From Sentiment Analysis to Preference Aggregation

Umberto Grandi

Department of Mathematics
University of Padova

22 November 2013

[Joint work with Andrea Loreggia, Francesca Rossi and Vijay Saraswat]

What is the collective sentiment about ...?

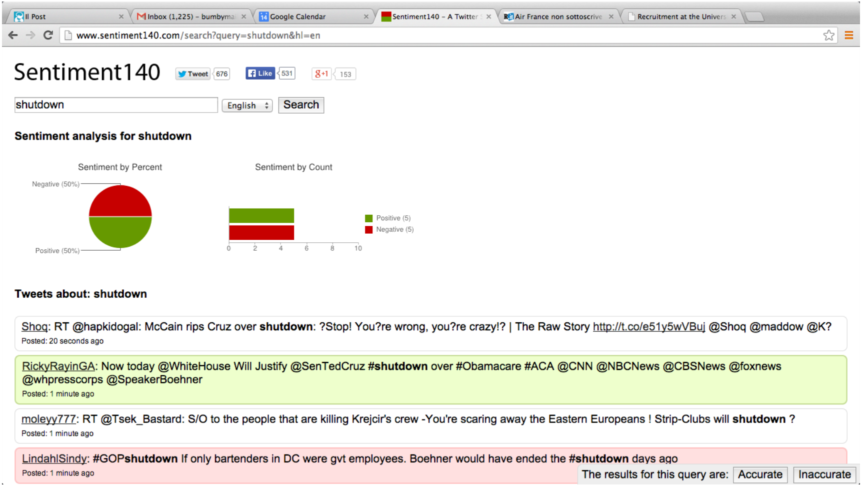

Aggregation of individual polarities (like/dislike)

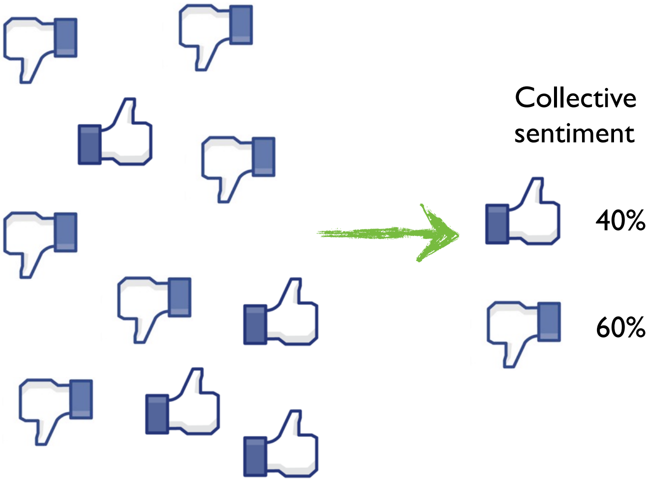

A problem: multiple alternatives

We extract the following sentiment about two candidates running for election:

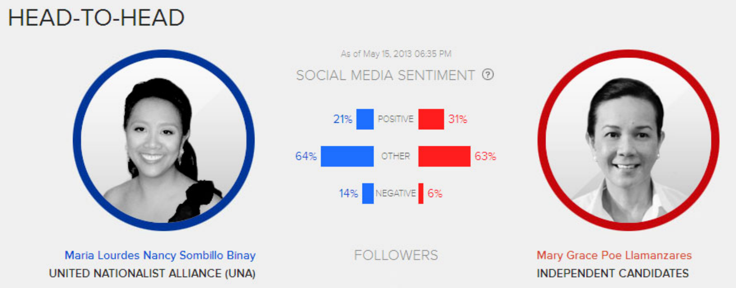

But what if preferences are the following:

$$
\begin{array}{lcc|cc}
\text{21 voters} & a & b & & \\
\text{10 voters} & & b & a & \\
\text{4 voters} & & & a & b
\end{array}
$$

Sentiment analysis predicts a different winner than majority aggregation!

Outline

1. Basic definitions: sentiment analysis and preference aggregation
2. Multiple alternatives:
   - Basic collective sentiment paradox
   - Counting paradoxes
3. Data structures from individual text:
   - pure sentiments (polarity)
   - pure preference (preorder)
   - sentiment and preferences (SP-structures)
4. Aggregation of SP-structures: Borda* rule
5. Open problems: six challenges in preference analysis

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 $\mathcal{T}$

Most common approach:
The percentage of positive expressions is the collective sentiment about x

Preference Aggregation

Ingredients:

- A set of candidates $\mathcal{X}$
- A set of individuals $\mathcal{I}$ expressing preferences as linear/weak orders on $\mathcal{X}$ or as sets of approved candidates in $\mathcal{X}$

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.

Part I:

The Problem

Basic collective sentiment paradox

Two candidates a and b are competing in an election:

- Sentiment analysis extracts 100% positive comments for b
- Majority rule elects a with a majority of 90 vs 10

Alternatives at the left of $|$ are positive, preferences from left to right:

90 voters	a	b	$\vert$	
10 voter		b	$\vert$	a

Majority rule winner: a
Collective sentiment predictor: b

Sentiment analysis can give the wrong result when predicting the majority rule!

More generally: sentiment analysis is problematic
in comparing more than two alternatives

Counting paradoxes: characterisation

A simple result to characterise collective sentiment paradoxes:

Proposition

A collective sentiment paradox with 2 *candidates occurs iff:*

$$N(a|b) \geqslant N(b|a)$$
$$N(ba|) + N(b|a) + N(|ba) > N(ab|) + N(a|b) + N(|ab)$$

or symmetrically for b *winning in SA.*

How to quantify the fraction of paradoxical profiles?

