Big Data Analysis of News and Social Media Content

M Vamsi
Student, Lingayas
University
vamsimadala02@gmail.com

B Sri Harsha
Student, Lingayas
University
harshacool2712@gmail.com

Priyanka Sharma
Professor,Lingayas
University
priyankasharmaiitd2@gmail.com

Abstract

The analysis of media content has been vital in social sciences, due to the key role that media plays in determining public opinion. This kind of analysis typically relies on the coding of the text being examined; a step that involves reading, annotating it, and that limits the sizes of the corpora that can be analyzed. The use of modern technologies from Artificial Intelligence allows researchers to computerize the process of applying different codes in the same text. Computational technologies also enable the automation of data collection, preparation, management and visualization. This provides opportunities for performing massive scale investigations, real time monitoring, and system-level modelling of the global media system. We describe how the analysis of content on Twitter can reveal mood changes of entire populations, how the political relations among US leaders can be extracted from large corpora, how we can determine what news people really want to read, how writing-style in articles change among different outlets etc. Most importantly, this paper aims to demonstrate some of the steps that can be automated, allowing researchers to access patterns that would be otherwise out of reach.

Introduction

The ready availability of masses of data and the means to exploit them is changing the way we execute science in many domains (Cristianini, 2010; Halevy et al., 2009). Molecular biology, astronomy and chemistry have already been transformed by the data revolution, undergoing an actual paradigm shift. In other domains, like social sciences (Lazer et al., 2009; Michel et al., 2011) and humanities (Moretti, 2011), data-driven approaches are just beginning to be deployed. This delay was caused by the density of the social interactions and the unavailability of digital data (Watts, 2007). Examples of works that deploy computational methods in social sciences include: the study of human interactions using mobile phones data (Onnela et al., 2007; González et al., 2008); the study of human interactions in online games and environments (Szell et al., 2010); or automating text annotation for political science research (Cardie et al., 2008).

In particular, the analysis of media content in the social sciences has been a very important concern, due to the important role that media play in reflecting and determining public opinion (Lewis, 2001). This includes the analysis of both traditional news outlets, such as newspapers and and modern social media, such as Twitter where content is generated by users. The analysis of news content is traditionally based on the coding of data by human coders. This is a labour-intensive process that restricts the possibility and potential of this kind of investigation. Computational methods from text mining and pattern recognition, originally developed for applications like e-commerce and information retrieval, can be deployed to automate many aspects of data collection, annotation, analysis and visualization.

Sentiment Analysis of Twitter Content

Since opinions are key influencers of human behavior, measuring the current public mood is an important and challenging task. The traditional approach would require questioning a large number of people about their feelings and opinions. Social media, such as Twitter or Facebook, can easily become a valuable source of information about the public due to the fact that people use them to express their feelings in public.

As demonstrated in the study by Lansdall-Welfare et al. (2012), it is feasible to capture the public mood by monitoring the stream of Twitter data. The dataset that was analyzed was comprised of 484 million tweets that were generated by more than 9.8 million users, between July 2009 and January 2012. The data were collected from the 54 largest cities in the UK. The research focuses on tracking the four moods which are "Fear", "Joy", "Anger" and "Sadness". For each and every mood, we track a long list of related words and we count the frequencies that these words appear in tweets. This process generates one timeline of the volume of connected tweets for each emotion. The further analysis of these timelines reveals that each of these emotions changes over time in a quite predictable manner. For an example, we found a periodic peak of joy around Christmas and a periodic peak of fear around Halloween. More surprisingly, we found that negative moods started to dominate the Twitter content after the announcement of enormous cuts in public spending on October 2010. Also there was a significant increase in anger in the weeks before the summer riots of August 2011.

Fig.1 we plot the mood levels for the period of study and we visualize them as a facial expression using the Grimace tool. In the project website we have an interactive demo that can be used to explore our data[2].

