

Geo-Temporal Visual Analysis of Customer Feedback Data Based on Self-Organizing Sentiment Maps

Halldór Janetzko, Dominik Jäckle and Tobias Schreck

University of Konstanz

Konstanz, Germany

Email: Halldor.Janetzko@uni.kn, Dominik.Jaekle@uni.kn, Tobias.Schreck@uni.kn

Abstract—The success of a company is often dependent on the quality of their Customer Relationship Management (CRM). Knowledge about customer's concerns and needs can be a huge advantage over competitors but is hard to gain. Large amounts of textual feedback from customers via surveys or emails has to be manually processed, condensed, and lead to decision makers. As this process is quite expensive and error-prone, CRM data is in practice often neglected. We therefore propose an automatic analysis and visualization approach helping analysts in finding interesting patterns. We combine opinion mining with the geospatial location of a review to enable a context-aware analysis of the CRM data. Instead of overwhelming the user by showing the details first, we visually group similar patterns together and aggregate them by applying Self-Organizing Maps in an interactive analysis application. We extend this approach by integrating temporal and seasonal analyses showing these influences on the CRM data. Our technique is able to cope with unstructured customer feedback data and shows location dependencies of significant terms and sentiments. The capabilities of our approach are shown in a case-study using real-world customer feedback data exploring and describing interesting findings.

Keywords—customer relationship management, review analysis, self-organizing maps, sentiment analysis.

I. INTRODUCTION

Many companies with business in the world wide web collect reviews and customer feedback of their products and services. One common way of assessing customer satisfaction are grading schemes (e.g., one to five stars) and free text forms allowing more detailed customer comments. But aside from showing the average rating or the distribution of ratings, more sophisticated and consequently also more expressive analyses are performed very rarely. This is surprising, as the free text provided by customers is a valuable source of hints with respect to customer needs and satisfaction levels, but a manual inspection is often not feasible. Modern approaches of text processing and visualization can help at this end, by summarizing important themes and sentiments in large amounts of text.

An effective analysis of textual customer feedback should involve and examine different aspects of the text content. The most obvious one is the *frequency of statements or terms*. Simple statistics and visualization methods like word clouds may help to get a first impression of most important keywords. But simple statistics do not help to analyze, whether the

customers liked or disliked these points. The next important aspect is the *sentiment* extracted from the context of the addressed keywords occurring in the text. E.g., customers may complain or praise products or services, and by using sentiment analysis, we aim at capturing this notion. From a company's point of view, negative statements are in many cases more important to analyze than the positive ones, to improve customer satisfaction. But the computation of one single sentiment score is not very expressive as customers might review more than one aspect, and different customers may have different opinions. Therefore, the challenge is to arrive at a fine-grained analysis of this complex data. The sentiment analysis should assign sentiment scores with respect to the attributes of the product or service, instead of computing one value. Customers, for instance, could like a certain bought product, at the same time complain about a too complicated ordering process.

Yet another key aspect holding valuable information in customer feedback data is the *geospatial location*. Customer feedback can be geolocated by several ways, including having the customer address in a corporate database, or by georesolving the IP address an anonymous web feedback was provided. From that we can derive the geospatial distribution of customer feedback, which is important for two reasons. First, for global companies, cultural differences may influence the customers' conception and country specific products or services should be offered. Second, besides cultural differences there is another aspect which may change customer's needs. The geographic location determines for instance the climate and may also impose delivery obstacles resulting from the geographic topology. In very dry areas, for example, it may be reasonable to leave a parcel outside the customer's housing, but in rainy areas the customer may complain about a soaked product. Concerning the topology, hard to reach customers (e.g., islands or exclaves) may complain about long delivery times, but there may be nothing the company could do about it.

This paper is an extended version of our previous work [1]. Compared to the previous work, we additionally take the time dimension of the customer feedback data into account. Performing temporal and seasonal analyses reveals further interesting patterns and insights.

Our main motivation for this work was the following starting hypothesis, to be explored on a real-world CRM data set: "The geographic position of reviewing customers correlates to



Figure 1. Visualization of customer feedback data using sentiment analysis and self organizing maps. Term groups with different sentiment scores are visible and additionally the geospatial distribution of the terms is displayed. The brightness of the background represents the coherence of the geographic distribution for each SOM node. [1]

their satisfaction levels and needs.” We wanted to see, whether there are differences in customer preferences caused by the geospatial location. The result of this analysis could help to improve the customer satisfaction by detecting differences in customer needs. Companies can therefore differentiate better among their customers and can easily focus and channel their efforts. Furthermore, we believe that the season of the review influences the sentiment for some terms. It is interesting to see which terms are affected by the temporal dimension and which not, to improve the company’s reactions. For instance, the need for support could be much higher after Christmas, as typically many technical devices are bought in this time period.

In this paper, we perform customer feedback analysis based on *sentiment maps*. Sentiment maps are the result of preceding opinion mining steps, where the occurrence of a term is drawn on a geographic map. The color used hereby depicts the sentiment and the sentiment map consequently shows not only the geospatial distribution of the term but simultaneously also the sentiment distribution. Following this approach leads to one sentiment map for each term. Further details of this approach can be found in the beginning of the analysis results section of this paper. A result of our technique is depicted in Figure 1.

We present in this paper our methodology analyzing customer feedback with respect to sentiment and geospatial customer location. Our contributions are the combined text and geospatial analysis of customer feedback data and the visual representation allowing a comparative analysis. Furthermore, we show that there are indeed frequent feedback terms (concepts) with a high geospatial dependency. The paper is structured as follows. First, we will give an overview to existing and related work in Section II, and then detail our approach in Section III. Findings from an application to a real-world data set will be discussed in Section IV. We will conclude with an outlook to future work.

