be seen in Figure 4; the drug Coumadin is selected when the mouse is over that section of the graph.

The natural language processing we are currently using is not sufficient to determine if a person's sentiment is towards the drugs mentioned within the message. However, we can gauge to some degree whether or not people are positive or negative over specific intervals of time. Attitude is an important aspect of treatment, especially for chronic illness. Even if we do not know whether the positive or negative attitude is attributable to a particular drug, it may be possible to make deductions in certain circumstances. "Attitude and understanding by heart failure patients can have a significant effect on how they function, live with their disease, and potentially

even affect their outcome" [17]. It is especially important to determine people's attitude in the chronic illness environment. It is estimated that 43% of CHF patients suffer from depression [17]. Therefore, the ability to track attitudes of large populations of CHF or other chronic illness patients is helpful to predict and improve care regimens.

The average sentiment in Figure 4 is very negative at the bottom of the line graph and neutral near the median. It is interesting to note that although 40% of people with CHF are thought to have depression [17], the sentiment analysis shows that on average, the sentiments in the group are neutral to positive. As a collective, they are rarely negative. As of yet, although there is no explicit evaluation of the sentiment analysis software, the scores appear consistent with those of a support group. That is, while a person may complain or seek support in a negative manner, many others may reply with positive messages of encouragement and support.

5. Methodology and Implementation

Another contribution of this work is the development of a toolkit to extract easily the information necessary for the visualizations mentioned previously. This extraction is facilitated by the creation of disk-based data structures and efficient algorithms that can scale to many messages.

Initially, archived messages were downloaded from the CHFpatients website using wget. The pages are html aggregates of individual messages. A parser was then created to extract individual messages from the archive. After the individual messages were extracted, header information containing the name of the person who sent the messages, as well as thread history, was analyzed. This information was used to generate a stop word list of people names that would be removed at index time.

The open source information retrieval package Lucene was used to create an index. This inverted index facilitated the extraction of useful information from the messages, such as drugs and side effects. The resulting data structure allows iterative passes through messages and message access via key words. Modifications to Lucene were made so that word-based ngrams could be indexed.

Ngrams are multiple word terms. Indexed ngrams range in size from 1 to 4grams; stop words such as *the/and/or* do not start or stop an ngram. In this manner, ngrams with little information content over smaller ngrams would be removed. For example, there is little difference between the ngram "the world trade center" and "world trade center." Figure

5 illustrates the resulting ngrams for the phrase “the glass of red wine,” which include: glass, red, wine, glass of red, glass of red wine, and red wine. Ngrams are useful because many drugs, such as *Beta Blocker* and *ACE Inhibitor*, or side effects, such as *painful urination*, consist of multiple words and are not necessarily meaningful if only one component is found.

... the glass of red wine ...

Figure 5 NGrams extracted for “...the glass of red wine...” are underlined.

The individual terms of the ngrams were stemmed using a Porter stemmer variant. Stemming involves normalizing words; for example, walks, walking, and walk are all about the same concept; these terms would all be stemmed to the same term. This removes some of the variation in the resulting lexicon.

According to Luhn, the most significant words are the mid-frequently occurring ones [12]. We used heuristics to remove the lowest frequency terms; for example lgrams that occurred in fewer than 8 messages and ngrams that occurred in fewer than 4 messages were removed. The high frequency terms were filtered using stop words, words that occur frequently and have little information content, such as *the, is, and, or*.

A taxonomy of common CHF drugs was created by a local medical doctor, consisting of 214 drugs in several categories, including: Beta Blockers, ACE Inhibitors, Diuretics, Vasodilators, Endothelin Blockers, Thrombin Inhibitor, Phosphodiesterase-III Inhibitors, Angiotensin II Receptor Antagonist, Calcium Sensitizer, Cholesterol Lowering Drugs, and Nitrate. These categories were selected because they represent a constructible, bounded-size, constrained vocabulary. Of these 214 drugs, 75 were found in the CHFpatients archive. One initial task of this research is to evaluate the viability of extracting known CHF drug regimens from the message group.

A list of 364 common side effects was also created. These were generated from both the messages themselves and the drug interaction/side effect information available through manufacturers of the drugs. The side effects were filtered for more colloquial terminology thought to be found within the messages as opposed to more technical medical terminology – for example, *chills* versus *temperature regulation problems*.