A detailed interpretation of the results of this study is a challenging task. A scholar needs to understand the causes of the measured statistical quantities and relate them with specific events that affect society. The interpretation of the results can not be automated by using computational methods and it is a task that needs to be undertaken by social scientists.
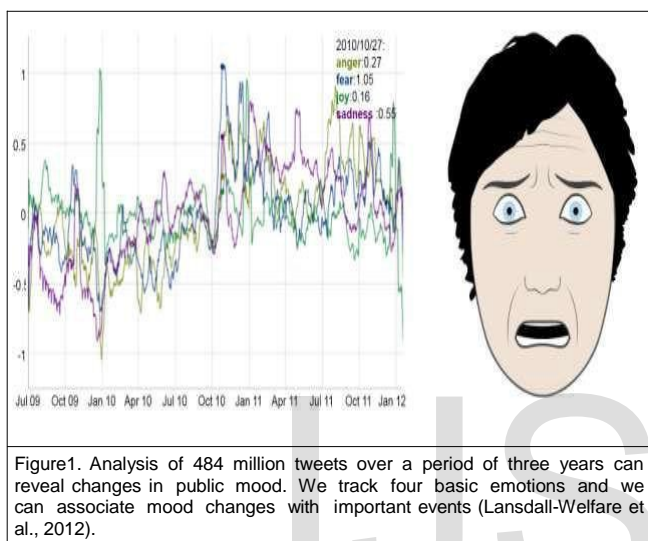

Figure1. Analysis of 484 million tweets over a period of three years can reveal changes in public mood. We track four basic emotions and we can associate mood changes with important events (Lansdall-Welfare et al., 2012).

Narrative Analysis of US Elections

Content analysis sometimes involves the identification of essential narrative information in a corpus. For example, determining the key actors, the key actions, and how actors interact. Quantitative Narrative Analysis (Franzosi, 1987; Earl et al., 2004) extracts Subject-Verb-Object (SVO) triplets and their study in order to analyze a corpus, while "distant reading" in humanities transforms the novels into networks as a first step of analysis (Moretti, 2011). The extraction of SVO triplets and turning them into a network is a process that can be automated and scaled up, enabling the analysis of big corpora.

In the work by Sudhahar et al. (2012), computational methods are used to perform narrative analysis of the US Elections 2012. The dataset we analyzed comprised of 125,254 articles that were automatically categorized as being associated to the US Elections using the machine learning Support Vector Machines (SVM) approach (Cristianini, 2000).The corpus text is pre-processed using GATE (Cunningham, 2002) tools and then we use Minipar parser (Lin, 1998) to extract Subject-Verb-Object triplets. We identified 31,476 actors from these triplets. A statistical analysis is used to identify the key actors and actions. Furthermore, we separate actions into two categories, endorsement and

opposition. Finally, the key actors and actions were used to conclude the network of the US Elections. This network is presented in Fig. 2 were the protagonists of the Elections, "Obama" and "Romney", are easily observed to dominate the networkExploring this network can give a useful insight on the topics that are most discussed during the elections, who the most dominant or least reported people are, how frequently the key actors are presented by media as endorsing or differing topics etc. Our results can be accessed online through an interactive website that is built for monitoring the elections[3]. The website is updated daily for the duration of the elections. A further analysis requires a social scientist to find and deeply understand the causes and relate them with raw data in order to interpret the patterns that surface.


Figure2. The key actors and their relations automatically extracted from a corpus of 125,254 articles associated to US Elections 2012 (Saatviga et al. 2012).

Comparison of News Outlets on the basis of the Topics, Writing Style and Gender Bias

Detection of differences or biases among news outlets is a trendy field of research in media studies. Nowadays, several aspects of content analysis can be automated: using machine learning techniques such as Support Vector Machines (Cristianini et al., 2000) we can evaluate text and detect general topics, such as Politics or Sports, that are present in a corpus; we can use natural language processing techniques to compute properties of the writing style of the text or detect mentions of people and their gender.

In the work by Flaounas et al. (2012), a dataset comprised of 2.5 million news articles collected from 498 news outlets over a period of 10 months is analyzed. For each article we identify the general topics that it covered, the people those are mentioned and their gender (Cunningham et al., 2002), as well as two fundamental writing style properties, namely the readability and linguistic subjectivity. Readability is a measure of how easy it is to read a text calculated by taking into account the length of sentences and the length of the words in syllables (Flesch, 1948). Linguistic subjectivity is a measure of the usage of words that carry sentiment, i.e. the more words with sentiment present in an article, the more linguistically subjective it is (Pang, 2008).