II. RELATED WORK

Our work relates to the wider area of visual data analysis. We discuss a body of work on visual cluster analysis with Self-Organizing Maps, on analysis of geospatial and time-oriented data, and on feature-based text visualization.

A. Self-Organizing Maps for Visual Data Analysis

Many problems in visual data analysis require the reduction of data to perform meaningful analysis on a reduced version of data. Clustering reduces the data to a smaller number of groups to easier analyze and compare; and dimensionality reduction reduces the number of dimensions of data items to consider, and to project data to 2D displays. The Self-Organizing Map (SOM) algorithm [2] is a well-known method, which provides both data reduction and projection in an integrated framework. As a neural-network type method it learns a set of prototype vectors arranged on a regular grid, typically embedded in 2D. The method typically provides robust results in both data clustering and 2D layouting. Using regular 2D grids as neural structures for the SOM training, visualization in form of heatmaps, component planes, and distance distributions comprise basic methods for visual exploration of data using SOM processing [3]. SOM-based Visual Analysis to date has considered different application domains, including financial data analysis based on multivariate data models [4], analysis of web clickstream data using Markov Chain models [5], trajectory-oriented data [6], or time-oriented data [7]. Image Sorter [8] proposed to visually analyze collections of images by training a SOM over color features extracted from the images. We here follow that idea, in that we analyze geospatial heatmaps of sentiment scores using SOM of respective color features as well.

B. SOM-Based Visual Analysis of Geospatial Data

Many application problems involve georeferenced data items, and visual analysis approaches have been identified as

very helpful also for geospatial data analysis processes [9]. Choropleth (or thematic) maps are a basic, popular technique to show the distribution of a scalar value over a land-covering map [10]. Also, SOM-based approaches have been studied in context of geospatial data analysis, and proven useful to this end. When considering georeferenced data with SOM, basically two approaches exist. First, in the *joint data model*, one single data representation is formed by combining spatial and other multivariate data into a single vector representation which is input to the SOM method. Examples include [11], where a joint vector representation for both geolocation and demographic data was formed for census data analysis. More methods can be found in [12]. As a second approach, *linked views* integrate visual data analysis of each data aspect (geolocation, time, multivariate measures, etc.) in separate views combined by Brushing & Linking. One example system is [13], where a linked view system proposed the joint visual analysis of geospatial and multivariate data. Also, in [14], we proposed to jointly analyze geospatial and temporal phenomena by a linked view. There, SOM clusters can be computed for either data perspective, and the correspondence of clusters to the other perspective is shown by an auxiliary view. In our approach we do not consider geolocation data explicit for the SOM generation, but implicitly by the spatial-sensitive color features extracted from sentiment heatmaps generated from text data (cf. also Section III for details).

C. Visual Analysis of Temporal Data

Also, the temporal dimension plays an important role in many data analysis problems. Typical tasks relating to analysis of time include the comparison of data across time, finding similarities, trends, and correlations among variables. The time dimension can be accommodated for analysis either interactively or automatically. On the interactive side, users may select time intervals of interest for navigation and drill-down operations. On the automatic side, many approaches exist for detecting interesting intervals in time series [15], or finding clusters of similar temporal patterns for seasonal comparison [16]. There exists a rich body of work on visualization of time series data. Specific visual representations have been proposed to cope with the many challenges of time-oriented data, including long, multivariate, unevenly spaced or categorical time series [17]. In one previous work, we proposed an interactive system for cluster-based analysis of time series data sets, which also included the possibility to search by visual example, e.g., allowing the user to sketch a target line chart to retrieve [18]. In that system, the SOM approach was found useful to support the exploration of large time series data. While the SOM method can be an effective layout generator to compare many time series, there exist alternative layout schemes which can rely on user-selected or content-based ordering schemes to create small multiple displays. In [19], a taxonomy of layout strategies is given. In this paper, we rely on the SOM for generating geospatial and text-oriented data displays for comparison of customer feedback data. We will make use of small multiple displays to compare situations across time for identification of trends and patterns.

D. Feature-based Text Visualization

Finally, we relate to a body of work in visual document analysis. In general, *feature-based document analysis* abstracts a document (or collection, or stream) by a set of features which are more easy to visualize, as compared to the content of the documents. Numerous document features for different applications have been studied to date. For example, features scoring the readability of documents have been proposed in [20], and features applicable to classify authorship of documents have been surveyed in [21]. Sentiment features rate the polarity (in terms of positiveness or negativeness of statements) in a given text. In combination with time-series analysis, sentiment features can be used, e.g., to detect critical customer opinions in near real time, as possibly arising from some feedback channel [22]. In [23], we applied sentiment analysis to customer feedback data and analyzed it by means of geospatial heatmaps generated for the sentiments. While in [23], we considered only small sets of such heatmaps which we sequentially inspected, the focus of our work here is the comparative analysis of large numbers of sentiment maps, based on the SOM method.

III. TECHNIQUE

Our approach enables the geospatial visual comparison of customer feedback sentiments by using a Self-Organizing overview display. Figure 2 shows the overall process that is divided into four steps: (1) First, we extract a color feature vector for each sentiment map. (2) Second, we train the SOM and assign every sentiment map exactly one node. In step (3), we aggregate all sentiment maps that are located on the same map node. (4) Finally, we calculate the coherence and enhance the aggregated sentiment map with the content terms from the represented customer review texts. To analyze temporal data, we partition the dataset according to the date and process the proposed pipeline multiple times for each partition separately. We afterwards use of small multiple displays to compare different dates. We next detail these steps.

