Towards a Framework for Detecting and Managing Opinion Contradictions

Mikalai Tsytsarau
University of Trento
tsytsarau@disi.unitn.eu

Themis Palpanas
University of Trento
themis@disi.unitn.eu

***Abstract*—Sentiment Analysis gains in interest due to the large amount of potential applications and the increasing number of opinions expressed in particular in the Web. The focus of this paper is the development of a framework on top of sentiment analysis for detecting contradictions. First, we introduce a statistical model of contradictions based on a mean value and the variance of sentiments among different posts. It can be used to analyze and track sentiment evolution over time, to identify interesting trends and patterns or even to enable argument extraction. Using synthetic datasets, we demonstrate the effectiveness of our method in capturing contradictions on noisy data. Inspired by this model, which has proven to be effective and efficient for numeric sentiments, we are trying to generalize it for arbitrary opinion data and outline a universal framework which can be efficiently used on a large scale. We discuss various problems and challenges of such a formulation and outline the scope of our future work in this direction.**

I. Introduction

The problem of contradictions, or sentiment diversity on some topic, has been studied in the context of different research areas, having a slightly varying notion in each case. For instance, in Information Retrieval opposite opinions and sentiments introduce noise to the fact-centric search and must be avoided [7]. In contrast, conflicting sentiments is one of the desired targets of mining of product reviews.

We say that we have a contradiction when there are conflicting opinions for a specific topic, which is a form of sentiment diversity. This kind of contradiction can occur at one specific point of time or throughout a certain time period. Furthermore, a contradiction can occur within one text when an author presents different opinions on the same topic, or across texts when different authors express different opinions.

Recently proposed methods can aggregate opinions expressed in customer reviews and extract a representative summary of sentiments on a feature-by-feature basis; or they can capture and aggregate sentiments on some topic among different texts [5]. However, recent sentiment aggregation methods are not designed to track the evolution of sentiments on a large scale, neither they are suitable for contradiction detection. Therefore, this problem essentially requires a consistent definition and new methods to deal with it.

In our recent publications [9], [10] we discuss some ideas and methods tackling this problem. More specifically, we define the concepts of aggregated sentiment, sentiment variance and contradiction with respect to the time dimension, and formulated relevant problems of contradiction discovery. Our method operates on sentence-level sentiments, which are represented in a continuous scale. This allows us to exploit different approaches for sentiment detection, which can be plugged in our framework. The use of mean and variance for contradiction detection allows our method to be fast and linearly scalable on the number of texts, which is an important feature for large-scale analysis. Tests on real datasets, as well as a user-study, demonstrate that our approach is able to efficiently and effectively identify contradictions. Apart from this, there is still a need for the unified framework to work with contradictions, and many of its desired properties remain undefined.

The objective of this paper is to exploit the current work on sentiment analysis and contradiction analysis, and extend the state of the art by proposing a more solid and uniform framework.

The remainder of this paper is structured as follows. In Section II we discuss the related work, and in Section III we present the current framework for detecting and storing contradictions and formally define the new problem in Section IV. We discuss our experiences in Section V and conclude in Section VI.

II. Related Work

In the past few years, we have witnessed an increasing research interest in the area of subjectivity analysis and specifically in opinion mining and contradiction analysis [8]. Contradiction analysis is a rather new research area, with different studies focusing on different aspects of the problem.

There are few approaches where contradictions are defined as form of textual inference (entailment) and analyzed using linguistic technologies. Exploiting the contradiction features developed in [1], [3], and supplementing them by the sentence alignment tool, Pado et al. introduced contradiction detection approach to a textual entailment application [6]. Ennals et al. [2] extend this technique to operate with web data, describing an approach that detects contradicting claims by checking the entailment to a database of disputed web claims. These studies show that contradictions may occur not only on the opinion level, but also on the topic level, requiring the development of efficient methods for detecting contradictions on various data. However, linguistic analysis and textual entailment remain computationally complex tasks, incompatible with the large-scale operation.

