

# Syndromic Classification of Twitter Messages

Nigel Collier and Son Doan

National Institute of Informatics, Tokyo, Japan

[collier@nii.ac.jp](mailto:collier@nii.ac.jp)

<http://born.nii.ac.jp/dizie>

**Abstract.** Recent studies have shown strong correlation between social networking data and national influenza rates. We expanded upon this success to develop an automated text mining system that classifies Twitter messages in real time into six syndromic categories based on key terms from a public health ontology. 10-fold cross validation tests were used to compare Naive Bayes (NB) and Support Vector Machine (SVM) models on a corpus of 7431 Twitter messages. SVM performed better than NB on 4 out of 6 syndromes. The best performing classifiers showed moderately strong F1 scores: respiratory = 86.2 (NB); gastrointestinal = 85.4 (SVM polynomial kernel degree 2); neurological = 88.6 (SVM polynomial kernel degree 1); rash = 86.0 (SVM polynomial kernel degree 1); constitutional = 89.3 (SVM polynomial kernel degree 1); hemorrhagic = 89.9 (NB). The resulting classifiers were deployed together with an EARS C2 aberration detection algorithm in an experimental online system.

**Keywords:** epidemic intelligence, social networking, machine learning, natural language processing.

## 1 Introduction

*Twitter* is a social networking service that allows users throughout the world to communicate their personal experiences, opinions and questions to each other using micro messages ('tweets'). The short message style reduces thought investment [1] and encourages a rapid 'on the go' style of messaging from mobile devices. Statistics show that Twitter had over 200 million users<sup>1</sup> in March 2011, representing a small but significant fraction of the international population across both age and gender<sup>2</sup> with a bias towards the urban population in their 20s and 30s. Our recent studies into novel health applications [2] have shown progress in identifying free-text signals from tweets that allow influenza-like illness (ILI) to be tracked in real time. Similar studies have shown strong correlation with national weekly influenza data from the Centers for Disease Control and Prevention and the United Kingdom's Health Protection Agency. Approaches like these hold out the hope that low cost sensor networks could be deployed as early

---

<sup>1</sup> <http://www.bbc.co.uk/news/business-12889048>

<sup>2</sup> <http://sustainablecitiescollective.com/urbantickurbantick/20462/twitter-usage-view-america>

warning systems to supplement more expensive traditional approaches. Web-based sensor networks might prove to be particularly effective for diseases that have a narrow window for effective intervention such as pandemic influenza.

Despite such progress, studies into deriving linguistic signals that correspond to other major syndromes have been lacking. Unlike ILI, publicly available gold standard data for other classes of conditions such as gastrointestinal or neurological illnesses are not so readily available. Nevertheless, the previous studies suggest that a more comprehensive early warning system based on the same principles and approaches should prove effective. Within the context of the DIZIE project, the contribution of this paper is (a) to present our data classification and collection approaches for building syndromic classifiers; (b) to evaluate machine learning approaches for predicting the classes of unseen Twitter messages; and (c) to show how we deployed the classifiers for detecting disease activity. A further goal of our work is to test the effectiveness of outbreak detection through geo-temporal aberration detection on aggregations of the classified messages. This work is now ongoing and will be reported elsewhere in a separate study.

## 1.1 Automated Web-Sensing

In this section we make a brief survey of recent health surveillance systems that use the Web as a sensor source to detect infectious disease outbreaks. Web reports from news media, blogs, microblogs, discussion forums, digital radio, user search queries etc. are considered useful because of their wide availability, low cost and real time nature. Although we will focus on infectious disease detection it is worth noting that similar approaches can be applied to other public health hazards such as earthquakes and typhoons [3,4].

Current systems fall into two distinct categories: (a) event-based systems that look for direct reports of interest in the news media (see [5] for a review), and (b) systems that exploit the human sensor network in sites like Twitter, Jaiku and Prownce by sampling reports of symptoms/GP visits/drug usage etc. from the population at risk [6,7,8]. Early alerts from such systems are typically used by public health analysts to initiate a risk analysis process involving many other sources such as human networks of expertise.

Work on the analysis of tweets, whilst still a relatively novel information source, is related to a tradition of syndromic surveillance based on analysis of triage chief complaint (TCC) reports, i.e. the initial triage report outlining the reasons for the patient visit to a hospital emergency room. Like tweets they report the patient's symptoms, are usually very brief, often just a few keywords and can be heavily abbreviated. Major technical challenges though do exist: unlike TCC reports tweets contain a very high degree of noise (e.g. spam, opinion, re-tweeting etc.) as well as slang (e.g. *itcy* for *itchy*) and emoticons which makes them particularly challenging. Social media is inherently an informal medium of communication and lacks a standard vocabulary although Twitter users do make use of an evolving semantic tag set. Both TCC and tweets often consist of short telegraphic statements or ungrammatical sentences which are difficult for uncustomised syntactic parsers to handle.

In the area of TCC reports we note work done by the RODS project [9] that developed automatic techniques for classifying reports into a list of syndromic categories based on natural language features. The chief complaint categories used in RODS were respiratory, gastrointestinal, botulinic, constitutional, neurologic, rash, hemorrhagic and none. Further processes took aggregated data and issued alerts using time series aberration detection algorithms. The DIZIE project which we report here takes a broadly similar approach but applies it to user generated content in the form of Twitter messages.

## 2 Method

DIZIE currently consists of the following components: (1) a list of latitudes and longitudes for target world cities based on Twitter usage; (2) a lexicon of syndromic keywords used as an initial filter, (3) a supervised machine learning model that converts tweets to a word vector representation and then classifies them according to six syndromes, (4) a post-processing list of stop words and phrases that blocks undesired contexts, (5) a MySQL database holding historic counts of positive messages by time and city location, used to calculate alerting baselines, (6) an aberration detection algorithm, and (7) a graphical user interface for displaying alerts and supporting evidence.