A. Process Pipeline

In order to process one partition of a temporal varying dataset, the introduced pipeline is processed exactly once. In this section, we describe the four different steps towards a sentiment SOM, before treating multiple temporal varying datasets.

Feature Vector Extraction. The feature vector we use as input to the SOM computation consists of localized RGB color values. We create a grid overlay for each sentiment map and calculate the color mean value for each cell. The mean value is determined by the color value of each single pixel contained in the corresponding grid cell. The representative feature vector for any sentiment map is created using the RGB color model. All RGB mean values are forming the feature vector:

$$\text{Picture1} = (R_{1,1} \ G_{1,1} \ B_{1,1} \ R_{1,2} \ G_{1,2} \ B_{1,2} \ R_{1,3} \ G_{1,3} \ B_{1,3} \ \dots)$$

$R_{i,j}$ represents the value of R for picture i and cell j . This format is used as feature vector representing one sentiment




Figure 2. Process pipeline to create self-organizing sentiment maps. The self-organizing sentiment map is determined by (1) extracting corresponding feature vectors, (2) calculating the best matching unit, (3) aggregating similar images according to the SOM, and finally (4) mapping the coherence and the additional information like the used terms to the images. [1]

map; each picture is assigned exactly one vector. Then, the extracted feature vectors are used to train the SOM using the SOMPAC implementation [24] (see also Figure 2 (1)).

Sentiment Map Classification. We apply a standard SOM training process following best practices suggested in [24]. Based on the defined SOM grid resolution, the prototype vectors are linearly initialized. Then, two learning phases are applied. First, a coarse learning is performed with a larger training radius, so that every considered node has a wide impact factor. Then, a fine-tuning training step is performed with a smaller training radius. Once the SOM-training has finished, the best matching prototype vector on the SOM grid (best matching unit, or *bmu*) is determined for each sentiment map by finding the node with the minimum distance (1).

$$bmu(SM) = \min_{k=1}^M \left(\sqrt{\sum_{i=1}^N (v(SM).i - v(node_k).i)^2} \right) \quad (1)$$

We iterate over all sentiment maps and calculate the best matching unit for each sentiment map SM . Then, we iterate all M trained SOM nodes and calculate the minimum Euclidean distance between the sentiment map and the trained SOM node. Therefore, the feature vector of the sentiment map and the vector of the SOM node are used. The corresponding vector is determined via the function $v()$ with size N . The control variable i addresses every single vector entry. Finally, the sentiment map is assigned to the SOM node with the minimum distance (see also Figure 2 (2)). The grid size can be chosen individually for each application.

Similarity-based Sentiment Map Aggregation. As the outcome of the SOM and *bmu* mapping, multiple sentiment images may share the same SOM node. Therefore, we need to provide aggregation of such sets of maps. To find a representative image for those sentiment maps different approaches are possible. We here chose to apply visual aggregation and merge all similar images into one. Therefore, every sentiment map is assigned a transparency value, so that we are able to create one image by lying one sentiment map upon each other. The resulting image visualizes all aggregated sentiment

areas. By adding multiple pictures on top of each other, the last added picture on top has the highest impact according to the process of alpha composition in terms of occlusion [25]. For that reason, we calculate the intersection of sentiment maps on our own based on the color, shown in Figure 2 (3).

Coherence Mapping and Map Enhancement. The last step of the pipeline is twofold: First, we map the background of the aggregated sentiment map to its coherence. Second, we enhance the aggregated sentiment map with additional information.

Aggregating multiple sentiment maps may result in an image showing a constant distribution. But when comparing all contained sentiment maps, the sentiment maps might be very diverse regarding to geospatial distributions. In order to understand the composition of those aggregated sentiment clusters it is important to define a quality criterion: the coherence of the sentiment maps. Thus, we make use of the background and define a coherence measure. The coherence measure (2) expresses how similar two sentiment maps are according to its feature vector. The coherence is mapped to the color range from black (high coherence) to white (very low coherence).

$$coherence(SMS) = \frac{\sum_{i=1}^N \sum_{j=1}^N dist(SMS_i, SMS_j)}{N \cdot (N + 1)} \quad (2)$$

The *coherence* is calculated for N sentiment maps (SMS) addressing the same SOM node. Summarizing, we build the average of all pictures including a distance function. We then sum the distance value of each sentiment map to all other sentiment maps. The distance function *dist* between two sentiment maps is defined in equation (3).

$$dist(p, q) = \frac{\sqrt{\sum_{k=1}^M (v(p).k - v(q).k)^2}}{|\{i \in 1..M : \neg(v(p).i = v(q).i = 0)\}|} \quad (3)$$

The distance function requires two sentiment maps as parameters with dimension M . The Euclidean distance is normalized

by the number of vector dimensions M excluding all dimensions with zero values in both dimensions. The problem of a possible low coherence raises with the algorithm of the SOM: To calculate the similarity in the second step the Euclidean distance combines sentiment maps that are very sparse, as the exact locations do not matter.