A possible way of reducing the amount of computations lies in applying only lightweight linguistic analysis, followed by statistical contradiction detection methods. Kim and Zhai [4] proposed the system allowing retrieval and comparison of contradicting opinions based on the measures of *representativeness* and *contrastiveness*, whose linear combination is used as a criterion for optimization problem. The first measure is based on the weighted sums of maximal content similarities, among positive and negative sets of sentences and their corresponding summaries. Representativeness reflects how well the summaries approximate the original text. Contrastiveness captures the similarity between positive and negative sentences in the summaries, but is computed based on the contrastive similarity (that is the same as content similarity, except that it is computed without taking into account sentiment words).

However, to overcome the inherent complexity of linguistic technologies and make the analysis more scalable and universal with respect to input data, this problem should be considered from the perspective of data mining. Therefore, some researchers attempted applying classical data mining methods to textual documents.

For example, Varlamis et al. [11] propose clustering accuracy as an indicator of the blogosphere topic convergence. Clustering accuracy (when represented by the utility function) measures the relative separation of the cluster centers with respect to cluster sizes and a number of unclustered blogs (noise). When the clustering is very good, this function reaches its maximum value. It is easy to demonstrate, that divergence in topics leads to greater separation of individual blogs in the feature space and, therefore, more reliable clustering (for a fixed number of clusters). By analyzing how accurate the clustering is in different time intervals, one can estimate how correlated or diverse the blog entries are. We note that this approach is relevant to the contradiction detection, in the sense that clustering is often defined as the process of finding distant (i.e., contradicting) groups of similar (i.e., non-contradicting) items. However, the type of contradictions that this approach discovers depends on the selection of features.

III. Sentiment Contradiction Framework

The framework we describe in this section addresses the efficient detection of contradicting sentiments in texts on specific topics in case they can be represented by numeric polarity values. When using the term 'text' we refer either to the entire web document or its individual sentences. Nevertheless, this framework can accommodate other information documents, depending on the specific requirements of each application.

For each of the topics T discussed in some text, we wish to identify the sentiment polarity expressed towards it[1].

Definition 1 (Sentiment): The sentiment S with respect to a topic T is a real number in the range $[-1, 1]$ that indicates the polarity of the author's opinion on T expressed in a text.

Negative and positive values represent negative and positive opinions respectively, while the absolute value of sentiment represents the strength of the sentiment. Apart from computing sentiments for individual texts, we also need to compute the polarity on some topic aggregated over multiple texts (that may span different authors, as well as time periods).

Definition 2 (Aggregate Sentiment): The Aggregate Sentiment μ_S expressed in a collection of documents $\mathcal{D}$ on topic T, is defined as the average value over sentiments in $\mathcal{D}$.

Definition 3 (Sentiment Variance): The Sentiment Variance σ_S^2 is defined as the variance of sentiments expressed on topic T among documents in $\mathcal{D}$.

By comparing the sentiment values of different collections of texts, contradictions are identified as follows.

Definition 4 (Sentiment Contradiction): There is a contradiction on a topic T between two groups of documents, $\mathcal{D}_1, \mathcal{D}_2 \subset \mathcal{D}$ in a document collection $\mathcal{D}$, where $\mathcal{D}_1 \bigcap \mathcal{D}_2 = \emptyset$, when the information conveyed about T is considerably more different between $\mathcal{D}_1$ and $\mathcal{D}_2$ than within each one of them.

We define contradiction on a *pairwise* basis, where we evaluate the disagreement between two groups of documents in a collection. In this case, the similarity of information within each group serves as a reference point, providing a basic disagreement level. This property can lead to different implementations, and, as we demonstrate later, can be generalized by using a metric space to represent opinions.

What is also interesting, is that in the case of sentiments (which we consider as real values), we can estimate the level of contradiction without implicitly identifying the opposite groups of documents (although this can also be done). Here, we propose to estimate the sentiment contradiction level based on the aggregate sentiment and variance. The intuition behind this measure is that when the aggregate sentiment is close to zero, while the sentiment variance is high, then the contradiction level should be high. Combining μ_S and σ_S^2 in a single formula, we propose the following measure:

$$C = \frac{\vartheta \cdot \sigma_S^2}{\vartheta + (\mu_S)^2} W, \ \ \mu_S = \frac{1}{n}\sum_{i=1}^{n} S_i, \ \ \sigma_S^2 = \frac{1}{n}\sum_{i=1}^{n}(S_i - \mu_S)^2 \quad (1)$$

where n is the cardinality of $\mathcal{D}$, and W is a weight function that takes into account the (varying) number of posts that may be involved in the calculation. Also, there is a small value added the denominator, $\vartheta \neq 0$, which allows to limit the level of contradiction C when $(\mu_S)^2$ is close to zero. The nominator is multiplied by ϑ to ensure that contradiction values fall within the interval $[0; 1]$.

