Content Analytics System for Social Customer Relationship Management

Meenakshi Nagarajan, Danish Contractor*, Stephen Dill, Jitendra Ajmera*, Hyung-il Ahn, Ashish Verma*, Matthew Denesuk

IBM Research Almaden, IBM India Research*

Abstract

We propose to demonstrate a system that helps enterprises manage their social Customer Relationship Management (CRM) presence. Our system performs deep analytics over consumer posts and authors to help companies monitor their CRM presence, and to identify and respond to actionable posts. In this paper we share challenges and experiences in building this system. Findings yield a step change in the ability to implement social systems for customer care and leverage communities as effective proxies for CRM. This work also identifies novel knowledge and data mining challenges, such as identifying content that requires enterprise action, and prioritizing actionable posts based on various types of intent.

Introduction

Customers are increasingly using social media platforms such as Facebook, Twitter and blogs to express opinions, concerns and complaints about products and services, and to engage in conversations and solicit feedback from fellow customers and enterprises. Figure 1 shows an example of a social conversation between customers and an enterprise agent on a Facebook page set up by an automotive company. First, a customer posts a sarcastic complaint. Next, other customers respond to learn more and try to solve the problem. Finally, the company and other customers continue the conversation and approach a resolution. Enterprises interested in engaging with consumers on social platforms receive several such posts everyday day and employ agents from different business functions to monitor and resolve issues. The kinds of analytics required to support social CRM tasks are vastly different from those associated with monitoring (Chen 2009) and Business Intelligence (BI) applications (Cognos). Common aggregate-level metrics such as extent of unhappy consumers, brand/product news coverage, trending topics within focus communities, etc. are relevant but insufficient for identifying *individual* posts that an agent ought to respond to.

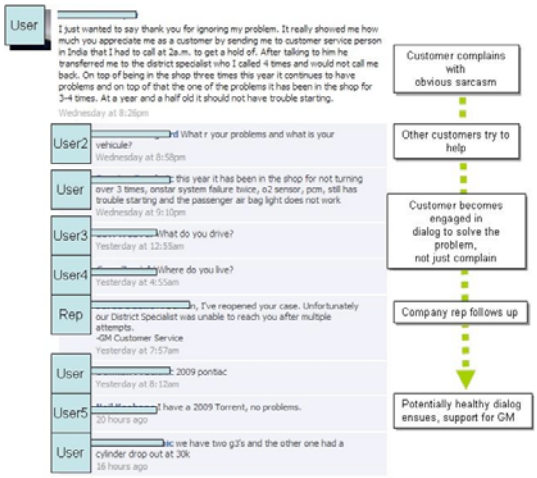

Figure 1 Conversation on a Social CRM Platform

A system assisting agents in CRM related activities must satisfy the following criteria.

1. Seamless Data Gathering and Ingestion
2. Identify 'actionable' posts in a timely fashion.
3. Assist enterprise agents in response generation
4. Provide a unified response interface and integration with traditional CRM data (email, call center etc.)
5. Provide a big picture view of a company's presence on all CRM channels.

In this demo, we will highlight *three scenarios* that highlight the core of our Content Analytics Platform: (a) Identifying and Ranking Actionable Posts (b) Assisting in making Responses to Actionable Posts (c) Channel Health Analytics. These scenarios are built over three core components of the system: (i) Data Crawl and Ingest (ii) Conversation and Author Analytics and (iii) User Interface components (see Figure 2). The scope of this demonstration will be limited to the social CRM presence of an enterprise in the *automotive industry*, as observed on their Twitter and Facebook pages.

