

A Borda Count for Collective Sentiment Analysis

Umberto Grandi

Department of Mathematics
University of Padova

10 February 2014

Joint work with Andrea Loreggia (Univ. Padova),
Francesca Rossi (Univ. Padova) and Vijay Saraswat (IBM Yorktown)

What is the collective sentiment about ...?

www.sentiment140.com/search?query=Paris&hl=en

Sentiment140


Tweet 697 Like 376 +1 157

Paris

English

Search

Sentiment analysis for Paris



Tweets about: Paris

ellie_mckone: @ParisElizabeth_ happy birthday paris!!! Have a great day babe


Posted: 20 seconds ago

kristenalperen: RT @KStewartINA: [latest] new picture of Kristen at LAX from

Posted: 22 seconds ago

thefreddiek: In Paris they dont even have human drivers, its all automatic & ins

Posted: 31 seconds ago



Closing the loop between governments & citizens


www.wegov-project.eu



on it #tubestrike


ery are: Accurate Inaccurate

Aggregation of individual polarities (like/dislike)



Multiple alternatives: SA is not enough!

Sentiment analysis (polarity aggregation) works well on a single issue
But what if we want to **compare multiple alternatives**?



20 people	<i>a</i>	<i>b</i>		
10 people		<i>b</i>		<i>a</i>
5 people				<i>a</i> <i>b</i>
<hr/>				
Sentiment analysis				<i>b</i>
Majority aggregation				<i>a</i>

Sentiment analysis predicts **a different winner** than majority aggregation!

Towards human-level decision making

What does sentiment analysis have to do with MAS?

- A wide **testing ground** for many MAS techniques: same methods used for agents coordination and collective decision making among humans

Computational Social Choice, Algorithmic Decision Theory...

- Our goal: Design **agents that are able to interpret correctly the collective sentiment** of a society or a multiagent system

Example: automated trading agents

Brandt, Conitzer, and Endriss. Computational social choice. In *Multiagent Systems*, 2013.

Wellman, Greenwald, and Stone. *Autonomous Bidding Agents: Strategies and Lessons from the Trading Agent Competition*, 2007.

Outline

1. Basic definitions: Sentiment analysis and preference aggregation
2. Challenge I: Data extraction and representation
 - pure sentiments (polarity)
 - pure preference (preorder)
 - sentiment and preferences (**SP-structures**)
3. Challenge II: Aggregation of individual opinions
 - the **Borda*** rule
4. Challenge III: Methods for validation
 - Predictive accuracy: Truth tracking power
 - Individuals' satisfaction: **Axiomatic properties**
5. Future work: Influence detection and strategic behavior

Sentiment Analysis

Ingredients:

- An **entity** x (no assumption about its structure)
- A **corpus** of individual expressions \mathcal{T} by a set of individuals \mathcal{I}
- A notion of **polarity**: $\{+, -, N\}$, 5-star scale or graded sentiment

Several NLP techniques used to extract the collective sentiment:

- entity extraction to find expressions mentioning x in \mathcal{T}
- word-count, Naive Bayes, and other **machine learning** techniques to extract the polarity of a single expression in \mathcal{T}

The collective sentiment about an issue x is positive if the **percentage of positive expressions** about x outnumbers that of negative ones.

Voting Theory

Ingredients:

- A set of alternatives \mathcal{X}
- A set of agents \mathcal{I} expressing preferences on \mathcal{X}

Voting rules are used to identify a set of **most preferred candidates**.
Several rules are possible!

We focus on two definitions of voting rules:

Borda rule - linear orders: if a voter ranks candidate c at j -th position this gives j points to c . The alternatives with highest score are the winners.

Approval voting - sets of candidates: the winners of approval voting are the candidates which receive the highest number of approvals.

Challenge I:

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

What kind of data structure can be extracted from text?

partial scores



4.5

5-star ratings



binary comparisons



-3



Pang and Lee. *Opinion mining and sentiment analysis*. Foundations and Trends in Information Retrieval, 2008.

Liu. *Sentiment Analysis and Opinion Mining*. Morgan & Claypool Publishers, 2012.