Sentiment Keyword Visualization. Every sentiment map corresponds to one term. As a consequence, if multiple sentiment maps are aggregated, the resulting image corresponds to multiple terms. Hence, we combine the aggregated sentiment maps with a simple but effective text representation: All terms are drawn semi-transparent with a gray border on top of the aggregated image. Also, the amount of sentiment maps that have been aggregated is indicated by a red number on the top left corner. Using an intelligent text layout algorithm, the analyst can easily identify the terms corresponding to the image; the text uses the full width and height to be easy to read. Figure 2 (4) illustrates the automatic labeling result.

Depending on the chosen grid size in the first step (feature vector extraction), the final result may differ. To allow data abstraction and overview large data sets, we typically chose a relatively small grid size, where the amounts of nodes is significantly smaller than the amount of considered sentiment maps.

B. Comparable Analysis of Temporal Data

For the analysis of temporal data, our technique needs to be extended as shown in Figure 3. In the previous section, the data according to only one time window has been processed in order to create the sentiment map. In case of a temporal dataset, we adapt our proposed technique: First, the temporal data is partitioned by date respectively. Second, the process pipeline is executed for each data partition separately; but the training phases change to obtain comparable layouts. Finally, all resulting sentiment maps are layouted and visualized. The visualization allows to compare multiple datasets by using the technique of small multiples.

Data Partitioning. To analyze the temporal dimension, we partition the data corresponding to predefined intervals by the user. For the analysis of seasonal trends, for example, the dataset is divided into the quarters of a year or years respectively. Each partition is treated separately.

Small Multiple Displays is a grid-like visualization technique allowing the comparison of similar graphics. We make use of this technique to visualize multiple sentiment map instances at the same time; this enables a seasonal comparison. We chose a square algorithm to layout all small multiples via a grid. Each grid cell is assigned one sentiment map. Therefore, we initially calculate the grid before assigning the sentiment maps.

$$rows = \left\lceil \sqrt{|SM|} \right\rceil \quad (4)$$

$$cols = \left\lceil \frac{|SM|}{rows} \right\rceil \quad (5)$$

We calculate the amount of rows ($rows$) by ceiling the result of the square root of the amount of all sentiment maps SM .




Figure 3. Adapted pipeline towards seasonal sentiment maps. First, the temporal data is partitioned according to user selected temporal intervals. Second, the process pipeline (see Figure 2) is processed for each partition separately in order to create small multiple displays. The sentiment maps are made comparable by performing the fine training on the same feature vector basis.

Then, we use this value to calculate the amount of columns $cols$. This step is especially necessary for an odd amount of sentiment map results. After the necessary amount of rows and columns has been calculated, a uniform width and height is assigned to each grid cell. In order to allocate resulting sentiment maps, we iterate through all columns and all rows assigning the corresponding spot to the sentiment map.