The computation of the aforesaid quantities allows different comparisons of sets of articles. For example, for the articles of each subject we calculated: a) the average readability - finding that articles about sports are easy to read while articles on politics are the hardest to read; b) Linguistic subjectivity - finding that articles about fashion and arts are the most biased while business articles were the most objective; c) the ratio of males over females – finding that articles about fashion and art mention the most women while articles about sports are male dominated.

The combination of large amounts of data in digital format with the application of machine learning and natural languages techniques can help social scientists answer a diverse range of research questions that could not be posed before.


Figure3. Newspapers compared based on their Readability and Linguistic Subjectivity. The illustration is based on the analysis of a dataset comprised of articles from US and UK newspapers (Flaounas et al. 2012).

The EU News Media Network

A recent technology that enables new avenues of research is machine translation. The technology is not mature enough to create human quality translations but nevertheless allows a scholar to access textual material written in most languages. Furthermore, machine translation can be combined with other computational methods in order to analyse texts written in non-English languages using computational tools developed for English. This reduces the costs of developing specialised tools for each language separately.

In the paper by Flaounas et al. (2010), it is shown how machine translation allowed the content analysis of a large corpus comprised of 1.3 million articles published by main news media in the 27 EU countries over a period of six months. Those articles were written in 22 different European languages. All non-English articles were machine-translated into English before further processing. In this research, we used the machine learning approach of clustering to identify sets of articles that refer to the same event. Then statistical approaches were used to infer the network of the EU news media. In that network, two outlets are connected if they tend to publish articles about the same events. In Fig. 4 we illustrate a sparse version of that network. It can be observed that news outlets, represented as nodes of the network, are organised in communities. An interesting and expected result is that outlets from the same country belong to the same community, i.e; they are interested in the same stories.

A deeper analysis of the structure of the network revealed some less expected results. For example, that the EU media system reflects the geographic, economic, and cultural relations between the EU countries. Also, we showed that the deviation of the content of an EU news outlet from the "average" EU media content is significantly correlated with the year of accession to the EU of the country that the outlet is based (correlation 48.94%, p- value<0.01). Interestingly enough, the editors of 215 different news outlets from 27 countries made their independent choices over what stories they will cover; but these choices shape the EU media sphere in a way that reflects relations between EU countries.

The detection of these subtle signals in a statistically rigorous way required the analysis of a massive dataset and the deployment of methodologies such as machine translation, data clustering and network analysis. To conclude, we showed that ideally, the language differences should not be an obstacle for a scholar to formulate research questions and analyze data.

Figure4. The communities of the network of EU news outlets. Nodes represent outlets and two outlets are connected if they tend to cover the

same stories. Outlets from the same country are coloured with the same colour. To infer this network we used more than 1.3 million articles written in 22 different languages (Flaounas et al., 2010)
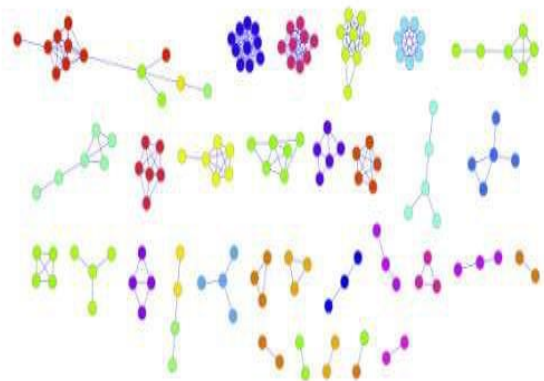

Figure 4.

Figure5. The word cloud presents in pink the words that are present in high-appeal articles published by "News.com.au" and in black the words that are present in low-appeal articles. The size of each word reflects the contribution of the word, either positive or negative, to appeal. The analyzed dataset contained 2.4 million of pairs of popular and non-popular articles (Hensinger et al. 2012b).

How News Outlets "Highlight" Stories

Some online news outlets highlight their "Most Popular" stories. This is a valuable source of information since it can be used to model the reading preferences of their audience. The modelling starts by forming a large number of pairs of articles. Each pair is comprised of one popular article and one non-popular article, i.e. an article that was published on the main page of the outlet, on the same day, but it didn't become popular. Then a machine learning algorithm, namely Ranking-SVM, can be used to create a statistical model that captures the words that tend to appear in popular articles (Hensinger et al., 2012a). The same model can be used to predict if any article will become popular or not.

