European Parliament elections on Twitter

Analysis of Twitter feeds

Jan Lasek, Bartosz Ziembiński

6 June 2014

Outline

The goal of our project is *to investigate any possible relations between public support towards two Polish major political parties - Platforma Obywatelska (PO) and Prawo i Sprawiedliwość (PiS) - and attitude towards them expressed in Twitter feeds.*

Schedule of our presentation:

1. Methodology of Profile of Mood States (POMS)
2. Expanding POMS lexicon
3. Experiments - linking the tweets with the polls

Twitter mood predicts the stock market

There is a known paper by Johan Bollen et al. about predicting the stock market with the usage of Twitter feeds titled *Twitter mood predicts the stock market*.

Behavioral economics tells us that emotions can profoundly affect individual behavior and decision-making.

The paper tries to answer the following questions:

- Does this also apply to societies at large, i.e. can societies experience mood states that affect their collective decision making?
- By extension is the public mood correlated or even predictive of economic indicators?

Methodology of Profile of Mood States

Measuring the public mood is not a simple task. Large surveys of public mood over representative samples of the population are generally expensive and time-consuming to conduct.

To extract mood variations from Twitter feeds authors use a tool named Google-Profile of Mood States.

It is a computational version of known and well-vetted psychometric instrument, namely the Profile of Mood States.

Profile of Mood States

POMS is a psychological test formulated by McNair et al. (1971) which is used to measure the current mood of a person. It is a 65-item inventory that assesses six dimensions of the mood construct: anger, confusion, depression, fatigue, tension and vigour.

	Not at All	A Little	Moderately	Quite a Bit	Extremely
1. Worn out	0	1	2	3	4
2. Resentful	0	1	2	3	4
3. Annoyed	0	1	2	3	4
4. Fatigued	0	1	2	3	4
5. Bitter	0	1	2	3	4
6. Exhausted	0	1	2	3	4
7. Helpless	0	1	2	3	4
8. Cynical	0	1	2	3	4
9. Irritated	0	1	2	3	4
10. Drained	0	1	2	3	4
11. Overwhelmed	0	1	2	3	4
12. Enraged	0	1	2	3	4

Profile of Mood States (2)

It is used for example to test athletes' moods before contests. Lower scores indicate people with more stable mood profiles.

In his paper P. Terry provides POMS norms for an athletic sample ($n = 2086$) grouped by level of competition (international standard athletes, club level athletes and recreational athletes)

Group	Tension	Depression	Anger	Vigour	Fatigue	Confusion
International	5.66	4.38	6.24	18.51	5.37	4.00
Club	9.62	8.67	9.91	15.64	8.16	7.38
Recreational	6.00	3.11	3.60	17.78	6.37	4.84

What Bollen et al. did in their work was to create a computional version of the test. They did it by expanding the basic lexicon of 65 adjectives from POMS with similar words, which can be obtained by analysing word co-occurrences in big collections of text.

The expanded lexicon consisted of 964 associated terms which were collected by analyzing word co-occurrences in a collection of 2.5 billion 4- and 5-grams computed by Google in 2006 from approximately 1 trillion word tokens observed in publicly accessible webpages.

Computational Profile of Mood States

Then, in order to measure the public mood, they matched the terms used in each tweet against the expanded lexicon.

Each tweet term that matched a term was mapped back to its original POMS term and via the POMS scoring table to its respective POMS dimension.

All matches were then summed for each dimension in order to obtain the scores.

Assessing Polish Twitter

In our work we wanted to use the same approach to assess Polish Twitter's mood in the context of the elections to European Parliament.
We wanted to see if there are some correlations between mood dimensions and support for main Polish political parties.

In order to do that we needed to translate 65 POMS adjective to Polish.

Problems with Polish POMS (1)

Simple translation would not be sufficient to measure the mood of a Polish person in a original POMS test.

Nevertheless, as we were going to expand the lexicon with similar words, we did not have to consider this problem (we were going to find proper translations anyway).

Problems with Polish POMS (2)

Polish adjectives can occur in different grammatical persons:

- "przyjacielski",
- "przyjacielska" etc.

and cases:

- "przyjacielskiego",
- "przyjacielskim" etc.

Problems with Polish POMS (2)

The solution is to transform the word from a tweet to its basic grammatical form before comparison with a term from the lexicon.

To achieve this we use *Morfeusz*, a morphological analyser for Polish.

For efficiency we don't pass every word through *Morfeusz*, but we compare first two letters of the given word and the term from the lexicon to find out if there is a chance for the match.

The rest of the algorithm is the same as the one used by Bollen.

Remark: Additional dimensions

In our approach we are measuring more dimensions than in original POMS. We are expanding basic 6: *anger, confusion, depression, fatigue, tension, vigour* with 7th dimension *friendliness* which is also measured by POMS, but not used because of the fact that it was not tested.