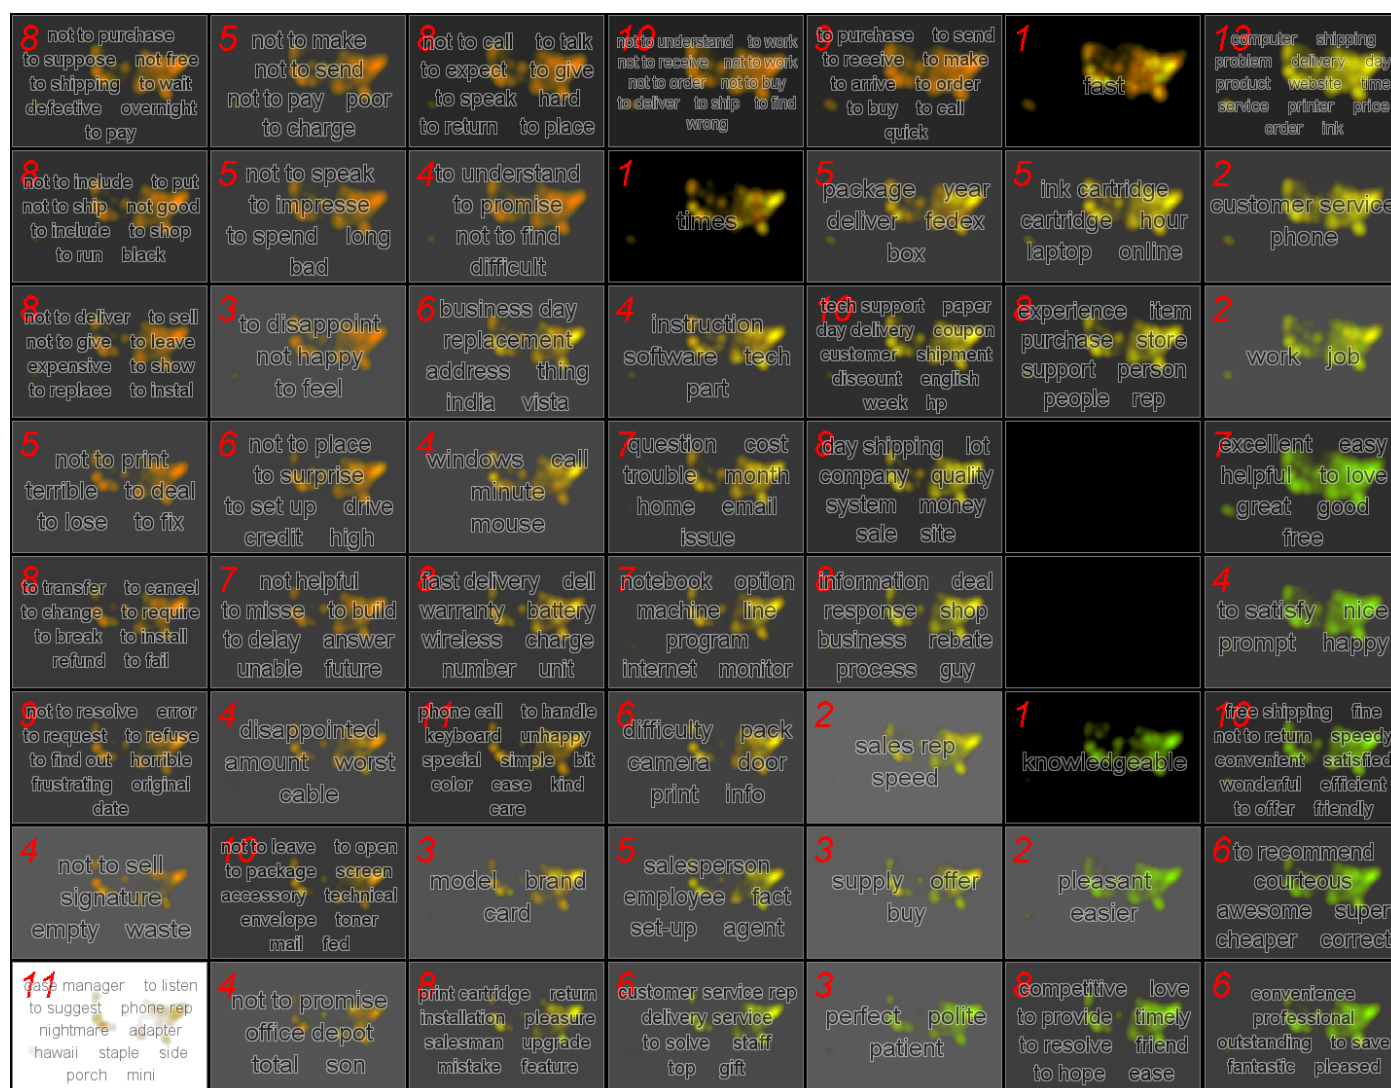


Figure 4. The location maps of positive (a) and negative (c) review terms are blurred to increase the visual saliency and to give a visual aggregation (b, d). Both blurred location maps are then combined (e) by using the RGB channels and show the distribution of positive (green), neutral (yellow) and negative (red) term occurrences. [1]



will result in yellow colors, while pure positive sentiments will result in green colors. We did not differentiate within negative sentiments or positive sentiments respectively as this differentiation is highly user and application dependent. But sentiment maps could be extended by this possibility. The final result of this technique can be seen in Figure 4(e).

Geospatial Analysis and Findings. We applied the technique described in Section III to a dataset consisting of 327 sentiment maps. These terms were found in preceding document mining steps and contains the words being nouns, verbs, and adjectives. Note that some of these terms are even compound nouns like "phone call" or negated verbs like "not to send". The resulting overview visualization can be seen in Figure 5.

On an abstract level there is a clear grouping and ordering of the sentiment maps visible. Terms with only negative occurrences (reddish images) are located in the upper left while positive terms (greenish sentiment maps) are located in the lower right. The first diagonal consists of terms being either mentioned negatively and positively equally often (upper right) and terms with a geospatial, diverse distribution (lower left). The SOM analysis enables the analyst to get a fast overview over terms being mentioned always positive or negative.

The strongly highlighted, white node of Figure 5 in the lower left contains eleven terms showing a very diverse geospatial distribution. As this is the node being highlighted most we will investigate this node in the following paragraphs. Detailed analysis via drill-down techniques are possible in our system and reveal the geospatial distribution for each single term. The visualization of all eleven contained terms is depicted in Figure 6.




Figure 6. Visual representation of all sentiment maps (in reading order: to suggest, staple, nightmare, to listen, hawaii, phone rep, side, case manager, porch, adapter, mini) contained in the lower left node of Figure 5. [1]

Inspecting the sentiment maps more in detail reveals that this node mainly contains sentiment maps with sparse and diverse geospatial distributions.

The most obvious sentiment map contained in this SOM node is the term "hawaii". This term is occurring mostly positively and collocated with the geospatial position of the Hawaiian islands. Inspecting the customer comments in detail, we found that customers liked the free shipping possibilities to Hawaii, which seems not to be taken for granted. Service managers can learn from this information that (Hawaiian) customers do care about the shipping procedure and that free shipping might be an advantage over competitors.

Also, the term "case manager" (third row, second column in Figure 6) shows an interesting pattern. Although mostly mentioned negatively because of language issues – the customer support was hard to understand because of foreign accents – there are many positive occurrences in Houston, Texas where customers liked the support regarding their printers. Service managers should now investigate further what the characteristics about the problems in Houston were.

Two further interesting sentiment maps are the ones of "nightmare" and "porch". Investigating the underlying reviews shows that the preceding sentiment analysis did not work correctly as all the reviews were purely negative. This is not a drawback of the method per se, but exemplifies the uncertainty of any sentiment analysis and the sensitivity of our method to the input data. The comments regarding the term "porch" were mentioning that the parcel was left unattended on the porch. The term "nightmare" was used in cases where the process of ordering and returning products did not go smoothly.

Temporal Analysis and Findings. We performed a temporal analysis of the sentiment distribution by partitioning the data set according to the time dimension. We used the quarters of a year in order to group the reviews into meaningful units. The analysis is initiated by computing the overall sentiment of the customer feedback messages. Then, we use the resulting sentiment score per message. Figure 7 shows the results. The time dimension is denoted on the x-axis and the sentiment scores on the y-axis. Unfortunately, the differences of the sentiment distributions are not significant as all boxes do overlap each other. Consequently, the overall sentiment distribution does not significantly vary over time when analyzed without any geospatial reference.