Individual Data

Two forms of opinions can be extracted with existing NLP techniques:

- Objective opinions: "Nikon is a good camera" → **score** of a single entity
- Comparative opinions: "I prefer Canon to Nikon" → **binary comparisons**

Let \mathcal{T}_i be the set of all expressions of individual's i about objects in \mathcal{X} :

Definition

The **individual data** extracted from \mathcal{T}_i is a tuple (σ_i, P_i, N_i) :

- $\sigma_i : D_i \rightarrow \mathbb{R}$ to represent objective opinions on $D_i \subseteq \mathcal{X}$
- subsets P_i and N_i of \mathcal{X} preordered by \leq_i^P and \leq_i^N , representing positive and negative comparative opinions

b 3

Example:

a -3 > c

The sentiment analysis approach

Definition

The *pure sentiment data* associated with raw data (σ_i, P_i, N_i) is a function $S_i : \{D_i \cup P_i \cup N_i\} \rightarrow \{+, -, 0\}$ defined as:

$$S_i(c) = \begin{cases} \text{sgn}(\sigma_i(c)) & \text{if } c \in D_i \setminus (P_i \cup N_i) \\ 0 & \text{if } \sigma_i(c) = 0 \\ + & \text{if } c \in P_i \\ - & \text{if } c \in N_i \end{cases}$$

Example:

b

a

c


The preference aggregation approach

Definition

The pure preference data associated with raw data (σ_i, P_i, N_i) is a *preordered* set $(\mathcal{D}_i, \leq_i^{\mathcal{D}})$ where $\mathcal{D}_i = D_i \cup P_i \cup N_i$ and

$$x \leq_i^{\mathcal{D}} y \Leftrightarrow \begin{cases} x \leq_i^P y \text{ and } x, y \in P_i & \text{or} \\ x \leq_i^N y \text{ and } x, y \in N_i & \text{or} \\ x \in N_i \text{ and } y \in P_i & \text{or} \\ \sigma_i(x) \leq \sigma_i(y) \text{ and } x, y \in D_i & \end{cases}$$

Example:



Our approach: SP-structures

Definition

An SP-structure over \mathcal{X} is a tuple $(\mathcal{P}, \mathcal{N}, \mathcal{Z})$ such that:

- \mathcal{P} , \mathcal{N} and \mathcal{Z} form a partition of \mathcal{X}
- \mathcal{P} and \mathcal{N} are ordered respectively by *preorders* $\leq^{\mathcal{P}}$ and $\leq^{\mathcal{N}}$

An SP-structure $(\mathcal{P}_i, \mathcal{N}_i, \mathcal{Z}_i)$ can be extracted from raw data (σ_i, P_i, N_i) :

- \mathcal{P}_i is the union of P_i and the set of entities with positive score
- Analogously for \mathcal{N}_i . \mathcal{Z}_i is the set of entities with zero or no score
- Preordered relations extracted from σ_i and copied from P_i and N_i

SP-structures combine (interpersonally non-comparable) scores with (incomplete) pairwise comparisons between entities

Challenge II:

How to aggregate individual information
into a collective opinion?

Aggregating SP-structures

Definition

The **Borda*** score of entity $c \in \mathcal{X}$ in SP-structure $(\mathcal{P}, \mathcal{N}, \mathcal{Z})$ is defined as:

$$s^*(c) = \begin{cases} 2 \times |\text{down}^{\mathcal{P}}(c)| + |\text{inc}^{\mathcal{P}}(c)| + |\mathcal{Z}| + 2 & \text{if } c \in \mathcal{P}_i \\ -2 \times |\text{up}^{\mathcal{N}}(c)| - |\text{inc}^{\mathcal{N}}(c)| - |\mathcal{Z}| - 2 & \text{if } c \in \mathcal{N}_i \\ 0 & \text{if } c \notin \mathcal{P}_i \cup \mathcal{N}_i \end{cases}$$

Given a profile \mathcal{S} of SP-structures, the **most popular candidates** are the ones maximising the sum of the individual Borda* score:

$$B^*(\mathcal{S}) = \operatorname{argmax}_{c \in \mathcal{X}} \sum_{i \in \mathcal{I}} s_i^*(c)$$

Example of using Borda*

Agent 1	Agent 2	Agent 3	
a			\mathcal{P}
$b \sim c$	b	a	
			\mathcal{Z}
	a		
			\mathcal{N}
	c		

Table: SP-structures extracted from the previous example.