Problem 1 (Contradiction Detection):
For a given set of documents $\mathcal{D}$, and topic T, identify whether a contradiction level for T is exceeding some threshold ρ.

In order to detect contradicting opinions in collections of texts, we first need to determine all the different topics and then calculate the corresponding sentiments. Assume that we want to look for contradictions in a shifting time window w. For a particular topic T, the set of documents $\mathcal{D}$, which we use for calculation, will be restricted to those, that were posted within the window w. We denote this set as $\mathcal{D}(w)$, and n as its cardinality, $n = |\mathcal{D}(w)|$.

[1] In the following, we refer to sentiment polarity simply as *sentiment*.

IV. Opinion Contradiction Framework

In the previous section, we only considered sentiment-based contradictions. We now turn our attention to other forms of opinion contradictions, and formulate the problem in a more general context of opinions.

Definition 5 (Opinion): The opinion O represents a statement or claim expressed by the author on topic T in a text.

Unlike *sentiment*, the opinion can be either an objective statement, e.g. "car is *black*", or a subjective statement, e.g. "war is *bad*". In fact, there exist a wide range of different types of opinions. In this work, we are interested in contradicting ones, i.e. those that have no sense together. For example, "car is *black* and *white*", or "war is *good* and *bad*".

To build a general framework, we purposely not specify the exact type of opinion, but rather propose to formulate all differences between opinions in a form of distance function.

Definition 6 (Opinion Distance): The opinion distance $d(x, y) = \parallel x - y \parallel \in \mathbb{R}^+$ is a positively-defined (multi-dimensional) function that satisfies to the conditions of semi-metric:

$$\begin{cases} d(x,y) \geq 0 \\ d(x,y) = 0 \text{ if and only if } x = y \\ d(x,y) = d(y,x) \end{cases}$$

Definition 7 (Aggregate Opinion): Aggregate Opinion $\overline{O}$ is an opinion with the closest accumulative distance to other opinions within a group:

$$\overline{O} = \underset{O}{\operatorname{argmin}} \sum_{O_i \in \mathcal{D}} \parallel O - O_i \parallel^2$$

Definition 8 (Opinion Variance): Opinion Variance σ_O^2 is the average distance between opinions in $\mathcal{D}$ and Aggregate Opinion $\overline{O}$:

$$\sigma_O^2 = \frac{1}{n} \sum_{O_i \in \mathcal{D}} \parallel O_i - \overline{O} \parallel^2$$

By comparing opinion values of different collections of texts contradictions are identified as follows:

Definition 9 (Opinion Contradiction): A collection $\mathcal{D}$ of texts talking about topic T, is considered contradictory, if it can be partitioned into several groups of texts $\mathcal{D}_i \subset \mathcal{D}$ such that the distance between aggregate opinions of any two groups is at least α times greater than the maximum opinion variance:

$$\min_{i \neq j} \parallel \overline{O}(\mathcal{D}_i) - \overline{O}(\mathcal{D}_j) \parallel^2 \; > \alpha \cdot \max_k \sigma_O^2(\mathcal{D}_k) \tag{2}$$

This definition allows us to detect contradictions, but does not assess their strength. For this purpose, we define our measure for contradiction C based on the number and size of contradicting groups: the largest contradiction occurs when there are many groups of equal sizes.

Problem 2 (Opinion Contradiction Detection): Partition a given collection of documents $\mathcal{D}$ into a minimal number of non-intersecting sub-groups $\mathcal{D}_i \cap \mathcal{D}_j = \emptyset$, such that Equation 2 still holds, and compute the level of contradiction:

$$C = - \sum_{\mathcal{D}_i \in \mathcal{D}} \frac{|\mathcal{D}_i|}{|\mathcal{D}|} \cdot \log \frac{|\mathcal{D}_i|}{|\mathcal{D}|} \tag{3}$$

V. Discussion

In this section we discuss various properties relevant to both of our frameworks, as well as indicate their key differences.

