From Sentiment Analysis to Preference Aggregation

Umberto Grandi,[1] **Andrea Loreggia,**[1] **Francesca Rossi**[1] **and Vijay A. Saraswat**[2]
[1] University of Padova, Italy
umberto.uni@gmail.com, andrea.loreggia@gmail.com, frossi@math.unipd.it
[2]IBM TJ Watson Research Lab
vijay@saraswat.org

Abstract

Current sentiment analysis techniques are good enough when predicting the opinion of a community of individuals over one item, but they may produce wrong or inaccurate results when several possibly correlated items are under consideration. We propose to exploit and adapt formalisms, methods, and tools from knowledge representation and voting theory to generalize sentiment analysis and obtain an accurate definition of collective sentiment about multiple items.

Our vision

Design agents that are able to interpret correctly the collective sentiment of a society or a multiagent system from individual opinions expressed in natural language.

Sentiment analysis

Sentiment analysis (Pang and Lee 2008; Liu 2012) combines the opinions of individuals about a collection of entities (such as products, candidates or stocks) into an overall sentiment. Opinions are extracted from textual representations in blogs, reviews, tweets, Facebook posts etc using information retrieval techniques and combined using a variety of aggregation rules.

The overall opinion of a text is usually identified in a positive, negative or neutral polarity. More complex approaches extend this 3-valued polarity to a scale, such as the "5-stars" approach used in several rating systems, defining different notions of graded polarity. Once a set of individual opinions has been extracted in the form of polarities or graded polarities, this information is aggregated into a collective sentiment.

In recent years, sentiment analysis has received significant attention in industry and in the media. A number of web tools, commercial products and applications have recently appeared in this space.[1]

However, in our view, basic foundational questions have not received enough critical attention.[2]

From sentiment analysis to preference aggregation

Sentiment analysis is adequate when summarizing collective sentiments about a single topic, either through a polarity (positive, negative or neutral), or through more complex approaches (such as a "5-stars" scale). More sophisticated techniques are called for, however, as soon as we wish to develop collective judgments comparing two or more objects.

If individuals express more than two options (like/dislike) over several items, then we believe *preferences* (Rossi, Venable, and Walsh 2011; van Harmelen et al. 2007) and voting theory (Arrow, Sen, and Suzumura 2002) are more appropriate than polarities. In our view, insisting in modeling opinions on several items via (possibly graded) polarities may lead to incorrect or very inaccurate results in terms of understanding the overall opinion of a community.

The literature on preference representation and aggregation provides a variety of definitions, tools and concepts which can be applied to the formulation of collective sentiment. More generally, we believe the area of *computational social choice* (Brandt, Conitzer, and Endriss 2013), concerned with computational properties of procedures for collective choice in Artificial Intelligence and multi-agent systems, is very relevant.

Motivating example

Consider the following simple problem: two candidates a and b are competing in an election which will be decided using the majority rule, i.e., the candidate receiving a majority of the votes is elected. Utterances by voters about the candidates are available on social media, blogs and news forums. We are interested in predicting the winner of the election based on these utterances.

Assume that the electorate is composed of 100 individuals and that 90 individuals express a positive opinion on both a and b, but when asked to choose would prefer a to b. The remaining 10 individuals have a positive opinion on b and a negative opinion on a (and thus will choose b over a at the

[1]See, e.g., www.sentiment140.com, www.social-trends.it, www.thestocksonar.com, www.nytimes.com/interactive/us/politics/2010t-twitter-candidates.html

[2]Though we note, in passing, the work of Metaxas *et al*, (2011) in sketching a road map for developing sentiment analysis as an alternative to opinion polls for the prediction of electoral results.

time of the election). The situation can be represented in the following table in which preferences are read from left to right and the bar signals the threshold of positive sentiment, i.e., candidates to the left of the bar are approved as positive.

90 voters	a	b	|	
10 voter		b	|	a

Majority rule winner: a
Collective sentiment predictor: b

In this example, sentiment analysis would detect 100% positive opinions for b, and thus predict that she will be the winner of the election. However, when individuals will be asked to vote, a will win over b by an overwhelming majority (90 vs 10).

Compact preference representation and voting rules

Motivated by the previous example, we list six challenges as a road map for collaboration between the areas of sentiment analysis and preference aggregation/voting theory.

We propose to use models of preferences that are more expressive than just polarities when representing individual opinions extracted from text. Examples include rankings or partial orderings, as well as more complex structures such as conditional preference networks. This requires a deep understanding of how to adapt existing known representation frameworks (see, e.g., Rossi *et al*, 2011) to modeling the kind of preferences and sentiments that can be expressed in textual corpora such as those collected on blogs or Twitter posts.