The score $S^*(a) = 6 - 2 + 4$, $S^*(b) = 4 + 2 + 0$, $S^*(c) = 2 - 4 + 0$
 The **most preferred candidate** under the Borda* rule is a .

Challenge III:


How should similar methods be validated?

Two avenues for validation

Test predictive ability
against real-world events



Axiomatic properties
(Social Choice Theory)



Desirable properties of collective sentiment definitions

An important property when dealing with big data: Divide and conquer approaches are possible:

Consistency: *For all profiles S_1 and S_2 , if $F(S_1) \cap F(S_2) \neq \emptyset$ then $F(S_1 + S_2) = F(S_1) \cap F(S_2)$.*

If we rename the items the result should be the renaming of the initial result:

Neutrality: *For any profile S and permutation ρ of entities in \mathcal{X} , we have that $F(\rho(S)) = \rho(F(S))$.*

Considering one more individual in the computation of the collective sentiment results in a candidate that is higher in her rank (non-classical interpretation):

Voters Participation: *For all profiles $S = (S_1, \dots, S_n)$ and SP-structures S_{n+1} , we have that $F(S + S_{n+1})$ is ranked above (or equal) $F(S)$ by voter $n + 1$.*

Rank-participation, cancellation, weak Pareto, anonymity...

Desirable properties II

More information on an individual's opinion will lead to results that the individual ranks higher:

Rank Participation: *For all profiles $S = (S_1, \dots, S_n)$, individual i and SP-structure $S' \subset S_i$ we have that $F(S)$ is ranked above (or equal) $F(S_{-i} + S')$ by agent i .*

In symmetric cases, the disagreement is so extreme that it makes sense to declare all items as the most preferred ones in the collective opinion.

Cancellation. *If a profile S is symmetric then all entities are in the winning set, i.e., $F(S) = \mathcal{X}$.*

Agreement among all individuals should be reflected in the collective opinion.

Weak-Pareto: *If S is a profile in which all individuals rank a above b then $b \notin F(S)$.*

Theoretical results

The rule we propose satisfies a number of axiomatic properties:

Theorem

The Borda rule satisfies consistency, anonymity, neutrality, voters and rank participation as well as cancellation.*

It is also consistent with sentiment analysis and preference aggregation:


Theorem

If a profile S is purely preferential, then $B^(S) = \text{Borda}(S)$.*

If a profile S is purely sentimental, then $B^(S) = \text{Approval}(S)$.*


Experimental results I

Our analysis is based on the assumption that **sentiment analysis and preference aggregation** yields different results. How often does this happen?



Experimental results II

Incompleteness is perhaps the most important problem of text-extracted information. How does Borda* behave in **incomplete domains**?



Future work: detecting influence and strategic behavior

GET GUARANTEED TWITTER FOLLOWERS NOW
100% MONEY BACK GUARANTEE!

Twitter Followers	Price	Delivery Time
5,000	\$77	12 to 20 Days
2,000	\$32	7 to 9 Days
1,000 (Limited Time Offer)	\$17	4 to 7 Days
10,000	\$127	20 to 30 Days
25,000	\$247	30 to 40 Days

We don't require your Twitter password

Classical studied in Computational Social Choice:

- Strategy-proofness (aka resistance to individual manipulation): not interesting, absence of incentives for individual deviations...
- **Cloning and bribery** are more likely to occur: an applicative setting!

From Sentiment Analysis to Preference Aggregation

1. What preferences/opinions can be **extracted** from the individuals' text?
Our proposal: sentiment score and pairwise comparison (raw data)
2. How to best **represent** (compactly) individual preferences and sentiments?
Our proposal: SP-structures based on preorders
3. How to **aggregate** the individual information into a collective opinion?
Our proposal: generalise Borda and Approval with the Borda rule*
4. Is it possible to identify influencers and prevent **strategic behaviour**?
Example: creation of fake accounts (cloning)...
5. How should preference aggregation methods be **validated**?
Predictive accuracy and Social Choice Theory
6. How to deal with **big data** in sentiment and preference analysis?

Umberto Grandi, Andrea Loreggia, Francesca Rossi, Vijay Saraswat. From Sentiment Analysis to Preference Aggregation. *Proceedings of ISAIM-2014*, 2014.