A. Detecting Contradictions

When identifying contradictions in a document collection, it is important to also take into account the time in which these documents were published. Let $\mathcal{D}_1$ be a group of documents published within some time interval t_1. Assume that t_1 is followed by time interval t_2, with the documents $\mathcal{D}_2$ containing a conflicting piece of information. In this case, we have a special type of contradiction, which we call *Asynchronous Contradiction*. Following the same line of thought, we say that we have a *Synchronous Contradiction* when both $\mathcal{D}_1$ and $\mathcal{D}_2$ are mixed in a single time interval, t.

Unlike with sentiment contradictions, where the *change of sentiment* can be detected by looking at the sign of aggregate sentiment, *change of opinion* is related to the change of the prevailing group of documents (the largest one).

When trying to detect contradictions, we would like to identify those that have a contradiction value above some threshold ρ. We refer to this solution as *fixed threshold*. In order to better fit a threshold to the nature of the data, we propose an *adaptive threshold* technique, which computes a different threshold for each topic and time window based on the global value for C as follows: $\varrho = p \cdot C$, $0 < p < 1$.

Note that we cannot achieve the same result by using *top-k* queries, since adaptive threshold does not impose a strict limit on the number of contradictions in the result, and can thus report the entire set of interesting contradictions within some time interval.

B. Properties of Sentiment Contradiction Measure

Figure 1 shows the operation of the proposed contradiction measure on synthetic data. Using this dataset, we verify the ability of the C function to capture the planted contradictions. The graph at the top (a) shows generated sentiments as points distributed around the planted trend, showing an initial positive sentiment that later changes to negative (at time instance t_1), which represents a change of sentiment. There is also a point around time instance t_2, where the sentiments are divided between positive and negative, a situation representing a synchronous contradiction. As can be seen in (b), μ_S closely captures the aggregate trend of the raw sentiments. The contradiction value (c), calculated using a sliding window, demonstrates the correct detection of two planted contradictions regardless of the added noise.

To detect a peak in contradiction we need to use a time window of a proper granularity. Smaller time windows will allow us to detect more simultaneous contradiction, while larger ones will reveal opposite opinions, which are sparse across time. Thus, we need to analyze time series of contradiction level using different granularities. To achieve an efficient computation of such time series, we need a data storage of a proper structure.

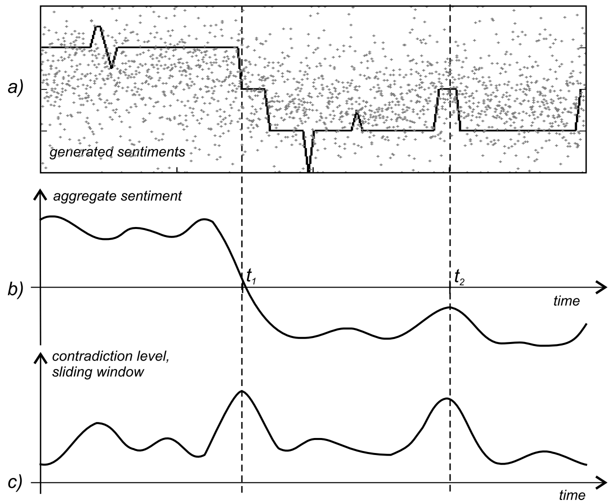

Fig. 1: Example of contradiction values computed from a synthetic dataset with two planted contradictions.

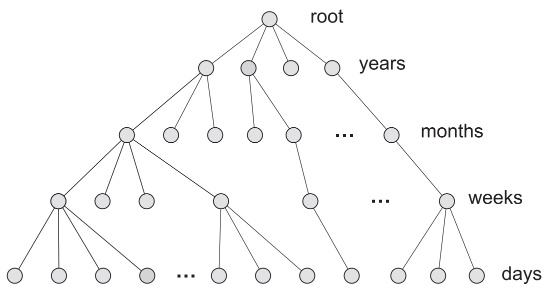

Fig. 2: Logical representation of the time hierarchy in CTree.

C. Storing Contradictions

We now turn our attention to the problem of organizing sentiment data in a way that will allow the efficient detection of contradictions in large collections of data that span very long time intervals.

The proposed solution stores levels of contradiction in a time-tree structure (CTree), demonstrated in Figure 2. Here, the nodes of each level of the tree correspond to a certain time granularity, resulting in a tree with fixed number of levels. Each node of the tree summarizes sentiments or opinion statistics of a given time interval. This gives us additional flexibility, since we can now compute the contradiction of a large time window by composing the corresponding values from the smaller windows contained in the large one.