We also propose to use voting rules to aggregate the preferences of the individuals and to be used as possible definitions of the sentiment of a collectivity (Brams and Fishburn 2002). Methods from the field of voting theory, knowledge representation and preference aggregation should therefore be adapted in order to model the opinion of each individual and aggregate these individual opinions into a collective one. A crucial aspect is the weight to be given at aggregation time to the information regarding individual sentiments, such as the polarity or the approvals, and the information regarding individual preference comparisons. Related work on this topic has been carried out in the literature on social choice theory (Brams and Sanver 2009) and can serve as a starting point for this line of research.

Challenge 1: What preferences/opinions can be extracted from the individuals' text?

Given a set of textual expressions written by one individual, there are two main forms of opinion that can be expressed concerning a set of items:

- *Regular* opinions, in which an individual expresses her degree of (dis)satisfaction on one single entity;
- *Comparative* opinions, in which the best entity is identified among a set of two or more alternatives.

Current sentiment analysis techniques can be used to extract a (graded) polarity identifying the opinion of an individual over the items under consideration (see, e.g., Pang *et al*, 2002, and Baccianella *et al*, 2010). Less attention has instead been given to the extraction of comparative opinions (see, e.g., Jindal and Liu,2006, and Ganapathibhotla and Liu, 2008).

It seems therefore that the basic information that can be extracted regarding individual opinions over a set of items consists of the two following objects:

- *Graded polarities* indicating the degree of positivity and negativity of a group of items (by for instance using a common scale, or numbers in a given interval, or unrestricted real numbers).
- A set of *binary comparisons* of the form $a \geq b$ stating which of two items is preferred to the other one.

A precise mathematical formulation of these two objects is needed. More work should also be done to evaluate the frequency of comparative as well as other forms of opinion in textual corpora, and refine existing natural language processing techniques for their extraction. Two additional parameters may play a role in the extraction of preferences and opinions:

- *Time*: individual opinions are expressed at a given point in time, and this information may play an important role in defining a coherent individual view. We point to the literature on knowledge representation (van Harmelen et al. 2007), in particular belief revision and merging, which provides tools for the analysis and summarization of conflicting information.
- *Features*: entities or items may be described by means of features, i.e., they may be elements of a product space. Techniques from natural language processing can be used to extract the relevant features and thus build the set of entities. However, in this setting preferences and opinions may compare features rather than entities, requiring a more elaborate framework for its extraction and representation. Moreover, the combinatorial explosion resulting from a large set of features may give rise to computational problems that require an adequate compact representation framework. The literature on social choice in combinatorial domains (Chevaleyre et al. 2008) is highly relevant to this problem.

Challenge 2: What are the best representations for individual preferences and sentiments?

Individual preferences and opinions can be represented with a wide variety of models (Rossi, Venable, and Walsh 2011). One possible parameter in this choice is the structure of the items under consideration. In case the set of items is too large then it might be infeasible to list them all in a ranked order, according to the preferences of the individual. This can happen when alternatives are described by means of multiple features such as in a combinatorial domain (Chevaleyre et al. 2008). For example, a car has features such as color, brand, model, engine, and so on, each of which may have several options.

Preferences may also be expressed directly on features, making it more difficult to identify the most preferred alternatives if for instance the set of preferential statements is cyclic or preferential dependencies are present. This happens in case the preference of an individual concerning a certain feature depends on the choice that has been done on a different feature. Continuing on the car example, one individual might prefer a red car but only if it is a sporty car, otherwise she prefers a white car. Several frameworks have been proposed in the literature to compactly represent preferences in these complex domains, notable examples being conditional preference networks and soft constraints.

Observe also that the preferences and opinions could be expressed over time (e.g., tweets) and could therefore contain contradicting or updating information. Preference modeling should take this into account to remove/mitigate the contradictions and generate a consistent model of the preferences of an individual.

We identify two features of preferences extracted through sentiment analysis techniques:

- *Interpersonal incomparability*: Since individuals have very different styles of writing or attitudes towards judging the entities under consideration, we believe that the scores of graded polarity cannot be compared across individuals. Thus, we argue in favor of an ordinal representation by means of binary comparisons to model both regular and comparative opinions about a given set of entities.
- *Incompleteness*: Since preferences and sentiments are observed from individual expressions and are not submitted voluntarily by individuals as happens in classical settings of social choice theory, we cannot assume this information to be complete and hence there will always exist incomparable alternatives (more precisely, alternatives for which no preference or sentiment is being expressed).