Awkward formula:

$$\sum_{l=\frac{n+1}{2}}^{n} \binom{n}{l} \sum_{t=0}^{n-l} \binom{l}{t} 2^{l-t} \sum_{m=t}^{n-l} \binom{n-l}{m} 2^{n-l-m}$$

Counting paradoxes: simulation

We performed experiments with 2 entities:

- sampling 10.000 profiles with the impartial culture assumption
- enumerating all paradoxical profiles up to $|\mathcal{I}| = 93$ (see figure below)

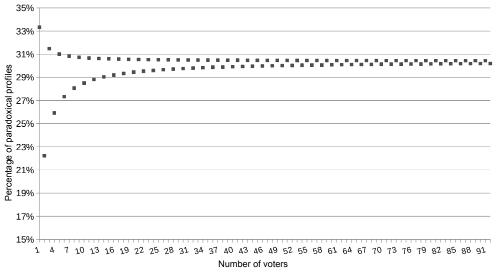

Figure: % of collective sentiment paradoxes

Sentiment analysis and preference aggregation differ in 30% of the profiles

Part II:
Data Structures

Preference Analysis

Mix the ingredients of sentiment analysis with those of preference aggregation:

- A set of entities/items/alternatives $\mathcal{X}$
- A corpus $\mathcal{T}_i$ of individual expressions for each i in a set $\mathcal{I}$ of individuals
- What is the most preferred entity?

Lesson learned from collective sentiment paradoxes:
Polarity extraction is not sufficient if we want to compare entities!

What data structure we can/want to extract from individual expressions?

- polarity/graded polarity/score
- only binary comparisons between alternatives
- a combination of sentiment and preference

Raw Data Extraction

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

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

Definition

The raw data extracted from individual expressions $\mathcal{T}_i$ 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* $\leqslant_i^P$ *and* $\leqslant_i^N$*, representing positive and negative comparative opinions*

Pang et. al., Thumbs up? Sentiment Classification Using Machine Learning Techniques, EMNLP-2002.

Ganapathibhotla and Liu, Mining Opinions in Comparative Sentences, COLING-2008.
Jindal and Liu, Mining Comparative Sentences and Relations, AAAI-2006.

Raw Data Extraction

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

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

Definition

The raw data extracted from individual expressions $\mathcal{T}_i$ 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* $\leqslant_i^P$ *and* $\leqslant_i^N$, *representing positive and negative comparative opinions*

Example

Three entities a, b *and* c, *and three individuals:*

- $\sigma_1(a) = 5, \sigma_1(b) = \sigma_1(c) = 4$ *and* $P_1 = N_1 = \emptyset$
- $\sigma_2(b) = 1$, $P_2 = \emptyset$, *and* $N_2 = \{a, c\}$ *with* $a \geqslant_2^N c$
- $\sigma_3(c) = 0$, $P_3 = \{a, b\}$ *with* $a \geqslant_3^P b$, *and* $N_3 = \emptyset$

Pure Sentiment Data

Definition

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

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

Example

Pure sentiment data only deals with polarities:

- $S_1(a) = S_1(b) = S_1(c) = +$
- $S_2(b) = +$ *and* $S_2(a) = S_2(c) = -$
- $S_3(a) = S_3(b) = +$ *and* $S_3(c) = 0$.

The most popular candidate using approval voting is b.

Pure Preference Data

Definition

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

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

Example

Pure preference data only deals with pairwise comparisons:

- $a \geqslant_1 b \sim_1 c$
- $b \geqslant_2 a \geqslant_2 c$
- $a \geqslant_3 b \geqslant_3 c$

The most preferred candidate using the Borda rule is a.

Sentiment Preference 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 $\leqslant^{\mathcal{P}}$ and $\leqslant^{\mathcal{N}}$*

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

Example

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

Table: SP-structures extracted from the previous example.

Part III:
Aggregation of SP-structures

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:

$$s^*(c) = \begin{cases} 2 \times |\,\mathrm{down}^{\mathcal{P}}(c)| + |\,\mathrm{inc}^{\mathcal{P}}(c)| + |\mathcal{Z}| + 1 & \textit{if } c \in \mathcal{P}_i \\ -2 \times |\,\mathrm{up}^{\mathcal{N}}(c)| - |\,\mathrm{inc}^{\mathcal{N}}(c)| - |\mathcal{Z}| - 1 & \textit{if } c \in \mathcal{N}_i \\ 0 & \textit{if } c \notin \mathcal{P}_i \cup \mathcal{N}_i \end{cases}$$

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

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

Example of using Borda*

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

Table: SP-structures extracted from the previous example.

The most preferred candidate under the Borda* rule is a.

What we know about Borda*

A profile is purely preferential if all comparisons are positive (negative) for all individuals. A profile is purely sentimental if only positive/neutral sentiment is expressed and no pairwise comparison.

Theorem

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

Axiomatic properties adapted from Social Choice Theory:

Theorem

The Borda rule satisfies consistency, faithfulness, neutrality and the cancellation property.*

Theorem

If $\boldsymbol{S}$ is a profile in which all individuals rank a above b then $b \notin B^*(\boldsymbol{S})$.

Part IV:
Open Problems

From Sentiment Analysis to Preference Aggregation

Six challenges to study the use of preference/voting tools in sentiment analysis:

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

From Sentiment Analysis to Preference Aggregation

Six challenges in the use of preference/voting tools in sentiment analysis:

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

Thank you for your attention!