In Fig. 5 we present as an example the word cloud of the model for "News.com.au" which shows the words that trigger the attention of the audience and the words that contribute less to the appeal of an article.

Similar approaches to the one we present can also be used to identify the preferences of people, by monitoring their online behavior - simply checking what they click on. Both approaches offer a coarse understanding of people's behavior. Of course, a deeper understanding would require a more traditional approach of questioning and interviewing people and reasoning about their answers.

Conclusions

The capability to gather vast amounts of digital data and the computational methodologies to analyze them are changing many aspects of science. The research questions that can be answered using data-driven approaches evolve and are multiplied based on the progress of pattern-analysis and data-mining technology. Machines provide the opportunities for extracting valuable information from big data and performing the repetitive and mundane tasks. Their capabilities to digest large amounts of data, far more than any researcher could analyze, are compensating for their inaccuracies and lack of reasoning. Of course, in many cases there is no substitute for human interpretation and judgment.

The boundary between tasks that can be solved by machines and those that can not is constantly changing. This is particularly true in the case of text analysis technologies. For example, just a decade ago machine translation could not be listed among the tools available to social scientists. The technologies of text-mining and text-analysis are evolving quickly, driven by strong commercial pressure, and as a side effect can lead to new tools for the social sciences. It is certainly worth keeping an eye on which new tools become available, as they could provide help in answering entirely new research questions. At the same time, entirely new types of data are constantly being released, ranging from news to social media content, from political proceedings to The content of millions of books (Michel et al., 2011). The opportunities for progress are real, when

we properly combine these two types of resources.
The work surveyed in this paper demonstrates technologies that can reliably detect topics, sentiment, persons, events and relations among actors, applied on millions of documents. Computational methods can also retrieve them, translate them and code them, all on a scale that cannot be matched by human analysts. We followed common practices in text analysis, by calibrating the tools we deployed on reference datasets that have been annotated by humans, so that we know the precision of our automatic annotation.

One important aspect of using big data in research has been mentioned only marginally in this survey: the need to visualize the data and the relations discovered in it, in order to help analysts make sense of the large scale datasets. We have demonstrated how political relations can be represented as networks, how mood can be represented as a face expression, but much more can be done in this direction.

While we have focused on the analysis of textual news content, and any other structure we derived from it (e.g. networks, maps and timelines), many researchers in the area of Computational Social Science (CSS) focus on different types of data: for example social networks, or search engine queries. In most of those cases, all data from CSS share the feature of being mostly unstructured, that is not naturally organized in rigorous structures, which is instead the norm in engineering. The exploitation of unstructured data is becoming a major area of concern in computer science and we can expect fast progress in the next few years.

As we move quickly towards a new way of doing science, we can see both opportunities and risks. It is important to keep both of them in mind, and to avoid unnecessary hype. There are many cases where forcing the culture of engineering on a domain such as the social sciences, but also the natural sciences, would be a mistake. We should keep in mind that these tools can only hope to complement – never to replace – the role of analysts, with their capability to understand context, to interpret and to incorporate their findings into a coherent narration. Data-driven technologies can be useful when they allow analysts to become aware of macroscopic trends that would be otherwise out of reach, but we should not delude ourselves that they have any chance of ever taking their place.

References

Peter Wlodarczak, Jeffrey Soar, Mustafa Ally "Big Data Analytics of Social Media"

Ralph Schroeder (2014) "Big Data and the brave new world of social media research"

Let data speak to data (2005), (Editorial), Nature 438, pp. 531.

Cardie, C, Wilkerson, J (2008) "Text annotation for political science research", Journal of Information Technology & Politics, 5(1), pp. 1–6.

Cristianini, N, Shawe-Taylor, J (2000) "An Introduction to Support Vector Machines and other Kernel-based learning methods", Cambridge University Press.

Cristianini, N (2010), "Are we there yet?", Neural Networks, Elsevier, 23(4), pp. 466-470.