Starting from these two principles we put forward the research problem of designing the most suitable representation that can combine the sentiment and preference information extracted from a set of individual expressions.

A simple proposal which generalizes the approach of Brams and Sanver (2009) could represent the sentiment as a partition of the alternatives in three subsets indicating the positive, neutral and negative elements, adding a preorder over these subset which specifies binary comparisons among alternatives. This can be visualized in the following figure, in which alternatives a_j at higher levels are preferred to alternatives at lower levels, and the three sets are separated by horizontal lines.

a_1, a_2 — a_3	Positive
a_4, a_5	Neutral
a_6 — a_7	Negative

Challenge 3: How to aggregate the individuals' preferences into a collective opinion?

Once the preferences of each individual have been extracted and represented they should be aggregated to derive the opinion of the entire community of individuals over the set of items under discussion. When only two items are considered then there is no question on what voting rule to use to perform preference aggregation: the majority rule has been proven to be the only "perfect" voting rule in such situations (May 1952). However, when considering preferences over more than two items there are many ways to perform the aggregation. The literature on voting theory and preference aggregation (see, e.g., Brams and Fishburn, 2002) contains a number of possible voting rules to be used (Plurality, Borda, Kemeny, Copeland, Approval, you name it), each with different desirable and undesirable properties, and each one possibly giving a different result.

Thorough studies should be performed to identify the most suitable voting rule for the definition of the collective sentiment in different domains. However, existing voting rules cannot be used directly, since they are usually defined on input composed of a collection of total orders, thus disregarding the polarity information about sentiments and assuming completeness of preferences. Previous work in the literature on voting with partial orders (see, e.g., Fagin *et al*, 2004, Cullinan, 2013) and on combining sentiment/approval with preference (Brams and Sanver 2009) should therefore be complemented to obtain meaningful aggregation procedures to be used on sentiment and preference information. For instance, if each individual submits two partial orders together with polarity information as in the previous example, a generalized notion of the Borda count can be devised to work on this more complex input.

Challenge 4: Is it possible to automatically identify "influencers" and prevent strategic behavior?

The individuals composing a society, as well as the agents in a multiagent system, are very often connected by interpersonal ties, e.g., when individuals are organised in a network. In this case, individual preferences and opinions are not only the result of personal reflection but may also take into consideration the position taken by influential individuals or simply by agents that are close to them in the network.

The field of *social network analysis* (see, e.g., Jackson, 2010, Easley and Kleinberg, 2010) is a burgeoning research area which has the potential of generating highly interesting results once combined with techniques of preference and sentiment analysis. For instance, one may ask whether it is possible to extract automatically the influence structure, e.g., a graph, from the collected set of individual expressions. Influential individuals may then be identified, and the definition of collective sentiment may be changed accordingly. Computation-friendly models of influence in situations of collective choice have been proposed (Maran et al. 2013), but more work is required to a deeper understanding of this problem.

Sentiment analysis techniques are moreover not immune to strategic manipulation. A rising phenomenon is the creation of web services proposing the opening of thousands of fake Twitter accounts to be used as followers of the manipulator's account, or the publishing of big volumes of positive posts related to the manipulator's products. This represents a prime example of strategic behavior in collective choice problems, and the whole body of literature published on this topic may be put to test with real world data once the two fields of sentiment analysis and preference aggregation have been put together to their full potential.

Challenge 5: How should preference aggregation methods be validated?

Since the variety of preference aggregation methods that can be defined is very large, a natural question is how to choose among them: What is the best method to aggregate individual opinions about a set of items into the collective opinion of the community of individuals?

- If methods of collective sentiment analysis are used over time, tested for several settings and items, and employed in the context of predicting the result of real-world processes (such as elections or the evolution of a market), then *machine learning* techniques can be used to learn the best aggregation method, that is, the one that has proven to be the most accurate.
- Work on voting rules seen as *maximum likelihood estimators* can also be useful in this respect (Conitzer and Sandholm 2005). Alternatively, as in classical voting theory, axiomatic properties as well as results about the computational complexity of aggregation rules could guide the choice of some aggregation methods over others.
- Comparative opinions and binary comparisons may not be very frequent in user generated content. Their presence and use should be tested on a corpus of textual expressions. Similar experiments may contribute to the problem of analyzing the *preference structure and distribution* expressed by individuals, an important open task in the research area on preferences (Mattei and Walsh 2013).

Challenge 6: How to deal with big data in sentiment and preference analysis?