While the Formula 1 that calculates the sentiment contradiction is based on the mean and variance, allowing us to compute it using incrementally updateable measures, it is not obvious that in the case of opinion contradictions the metric space would allow hierarchical aggregation of clusters so that Equation 2 will hold. Therefore, we should study and evaluate different metric functions, which not only allow computing Equation 2 from aggregate data, but also provide means to verify this inequality. Provided that the above properties hold, opinion contradiction level C can be efficiently computed from the aggregate group size counts using Formula 3.

VI. Conclusion

In this paper we describe the framework for mining sentiment contradictions and outline its more general version for opinion contradictions. We discuss whether we can use the same or similar techniques for the generalized contradiction problem, as the ones we used in our previous work.

According to our evaluation, some of the recent methods can be seamlessly applied to solve similar problems in the new framework (e.g. adaptive threshold, synchronous/asynchronous contradiction types), while some others (CTree aggregating storage, updateable measures, aggregate opinion) require a more careful modeling of opinions.

The most challenging problems in this direction include providing an opinion contradiction definition, which accepts an arbitrary number of document groups, and studying various metric spaces (and their properties), which allow effective comparison and contrasting of opinions, as well as the storage of aggregate opinion and opinion variance as updateable moments. Another interesting direction for our research would be evaluating the existing clustering frameworks and their applicability to our problems.

In our further investigation, we are going to address the above problems and refine the proposed framework. Moreover, we would like to draw a mathematical connection between the two frameworks, so that it would be possible to consider one as a restricted version of another, and better understand their properties.

References

[1] M. C. de Marneffe, A. N. Rafferty, and C. D. Manning. Finding contradictions in text. In *Proceedings of ACL-08: HLT*, pages 1039–1047, Columbus, Ohio, June 2008. ACM.

[2] R. Ennals, B. Trushkowsky, and J. M. Agosta. Highlighting disputed claims on the web. In *Proceedings of the 19th International Conference on World Wide Web, Raleigh, USA, April 26-30*, 2010.

[3] S. Harabagiu, A. Hickl, and F. Lacatusu. Negation, contrast and contradiction in text processing. In *Proceedings of the 21st National Conference on Artificial Intelligence*, pages 755–762. AAAI Press, 2006.

[4] H. D. Kim and C. Zhai. Generating comparative summaries of contradictory opinions in text. In *Proceedings of the 18th ACM conference on Information and knowledge management*, pages 385–394. ACM, 2009.

[5] K. Lerman, S. Blair-Goldensohn, and R. Mcdonald. Sentiment summarization: Evaluating and learning user preferences. In *12th Conference of the European Chapter of the Association for Computational Linguistics (EACL-09)*, 2009.

[6] S. Pado, M.-C. de Marneffe, B. MacCartney, A. N. Rafferty, E. Yeh, and C. D. Manning. Deciding entailment and contradiction with stochastic and edit distance-based alignment. In *Proceedings of the First Text Analysis Conference (TAC 2008)*, 2009.

[7] E. Riloff, J. Wiebe, and W. Phillips. Exploiting subjectivity classification to improve information extraction. In *Proceedings of the 20th National Conference on Artificial Intelligence*, pages 1106–1111. AAAI Press, 2005.

[8] M. Tsytsarau and T. Palpanas. Survey on mining subjective data on the web. *Data Mining and Knowledge Discovery, Special Issue on 10 Years of Mining the Web, accepted for publication*, pages 1–40, 2011.

[9] M. Tsytsarau, T. Palpanas, and K. Denecke. Scalable discovery of contradictions on the web. In *Proceedings of the 19th International Conference on World Wide Web, WWW 2010, Poster Track*, 2010.

[10] M. Tsytsarau, T. Palpanas, and K. Denecke. Scalable detection of sentiment-based contradictions. In *First International Workshop on Knowledge Diversity on the Web, Colocated with WWW 2011*, 2011.

[11] I. Varlamis, V. Vassalos, and A. Palaios. Monitoring the evolution of interests in the blogosphere. In *ICDE Workshops*, pages 513–518. IEEE Computer Society, 2008.