We further statistically analyzed and computed the sentiment distribution over time per term. We consequently selected all significant terms contained in the data set and inspected the occurrences. For each of them we assigned the sentiment. The sentiment scores were grouped by the quarter of the year they were received. Then, for each of these groups, we computed the average sentiment. In our case, this lead to the quarters Q3/2007 to Q1/2010 with the related average sentiment scores. As a last step, we compute the variance of these average sentiment scores resulting in a number, that describes, how much the sentiment score changed over time. Sorting this list of terms and sentiment changes over time leads to the most interesting terms, which can be inspected manually.

The simple temporal, statistical analysis described above, helped us in assessing which of the terms show interesting




Figure 7. Statistical analysis of the sentiment scores distribution over the different quarters. No significant difference between the distributions is observable.

behavior over time. But these statistical approaches do not reveal interesting findings per se as the geospatial reference is not regarded at all. As described in the technique section above, we partitioned the data set into meaningful time units – in our case into quarters of a year – and computed for these time units sentiment maps. In the following paragraphs, we will show interesting terms found by computing the variance in sentiment scores over the time quarters.

We investigated the top candidates found by the method above and present the corresponding sentiment maps. The first one we inspected is the term *camera*, shown in Figure 8, as it has one the highest variance scores. Comparing the different quarters of the year, it is obvious that only in the first quarter the term *camera* is mentioned positively in California. Interestingly, the second and forth quarters seem quite similar in the sentiment distribution. With this information, the CRM expert can inspect the corresponding feedback messages and plan special actions, such as special offers or sending apology messages.

The second term we present here, is the term *salesman* in Figure 9. Again, this term seems also being very sensitive to the time dimension and is not uniformly distributed. Furthermore, the term *salesman* seems to be more frequent in the first and last quarter compared to the other quarters. The reason might be that there is more need for guidance in the time around Thanksgiving and Christmas because of presents and salesman are important then.

In order to analyze the temporal development of the sentiment of all terms, we applied the Self-Organizing Map layout to the sentiment maps of one year and place them as small




Figure 8. The term *camera* shows one of the highest sentiment variation over time and reveals interesting patterns.

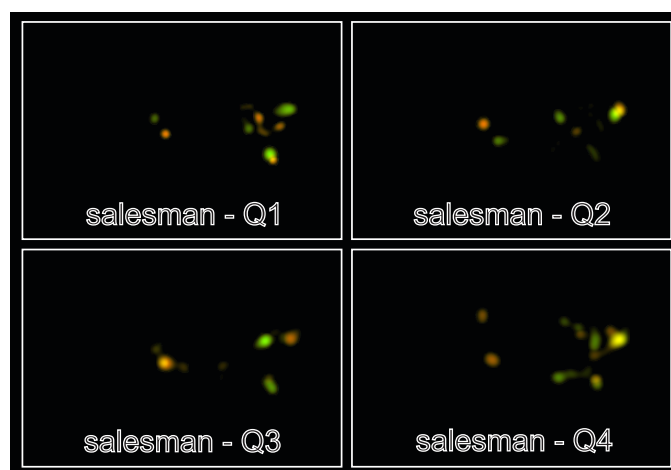


Figure 9. The term *salesman* shows one of the highest sentiment variation over time and reveals interesting patterns.

multiple side-by-side. By comparing the positions of the same term over the different years, a shift in sentiment becomes obvious. We exemplified one of these changes over time in Figure 10. We highlight the SOM node of the same term by Brushing & Linking and investigated in this example the term *customer service*. The term is mentioned quite ambiguous in the first year, but has a quite good overall sentiment distribution in the second year. Unfortunately, it is mentioned worse every year than, with the worst opinion in the last year. This shift of sentiment is visible by the movement of the term from the right side (mostly positive opinions) to the left side (mostly negative comments).

V. CONCLUSION

We presented our approach to visually compare and inspect large sets of textual customer feedback with respect to sentiment expressed regarding key concepts, and geographic distribution. For each concept, a sentiment map was rendered and