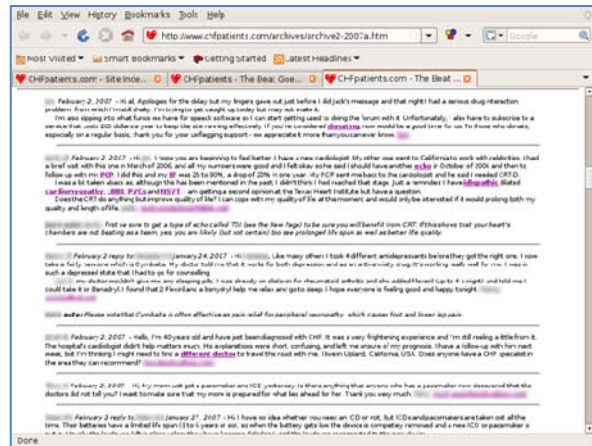


Figure 6 Snapshot of conversation on CHFpatients.com

The messages were analyzed for sentiment using a Naïve Bayesian classifier built in Weka,⁵ a data mining and machine learning toolkit. The classifier uses ngram-based features and is trained on the Rotten Tomatoes movie review corpus.⁶ The Rotten Tomatoes movie review corpus consists of 1,000 positive and 1,000 negatively processed one-sentence reviews of movies.

It is questionable whether or not the training data are applicable to the health message domain. However, the data are easily accessible and were thus used for this preliminary work to build the system. We are currently in the process of creating domain-specific training data.

6. Discussion and Future Work

The visualizations presented in this paper were created from patient-derived data from the forum CHFpatients.com. The interface for the conversation forum can be seen in Figure 6. The data in this forum are public. Although the information is easily accessible, the visualizations make some information salient that might not otherwise be observable. For example, the visualization allows one to see the number of people and the popularity of drugs at a glance, and provides the ability to locate cohorts who take similar sets of drugs. The point of the visualizations is to explore possible correlations and patterns in the community without having to remember everything that was already stated in conversation over a period of several years.

⁵ <http://www.cs.waikato.ac.nz/ml/weka/>

⁶ <http://www.cs.cornell.edu/people/pabo/movie-review-data/>

be satisfied with it until another drug in the same family appears, or a new regimen appears. Alternatively, participants could talk about a drug they stopped taking long ago in response to questions others have about it. We do not label any durations of drug use in the visualization, because we cannot conclusively prove them. We allow the users to make their own assumptions as they explore deeper into the data. We hope that the visualization further catalyzes questions and discussions in the forum.

In future visualizations, we will allow viewers to choose whether they would like to see the first mention of a drug or all mentions of a drug.

This visualization has limitations in its field of view. One cannot view a seven-year period at one glance; the user must scroll through the data. We are considering a syncing feature, where one can view first mentions of the same drug by several people at the same time to observe possible patterns over varied dates over the years.

6.3 Sentiment View

Further work on testing and developing training data specific to health messages is necessary. Additionally, we would like to extend the visualization to show time periods below the graph, so that one might be able to make hypothesis about the correlation of sentiment with a given time of year, for example, to see if average group sentiment declines during holiday seasons or if the opposite is true.

We are also interested in tracking the sentiment of individual users over time. It is hard to compare the sentiment of one user to another, and group sentiment analysis may not be accurate. We hope to monitor the sentiment analysis of a user over time and determine if it correlates to their feelings about their current treatment situation, and how the two co-vary.

We are currently developing named entity recognizers for treatments and side effects so that we do not have to rely on hand generating lists of treatments and side effects. We have also begun preliminary work on much larger datasets consisting of more messages and numerous groups.

With these additions, we plan to conduct a series of user studies with members of CHFpatients and local medical doctors to assess the efficacy and usefulness of the visualizations.

The full value of free text will become apparent when the natural language processing becomes more powerful. Medical questionnaires are limited to extracting answers from the questions asked, whereas health messages may contain information on a wide variety of topics. Currently, we extract exact names

of drug regimens and side effects, in a similar fashion to existing questionnaire-based systems. This is a measure of medical outcomes, in an objective fashion. The sentiment analysis is a subjective measure of how the patient is feeling, e.g. there may be major side effects but they are still happy to have reduced levels of pain.

As the depth of extraction increases, the sentiment view will evolve into broad measures of outcomes of medical treatments. The types of outcomes might range greatly, from reduction of pain to clarity of thought, and the types of treatments might range greatly, from diet and exercise to stress-reduction and spirituality. The effective path towards chronic illness requires monitoring the full range of lifestyle conditions and managing the full range of lifestyle changes [3].

Enabling patients to describe their own conditions in their own language is greatly superior to electronic medical records in hospitals or even personal health records in homes. As customary with information systems supporting persistent conversation, the topics and the moods can situationally vary for the environments and the individuals.