When the aggregation operation is relatively simple (e.g., the majority rule), it is possible to use straightforward techniques such as Hadoop MapReduce (Dean and Ghemawat 2008) to perform computations in parallel. The mapping phase can be used to run sentiment classifiers on text corpora in parallel; the resulting data objects can be combined/reduced in parallel.

However, with more complex structures (e.g., conditional preference networks), the combination procedure may be more combinatorial in nature and may require non-trivial parallel processing. In this context, modern scale-out programming languages such as X10 can be particularly valuable (Charles et al. 2005), making it easy to write code that runs over thousands of cores and deals with hundreds of gigabytes of main memory data. Of particular interest is developing incremental parallel algorithms that can update collective sentiments as new utterances stream in and need to be processed.

Conclusion

In sum, we contend that the field of preferences, voting theory and computational social choice provide a rich framework for sentiment analysis and suggest several directions for future work.

References

Arrow, K. J.; Sen, A. K.; and Suzumura, K., eds. 2002. *Handbook of Social Choice and Welfare*. Elsevier.

Baccianella, S.; Esuli, A.; and Sebastiani, F. 2010. Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In *Proceedings of LREC-2010*.

Brams, S. J., and Fishburn, P. C. 2002. Voting procedures. In Arrow, K.; Sen, A.; and Suzumura, K., eds., *Handbook of Social Choice and Welfare*.

Brams, S., and Sanver, M. R. 2009. Voting systems that combine approval and preference. In Brams, S.; Gehrlein, W.; and Roberts, F., eds., *The Mathematics of Preference, Choice and Order*.

Brandt, F.; Conitzer, V.; and Endriss, U. 2013. Computational social choice. In Weiss, G., ed., *Multiagent Systems*. MIT Press.

Charles, P.; Grothoff, C.; Saraswat, V.; Donawa, C.; Kielstra, A.; Ebcioglu, K.; von Praun, C.; and Sarkar, V. 2005. X10: an object-oriented approach to non-uniform cluster computing. In *Proceedings OOPSLA-2005*.

Chevaleyre, Y.; Endriss, U.; Lang, J.; and Maudet, N. 2008. Preference handling in combinatorial domains: From AI to social choice. *AI Magazine* 29(4):37–46.

Conitzer, V., and Sandholm, T. 2005. Common voting rules as maximum likelihood estimators. In *Proceedings of UAI-2005*.

Cullinan, J.; Hsiao, S.; and Polett, D. 2013. A Borda count for partially ordered ballots. *Social Choice and Welfare* 1–14.

Dean, J., and Ghemawat, S. 2008. MapReduce: simplified data processing on large clusters. *Communications of the ACM* 51(1):107–113.

Easley, D., and Kleinberg, J. 2010. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press.

Fagin, R.; Kumar, R.; Mahdian, M.; Sivakumar, D.; and Vee, E. 2004. Comparing partial rankings. *SIAM Journal on Discrete Mathematics* 20:47–58.

Ganapathibhotla, M., and Liu, B. 2008. Mining opinions in comparative sentences. In *Proceedings of COLING-2008*.

Jackson, M. O. 2010. *Social and Economic Networks*. Princeton University Press.

Jindal, N., and Liu, B. 2006. Mining comparative sentences and relations. In *Proceedings of AAAI-2006*.

Liu, B. 2012. *Sentiment Analysis and Opinion Mining*. Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers.

Maran, A.; Maudet, N.; Pini, M. S.; Rossi, F.; and Venable, K. B. 2013. A framework for aggregating influenced cp-nets and its resistance to bribery. In *Proceedings of AAAI-2013*.

Mattei, N., and Walsh, T. 2013. Preflib: A library of preference data. In *Proceedings of ADT-2013*.

May, K. O. 1952. A set of independent necessary and sufficient conditions for simple majority decision. *Econometrica* 20(4):680–684.

Metaxas, P.; Mustafaraj, E.; and Gayo-Avello, D. 2011. How (not) to predict elections. In *Proceedings of PASSAT-2011 and SOCIALCOM-2011*.

Pang, B., and Lee, L. 2008. Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval* 2(1-2):1–135.

Pang, B.; Lee, L.; and Vaithyanathan, S. 2002. Thumbs up? sentiment classification using machine learning techniques. In *Proceedings of EMNLP-2002*, 79–86.

Rossi, F.; Venable, K. B.; and Walsh, T. 2011. *A Short Introduction to Preferences: Between Artificial Intelligence and Social Choice*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers.

van Harmelen, F.; van Harmelen, F.; Lifschitz, V.; and Porter, B. 2007. *Handbook of Knowledge Representation*. Elsevier.