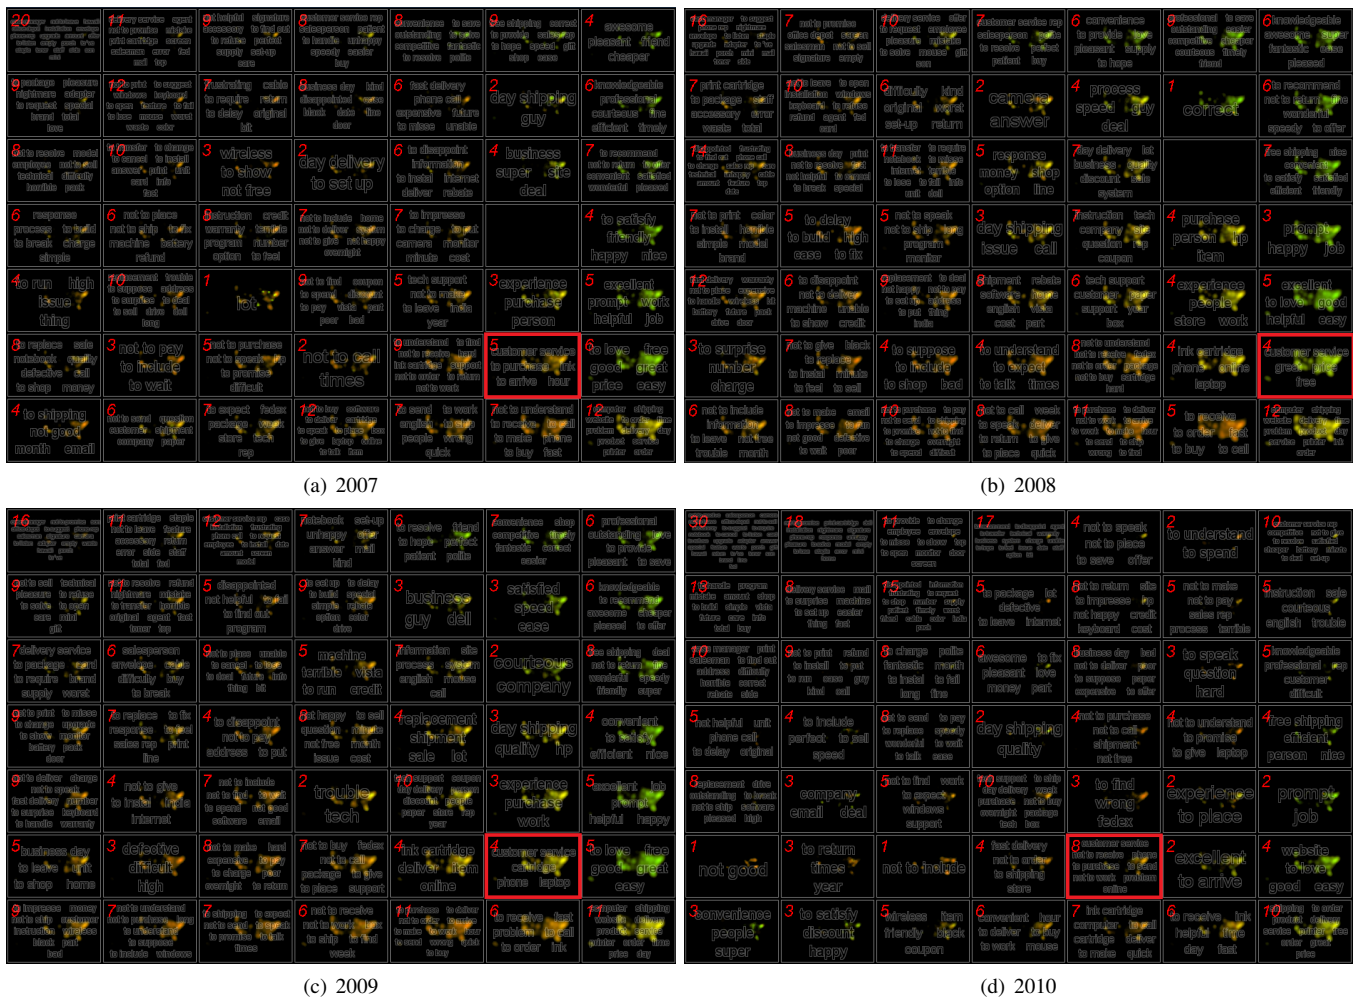


Figure 10. Comparison of the SOM layouts of the sentiment maps of different years. We highlighted the SOM node containing the term *customer service* by a red rectangle. The sentiment development over time is very good visible, as 2008 is the year with the best opinion about the customer service.

all maps were visually clustered and aggregated by the SOM approach. Interaction methods allow navigating the overview visualizations and drilling down for detailed inspection. Further temporal analyses of the customer feedback revealed seasonal influences on the customer opinion. Application findings presented indicate that key concepts and their sentiment scores being highly dependent on geographic position. Such findings can be very helpful in analyzing service levels across locations, products, customers, and similar applications in CRM. Our analysis system can be easily used on top of existing CRM systems, as we only need the text and the geographic position of the reviews.

We have several ideas to extend our work in future for improved analysis. One possibility improving the visual representation is the integration of semantic zoom approaches. Semantic zoom can allow merging neighboring SOM nodes to reduce the level of detail. Additionally, semantic zoom can be applied to the shown terms by using an ontology. The terms will be grouped by the ontology and only the

common parent of a set of related concepts will be visualized. The ontology also leads to another extension possibility we are going to integrate in future. We plan to show the hierarchic relationships between terms directly on the SOM representation. Last but not least, we want to consider more detailed map visualizations concerning production facilities and income distributions among different cities correlating geospatial dependent properties with the text features.

ACKNOWLEDGMENT

This work has been supported by the Consensus project and has been partly funded by the European Commission's 7th Framework Programme through theme ICT-2013.5.4 ICT for Governance and Policy Modelling under contract no.611688.

We are thankful to Ming Hao of Hewlett-Packard Labs for fruitful collaboration on the analysis of geo-referenced customer review report data. We also thank the Helsinki University of Technology for providing the SOMPak implementation which we used in this study.