We also measure two basic dimensions: *positive* and *negative*, using different lexicon of adjectives.

Extending POMS lexicon (1)

The list of POMS is relatively short so we shall enrich it in a way. We decided to use a very large and easily accessible collection of documents with an easy query interface - Web search engine (Bing).

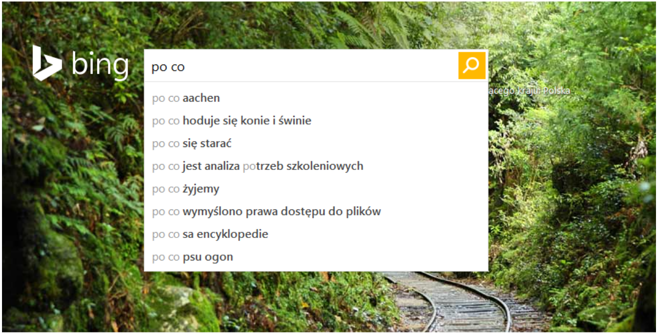

The solution is limited - Bing's API allows you to emit 5000 queries a month for free (each query gives you 50 results).

Extending POMS lexicon (2)

To generate cadidate words we issue queries of the type "*jest* $\langle adj \rangle$ *i*" and "*był* $\langle adj \rangle$ *i*", where $\langle adj \rangle$ denote a particular adjectives to which we want to find similar words. Here is a sample of more curious results, e.g., **"przyjacielski"**:

Extending POMS lexicon (2)

To generate cadidate words we issue queries of the type "*jest* $\langle adj \rangle$ *i*" and "*był* $\langle adj \rangle$ *i*", where $\langle adj \rangle$ denote a particular adjectives to which we want to find similar words. Here is a sample of more curious results, e.g., **"przyjacielski"**:

- "kontaktowy"

Extending POMS lexicon (2)

To generate cadidate words we issue queries of the type "*jest* $\langle adj \rangle$ *i*" and "*był* $\langle adj \rangle$ *i*", where $\langle adj \rangle$ denote a particular adjectives to which we want to find similar words. Here is a sample of more curious results, e.g., **"przyjacielski"**:

- "kontaktowy"
- "ufny"

Extending POMS lexicon (2)

To generate cadidate words we issue queries of the type "*jest* ⟨*adj*⟩ *i*" and "*był* ⟨*adj*⟩ *i*", where ⟨*adj*⟩ denote a particular adjectives to which we want to find similar words. Here is a sample of more curious results, e.g., **"przyjacielski"**:

- "kontaktowy"
- "ufny"
- "wierny"

Extending POMS lexicon (2)

Let's see one more example: **"wykończony"**:

Extending POMS lexicon (2)

Let's see one more example: **"wykończony"**:

- "gotowy"

Extending POMS lexicon (2)

Let's see one more example: **"wykończony"**:

- "gotowy"
- "ogrzewany"

Extending POMS lexicon (2)

Let's see one more example: **"wykończony"**:

- "gotowy"
- "ogrzewany"
- "umeblowany"

Extending POMS lexicon (3)

Further on the list we can place: "jest przyjacielski i wykształcony", "był nieszczęśliwy i traktowany", "rześki i przytomny", "taktowny i kompetentny", "zaniepokojony i smuuuuutny"...

Some postprocessing is necessary on the results but in general they are satisfactory. Expanded POMS lexicon consists of over 900 words.

The poll results can be found on the Web. However, over the period of campaign (March - May) they are issued irregularly. We need a way to account for this and also for their inherent noise.

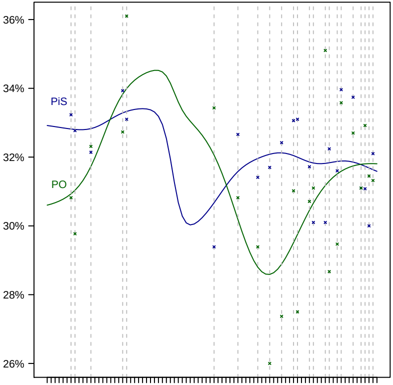

Figure: Public support curves approximated from poll results using Nadaraya-Watson estimator.

Another source of data for our study is Twitter. It should be noted that substantial amount of research have been done in analysing tweets. However, Polish Twitter have not gained much attention so far.

Certain things should be stressed when it comes to the analysis of tweets:

- There is some sample bias in tweets - most of them are produced by a group of young people (less than 25 years)
- The most popular channel on Twitter is @MTVPolska (highly much popluar than, e.g., @tvn24 or @gazeta_wyborcza)

Our data constitute for tweets harvested from 700 “political accounts”. Special acknowledgements to dr Aleksander Wawer for his help accessing the data.