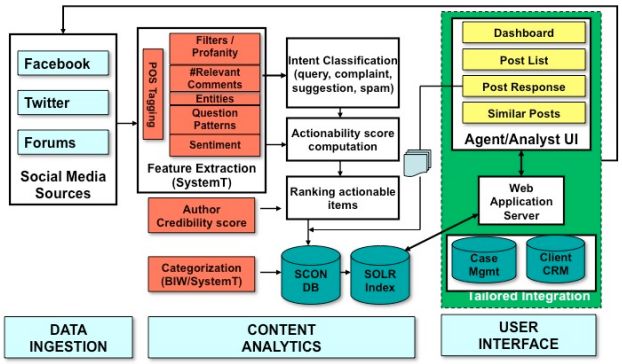

Figure 2 System Architecture and Components

Demo Scenarios

1. Identifying and Ranking Actionable Posts: For every conversation thread or main post that appears on a company's page, our system address two main goals: (i) Is this post or thread actionable, i.e. does it require a response? If so, which company business function (marketing, warranty, product related etc.) should respond to this post? (ii) How should actionable posts be ranked or prioritized in order to maximize agent productivity? Above goals are realized using a topic categorizer, an intent detection module and a ranking module (see Figure 2,3).

The topic categorizer uses a combination of bottom-up unsupervised clustering and top-down rule-based matching techniques to identify the business function that is best suited to respond to a post. **The intent detection module** identifies if a post is a query, complaint, suggestion, opinion or spam in order to deem actionability. Additionally, this module also identifies if a post has already been responded to or resolved, either by an agent or a member of the community. **The ranker module** uses timeliness of the post; severity and sentiment of a post such as in a complaint or query; 'temperament' characteristics of the author and other business function rules to rank posts within a business function category.

The 'actionability' and ranking criteria are configurable to meet changing business needs. For example, a business may want to rank posts differently around the time of a new product release vs. during normal business operation.

2. Assisting Response to Actionable Posts: In addition to getting posts to an agent, our system also assists them in constructing a response by showing previously resolved issues that are contextually similar. The infrastructure to support this process involves building a rich FAQ repository of past customer-agent interactions comprising of issues and resolutions. Analytics over a Solr-based index are implemented to support the matching of new incoming posts to issues in the repository. When similar posts are retrieved, they are ranked based on their similarity to the original post (e.g., content matching using

Figure 3 Snapshot of Ranked Posts and Metadata

TFIDF scores). Responses are presented to the agent to select or edit as appropriate. In addition, the agent can search the knowledge base using keywords to retrieve and peruse through prior posts. Our interactions with industry experts indicate that such a component is highly valuable in saving an agent time and effort.

3. **Channel Health Analytics:** In addition to providing post and thread level analytics, our system also aggregates activity information to present in several dashboard views. The goal is to be able to quickly spot areas of customer dissatisfaction (e.g., significantly higher number of complaints directed at the warranty business function), identify problem areas specific to spatio-temporal boundaries (e.g., Los Angeles dealerships have the most complaints for the first week of February), most discussed concepts, most active consumers on a channel etc. One such dashboard that we would like to highlight in the demo is the Channel Health Monitor. The motivation is to allow enterprises to see how their presence can be summarized across channels or data sources (Facebook, Twitter, Blogs, enterprise managed forums etc.) and across business functions or topics. We would like to show combinations of channels, topics and categories to view summaries of the kind of attention a company is receiving in terms of metadata inferred by the system, such as complaints, queries, sentiments and number of resolved issues.

A detail description of the system and algorithmic components can be found in this work (Ajmera 2013).

References

Chen, Ying; Proctor, Larry; Lelescu, Ana; Behal, Amit; He, Bin; Griffin, Thomas D.; Liu, Anna; Wade, Bradford W.; Davis, Trevor 2009. COBRA - mining web for COrporate Brand and Reputation Analysis. Web Intelligence and Agent Sys.

Cognos. Cognos Business Intelligence and Financial Performance Management, http://www-01.ibm.com/software/analytics/cognos/

Ajmera, Jitendra; Nagarajan, Meena; Ahn, Hyung-il; Verma, Ashish; Contractor, Danish; Dill, Stephen. System for Social CRM Analytics. World Wide Web Conferece, May 2013