REFERENCES

- [1] H. Janetzko, D. Jäckle, and T. Schreck, "Comparative visual analysis of large customer feedback based on self-organizing sentiment maps," *Proc. International Conference on Advances in Information Mining and Management*, 2013, pp. 12–17.
- [2] T. Kohonen, *Self-Organizing Maps*. Springer-Verlag, 2001.
- [3] J. Vesanto, "SOM-based data visualization methods," *Intelligent Data Analysis*, vol. 3, no. 2, 1999, pp. 111–126.
- [4] G. J. Deboeck and T. K. Kohonen, Eds., *Visual Explorations in Finance: with Self-Organizing Maps*. Springer, 1998.
- [5] J. Wei, Z. Shen, N. Sundaresan, and K.-L. Ma, "Visual cluster exploration of web clickstream data," in *IEEE VAST*, 2012, pp. 3–12.
- [6] T. Schreck, J. Bernard, T. von Landesberger, and J. Kohlhammer, "Visual cluster analysis of trajectory data with interactive kohonen maps," *Information Visualization*, vol. 8, no. 1, 2009, pp. 14–29, accessed June 13, 2014. [Online]. Available: <http://dx.doi.org/10.1057/ivs.2008.29>
- [7] J. Bernard, J. Brase, D. Fellner, O. Koepler, J. Kohlhammer, T. Ruppert, T. Schreck, and I. Sens, "A visual digital library approach for time-oriented scientific primary data," *Springer International Journal of Digital Libraries*, ECDL 2010 Special Issue, 2011, pp. 111–123.
- [8] K. U. Barthel, "Improved image retrieval using automatic image sorting and semi-automatic generation of image semantics," *Image Analysis for Multimedia Interactive Services*, International Workshop on, vol. 0, 2008, pp. 227–230.
- [9] N. Andrienko and G. Andrienko, *Exploratory Analysis of Spatial and Temporal Data: A Systematic Approach*. Springer-Verlag, 2005.
- [10] A. M. MacEachren, *How Maps Work - Representation, Visualization, and Design*. Guilford Press, 2004.
- [11] F. Bao, V. Lobo, and M. Painho, "The self-organizing map, the geo-som, and relevant variants for geosciences," *Computers & Geosciences*, vol. 31, no. 2, 2005, pp. 155 – 163, accessed June 13, 2014, *geospatial Research in Europe: AGILE 2003*. [Online]. Available: <http://www.sciencedirect.com/science/article/B6V7D-4F0851H-2/2/ac6bd99e54c5ff81db9a173604b0d3aa>
- [12] P. Agarwal and A. Skupin, Eds., *Self-Organising Maps: Applications in Geographic Information Science*. Wiley, 2008.
- [13] D. Guo, J. Chen, A. M. MacEachren, and K. Liao, "A visualization system for space-time and multivariate patterns (VIS-STAMP)," *IEEE Transactions on Visualization and Computer Graphics*, vol. 12, no. 6, 2006, pp. 1461–1474.
- [14] G. Andrienko, N. Andrienko, S. Bremm, T. Schreck, T. von Landesberger, P. Bak, and D. A. Keim, "Space-in-time and time-in-space self-organizing maps for exploring spatiotemporal patterns," *Computer Graphics Forum*, vol. 29, no. 3, 2010, pp. 913–92.
- [15] R. Kincaid, "SignalLens: Focus+context applied to electronic time series," *IEEE Trans. on Visualization and Computer Graphics*, vol. 16, 2010, pp. 900–907.
- [16] J. Van Wijk and E. Van Selow, "Cluster and calendar based visualization of time series data," in *Proc. IEEE Symposium on Information Visualization*. IEEE Computer Society, 1999, pp. 4–.
- [17] W. Aigner, S. Miksch, H. Schumann, and C. Tominski, *Visualization of Time-Oriented Data*. Springer-Verlag New York Inc, 2011.
- [18] J. Bernard, J. Brase, D. Fellner, O. Koepler, J. Kohlhammer, T. Ruppert, T. Schreck, and I. Sens, "A visual digital library approach for time-oriented scientific primary data," *Springer International Journal of Digital Libraries*, ECDL 2010 Special Issue, 2011, pp. 1–13.
- [19] M. Ward, "A taxonomy of glyph placement strategies for multidimensional data visualization," *Information Visualization*, vol. 1, no. 3/4, Dec. 2002, pp. 194–210.
- [20] D. Oelke, D. Spretke, A. Stoffel, and D. A. Keim, "Visual readability analysis: How to make your writings easier to read," *IEEE Transactions on Visualization and Computer Graphics*, vol. 18, no. 5, May 2012, pp. 662–674.
- [21] P. Juola, "Authorship attribution," *Found. Trends Inf. Retr.*, vol. 1, no. 3, Dec. 2006, pp. 233–334, accessed June 13, 2014. [Online]. Available: <http://dx.doi.org/10.1561/1500000005>
- [22] C. Rohrdantz, M. C. Hao, U. Dayal, L.-E. Haug, and D. A. Keim, "Feature-based visual sentiment analysis of text document streams," *ACM Transactions on Intelligent Systems and Technology*, Special Issue on Intelligent Visual Interfaces for Text Analysis, vol. 3, no. 2, 2012, pp. 26:1–26:25.
- [23] M. C. Hao, C. Rohrdantz, H. Janetzko, D. A. Keim, U. Dayal, L.-E. Haug, M. Hsu, and F. Stoffel, "Visual sentiment analysis of customer feedback streams using geo-temporal term associations," *Information Visualization*, Jun. 2013.
- [24] T. Kohonen, J. Hynninen, J. Kangas, and J. Laaksonen, "SOM_PAK: The self-organizing map program package," *Helsinki University of Technology*, Tech. Rep. A31, 1996.
- [25] T. Porter and T. Duff, "Compositing digital images," in *Proceedings of the 11th annual conference on Computer graphics and interactive techniques*, ser. SIGGRAPH '84. New York, NY, USA: ACM, 1984, pp. 253–259, accessed June 13, 2014. [Online]. Available: <http://doi.acm.org/10.1145/800031.808606>