Some sweets

We also looked at the stream of Polish tweets ($\approx 1\%$ sample) available via a dedicated API. Here are some of our favourites:

"twitterze zjedz snickersa bo zaczynasz znów gwiazdorzyc"

"Potop czyli lanie wody #matura2014"

"Dzisiaj Wesele, a w sierpniu Poprawiny. #matura2014"

"jem kolacje"

"@jmiziolek zobaczyłem prognozę pogody na niedzielę u @AGozdyra i od razu stawiam na PIS;)"

"Nie idę jutro do szkoły :)))"

In total from 1 March and 23 May we gathered 344710 tweets. This gives us an average of 4153 tweets per day. However, we have only 5287 tweets relevant for either PO or PiS in the considered period - this gives us an average of 31.5 tweets per day per party.

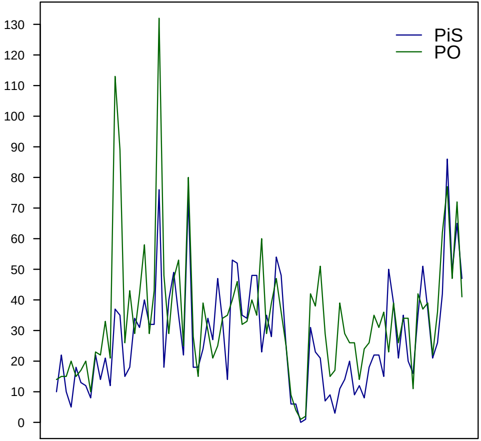

Experiments - method

To explore relations with different type of attitude toward parties we use Kendall's tau correlation coefficient. We correct for multiple hypothesis testing with Bonferroni's method and assume $\hat{\alpha} = 0.05/(10 \cdot 4) = 0.00125$ level of significance.

Moreover, we exclude 4 days from consideration as outliers (identified visually at the plot).

Results

Table: Kendall's correlation values for different dimensions and support for PO.

mood/lag	0	3	7	10
tension	-0.027	-0.008	0.027	0.048
depression	0.043	0.032	0.009	0.021
anger	0.112	0.092	0.047	0.032
vigour	0.005	0.029	0.055	0.047
fatigue	0.139	0.135	0.141	0.099
confusion	0.126	0.076	-0.008	-0.076
friendliness	-0.018	0.009	0.052	0.09
positive	0.142	0.122	0.06	-0.005
negative	0.213	0.17	0.093	0.021
total	**0.25**	0.205	0.101	-0.02

Correlation of the support with the total number of tweets on a given day is found significant (p-value of 0.0012).

Results

Table: Kendall's correlation values for different dimensions and support for PiS.

mood/lag	0	3	7	10
tension	-0.057	-0.04	0.035	0.098
depression	-0.129	-0.202	-0.203	-0.165
anger	-0.131	-0.098	-0.101	-0.037
vigour	-0.163	-0.175	-0.121	-0.089
fatigue	-0.12	-0.117	-0.015	-0.025
confusion	-0.027	0.011	-0.004	-0.004
friendliness	-0.114	-0.175	-0.179	-0.179
positive	-0.166	-0.134	-0.115	-0.128
negative	-0.153	-0.119	-0.026	0.032
total	-0.21	-0.147	-0.101	-0.096

No significant correlations found.

Summary and future work

- We repeated the experiment with the Pearson correlation coefficient and obtained the same result - the only significant correlation was the pair (support, total) in case of PO (correlation of 0.3722, p-value equal 0.0007)
- The more tweets about PO, the higher the support - it is not important how people speak about the party, it is only important that they do
- We need to double-check the result and hypothesise its origin
- A way to account for negations of POMS adjectives is needed
- We may gather more polls data (our calculation is based on 34 polls over 83 days)

Questions

?

Bibliography and further materials (1)

- J. Bollen, H. Mao, and X. Zeng. Twitter mood predicts the stock market. Journal of Computational Science, 2011
- D. McNair, M. Lorr, and L. Droppleman. Profile of mood states (POMS), 1989
- Terry, P. Normative Values for the Profile of Mood States for Use with Athletic Samples, (n. d.)
- M. Woliński: Morfeusz - A practical tool for the morphological analysis of polish. In: Intelligent information processing and web mining, Springer, 2006
- V. Hatzivassiloglou and K. R. McKeown. Predicting the Semantic Orientation of Adjectives. Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics, 1997

Bibliography and further materials (2)

- P. D. Turney. Thumbs Up or Thumbs Down? Semantic Orientation Applied to Unsupervised Classification of Reviews. Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, 2002
- B. O'Connor, R. Balasubramanyan, B. Routledge and N. Smith. From Tweets to Polls: Linking Text Sentiment to Public Opinion Time Series. International AAAI Conference on Weblogs and Social Media, 2010
- A. Wawer: Methods of Sentiment Detection in Polish, PhD Thesis
- http://twistedsifter.com/2014/05/nat-geo-traveler-photo-contest-2014-notable-entries/