Working with health messages simulates the future world of health monitors, where persons respond continuously to questions in their everyday lives. The monitors record vectors for each person, which can be assembled into groups across populations, so that persons can manage their health based on actual experiences from similar persons in similar situations. The present research is thus just a taste of the future of social networking in healthcare infrastructure.

7. Acknowledgements

We would like to thank Richard Berlin, M.D. for the assistance in generating the CHF drug taxonomy and evaluating possible drug regimens. We would also like to thank Eric Gilbert for the use of his trained sentiment model and pointers to sentiment corpora. This work was partially funded by the National Science Foundation through FIBR-BIO grant 0425852 for the BeeSpace project www.beespace.uiuc.edu.

8. References

- [1] J. B. Braunstein, G. F. Anderson, G. Gertenblith, W. Willer, M. Niefeld, R. Herbert, and A. W. Wu, "Noncardiac comorbidity increases preventable hospitalizations and mortality among medicare beneficiaries with chronic heart failure," *J. Am. Coll. Cardiol.*, 2003, 42 pp. 1226-1233.

- [2] R. Berlin M.D. Personal communication, August 25, 2008.
- [3] R. Berlin and B. Schatz, "Internet health monitors for outcomes of chronic illness," *Medscape General Medicine*, 1999, Sep. 2:E9.
- [4] R. Berlin and B. Schatz, "Monitoring of quality of life for congestive heart failure populations," *Congestive Heart Failure*, Jan/Feb 2001, 7(1) pp. 13-23.
- [5] L. Chittaro, "Information visualization and its application to medicine," *Artificial Intelligence in Medicine*, 2001, 22(2) pp. 81-88.
- [6] N.A. Christakis and J.H. Fowler, "Collective dynamics of smoking behavior," *New England Journal of Medicine*, 2008, 358(21) pp. 2249-2258.
- [7] N.A. Christakis and J.H. Fowler, "The spread of obesity in a large social network over 32 years," *New England Journal of Medicine*, 2007, 357(4) pp. 370-379.
- [8] S. Fox, "E-patients with a disability or chronic disease," *The Pew Internet & American Life Project*. <http://www.pewinternet.org> accessed, August 28, 2008.
- [9] S. Fox, "Health information online," *The Pew Internet & American Life Project*. <http://www.pewinternet.org> accessed, August 28, 2008.
- [10] T.M.J. Fruchterman and E.M. Reingold, "Graph drawing by force-directed placement," *Software-Practice and Experience*, 1991, 21(11) pp. 1129-1164.
- [11] E. Gilbert and K. Karahalios, "CodeSaw: A social visualization of distributed software development," In *Proceedings of INTERAC*, 2007.
- [12] H. P. Luhn, "The automatic creation of literature abstracts," *IBM Journal*, 1958, 4 pp. 159-165.
- [13] B. Pitt and S. Goldstein, "The captopril-digoxin multicenter research group study on the comparative steps of captopril and digoxin in patients with mild-moderate heart failure: Implications for therapy," *Cardiovascular Drugs and Therapy*, 1989, 3 pp. 7-9.
- [14] K. Karahalios and F. Viégas, "Social visualization: Exploring text, audio, and video interaction," *ACM Computer-Human Interaction*, 2006.
- [15] B. Pang and L. Lee, "Opinion mining and sentiment analysis", *Foundations and Trends in Information Retrieval*, 2008, 2(1-2) pp 1-135.
- [16] C. Plaisant, R. Mushlin, A. Snyder, J. Li, D. Heller, B. Shneiderman, "LifeLines: Using visualization to enhance navigation and analysis of patient records," *Proc American Medical Informatics Association Annual Fall Symposium, AMIA*, 1998, pp. 76-80.
- [17] M. Silver, *Success with Heart Failure*, Perseus Publishing, Cambridge, MA, 2002.
- [18] H. Taylor and R. Leitman, "Study reveals big potential for the internet to improve doctor-patient relations," *Health Care News*; Harris Interactive, 8 Jan 2001, 1(1) pp. 1-3.
- [19] J.M. Travaline, R. Ruchinskas, G.E. D'Alonzo, "Patient-physician communication: Why and how," *JAOA*, Jan 2005, 105(1) pp. 13-18.
- [20] F. Viégas and M. Smith, "Newsgroup crowds and authorlines: Visualizing the activity of individuals in conversational cyberspaces," *HICSS-37, Hawaii, HI*, January 5-8, 2004.