Cunningham H, Maynard D, Bontcheva, K, and Tablan, V (2002) "GATE: A Framework and Graphical Development Environment for Robust NLP Tools and Applications", In Proc. of the 40th Anniversary Meeting of the Association for Computational Linguistics, pp. 168– 175.

Earl J, Martin A, McCarthy J, Soule S (2004) The use of newspaper data in the study of collective action. Annual Review of Sociology 30:65–80.

Halevy, A, Norvig, P and Pereira, F (2009), "The unreasonable effectiveness of data", Intelligent Systems, IEEE, vol. 24, pp. 8-12.

Flaounas, I, Turchi, M, Ali, O, Fyson, N, Bie, TD, Mosdell, N, Lewis, J & Cristianini, N (2010), 'The Structure of EU Mediasphere' PLoS ONE, vol 5 (12).

Flaounas, I, Ali, O, Turchi, M, Snowsill, T, Nicart, F, De Bie, T, Cristianini, N (2011) "NOAM: news outlets analysis and monitoring system", SIGMOD Conference, ACM, pp. 1275- 1278.

Flaounas, I, Ali, O, Lansdall-Welfare, T, De Bie, T, Mosdell, N, Lewis, J & Cristianini, N (2012), 'Research Methods in the Age of Digital Journalism: Massive-scale automated analysis of news-content: topics, style and gender', Digital Journalism, Routledge, vol 1, no. 1.

Flesch, R (1948) "A New Readability Yardstick", *Journal of Applied Psychology* 32, pp. 221–233.

Franzosi R (2010) Quantitative Narrative Analysis. Sage Publications Inc, Quantitative Applications in the Social Sciences, pp. 162:200.

González, M, Hidalgo, C, and Barabási, A (2008), Understanding individual human mobility patterns. Nature, 453(7196), pp. 779–782.

Hensinger, E, Flaounas, I, Cristianini, N (2012a) "What makes us click? Modelling and Predicting the Appeal of News Articles", Proc. of International Conference on Pattern Recognition Applications and Methods, pp. 41-50.

Hensinger, E, Flaounas, I, Cristianini, N (2012b) "The Appeal of Politics on Online Readers", In: Internet, Politics, Policy 2012:Big Data, Big Challenges?, Oxford,

Lansdall-Welfare, T, Lampos, V & Cristianini, N (2012), 'Effects of the Recession on Public Mood in the UK'. In Mining Social Network Dynamics (MSND) session on Social Media Applications in News and Entertainment at WWW '12., ACM, pp. 1221 – 1226.

Lazer, D, Pentland, A, Adamic, L, Aral, S, Barabasi, A-L, Brewer, D, Christakis, N, et al. (2009) "Computational Social Science", science 323, pp. 721-723.

Lewis, J (2001) "Constructing Public Opinion, New York: Columbia University Press.

Lewis, D, Yang, Y, Rose, T G, and Li, F (2004) "RCV1: A New Benchmark Collection for Text Categorization Research", *Journal of Machine Learning Research* 5, pp. 361–397.

Lin D, (1998), "Dependency-Based Evaluation of Minipar", Text, Speech and Language Technology 20 pp. 317–329.

Michel, J-B, Shen, Y K, Aiden, A P, Veres, A, Gray, M K., et al. (2011) "Quantitative

Analysis of Culture Using Millions of Digitized

Books", Science 331, pp. 176–182. Moretti, F

(2011) "Network theory, plot analysis". New Left

Review 68, pp. 80–102.

Onnela, J, Saramäki, J, Hyvönen, J, Szabó, G, Lazer, D, Kaski, K, Kertész, J, and Barabási, A (2007) "Structure and tie strengths in mobile communication networks", Proceedings of the National Academy of Sciences, 104(18), pp. 7332–7336.

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

Sandhaus, E (2008) "The New York Times Annotated Corpus", The New York Times Company, Research and Development.

Snowsill, T, Flaounas, I, De Bie, T, Cristianini, N (2010), 'Detecting events in a million New York Times articles'. In Proc. of Machine Learning and Knowledge Discovery in Databases, Springer, pp. 615-618.

Sudhahar, S, Lansdall-Welfare, T, Flaounas, I, Cristianini, N (2012), 'ElectionWatch: Detecting Patterns In News Coverage of US Elections'. In Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics (EACL), pp. 82-86.