After an initial survey of high frequency Twitter sources by city location we selected 40 world cities as candidates for our surveillance system. Sampling in the runtime system is done using the Twitter API by searching for tweets originating within a 30km radius of a city's latitude and longitude, i.e. a typical commuting/shopping distance from the city centre. The sampling rate is once every hour although this can be shortened when the system is in full operation. In this initial study we focussed only on English language tweets and how to classify them into 6 syndromic categories which we describe below.

Key assumptions in our approach are that: (a) each user is considered to be a sensor in the environment and as such no sensor should have the capacity to over report. We controlled over reporting by simply restricting the maximum number of messages per day to be 5 per user; (b) each user reports on personal observations about themselves or those directly known to them. To control (a) and (b) and prevent over-reporting we had to build in filtering controls to mitigate the effects of information diffusion through re-reporting, particularly for public personalities and mass events. Re-tweets, i.e. repeated messages, and tweets involving external links were automatically removed.

### 2.1 Schema Development

A syndrome is a collection of symptoms (both specific and non-specific) agreed by the medical community that are indicative of a class of diseases. We chose six syndrome classes as the targets of our classifier: constitutional, respiratory, gastrointestinal, hemorrhagic, rash (i.e. dermatological) and neurological. These were based on an openly available public health ontology developed as part

of the BioCaster project [10] by a team of experts in computational linguists, public health, anthropology and genetics. Syndromes within the ontology were based on RODS syndrome definitions and are linked to symptom terms - both technical and laymen's terms - through typed relations. We use these symptoms (syndromic keywords) as the basis for searching Twitter and expanded them using held out Twitter data.

## 2.2 Twitter Data

After defining our syndromes we examined a sample of tweets and wrote guidelines outlining positive and negative case definitions. These guidelines were then used by three student annotators to classify a sample of 2000 tweets per syndrome into positive or negative for each of the syndrome classes. Data for training was collected by automatically searching Twitter using the syndromic keywords over the period 9th to 24th July 2010. No city filtering was applied when we collected the training data. Typical positive example messages are: "Woke up with a stomach ache!", "Every bone in my body hurts", and "Fever, back pain, headache... ugh!". Examples of negative messages are: "I'm exercising till I feel dizzy", "Cabin fever is severe right now", "Utterly exhausted after days of housework". Such negative examples include a variety of polysemous symptom words such as *fever* in its senses of raised temperature and excitement and *headache* in its senses of a pain in the head or an inconvenience. The negative examples also include cases where the context indicates that the cause of the symptom is unlikely to be an infection, e.g. headache caused by working or exercising. The training corpus is characterised using the top 7 terms calculated by mutual association score in Table 1. This includes several spurious associations such as 'rt' standing for 'repeat tweet', 'botox' which is discussed extensively as a treatment for several symptoms and 'charice' who is a new pop idol.

The final corpus was constructed from messages where there was total agreement between all three annotators. This data set was used to develop and evaluate supervised learning classifiers in cross-fold validation experiments. A summary of the data set is shown in Table 2. Inter-annotator agreement scores between the three annotators are given as Kappa showing agreement between the two highest agreeing annotators. Kappa indicates strong agreement on most syndromic classes with the notable exception of gastrointestinal and neurological.

## 2.3 Classifier Models

DIZIE employs a two stage filtering process. Since Twitter many topics unrelated to disease outbreaks, DIZIE firstly requests Twitter to send it messages that correspond to a set of core syndromic keywords, i.e. the same sampling strategy used to collect training/testing data. These keywords are defined in the BioCaster public health ontology [10]. In the second stage messages which are putatively on topic are filtered more rigorously using a machine learning approach. This stage of filtering aims to identify messages containing ambiguous

**Table 1.** Top 7 terms by syndrome calculated by mutual information score. \* indicates a spurious association.

Resp	Gastro	Const	Hemor	Rash	Neuro
throat	stomach	botox*	pain	road	headache
sore	ache	body	hemorrhage	heat	coma
cough	gib	charice*	muscle	arm	worst
flu	feel	jaw	tired	tired	gave
nose	rt*	hurts	pray	rash	giving
rt*	bad	stomach	brain	itcy	vertigo
cold	worst	sweating	guiliechelon*	face	pulpo*

**Table 2.** Structure of the annotated syndrome corpus of Twitter messages

Syndrome	Positives (P)	Negatives (N)	P/N	Kappa
Respiratory	627	738	0.85	0.67%
Gastrointestinal	489	676	0.72	0.49%
Neurological	549	434	1.26	0.42%
Rash	914	592	1.54	0.86%
Hemorrhagic	320	711	0.45	0.92%
Constitutional	1043	338	3.09	0.78%

words whose senses are not relevant to infectious diseases and messages where the cause of the symptoms are not likely to be infectious diseases. About 70% of messages are removed at this second stage.

To aid in model selection our experiments used two widely known machine learning models to classify Twitter messages into a fixed set of syndromic classes: Naive Bayes (NB) and support vector machines (SVM) [11] using a variety of kernel functions. Both models were trained with binary feature vectors representing a dictionary index of words in the training corpus. i.e. a feature for the test message was marked 1 if a word was present in the test message which had been seen previously in the training corpus otherwise 0. No normalisation of the surface words was done, e.g. using stemming, because of the high out of vocabulary rate with tools trained on general language texts.

Despite the implausibility of its strong statistical independence assumption between words, NB tends to perform strongly. The choice to explore keywords as features rather than more sophisticated parsing and conceptual analysis such as MPLUS [12] was taken from a desire to evaluate less expensive approaches before resorting to time consuming knowledge engineering.

The NB classifier exploits an estimation of the Bayes Rule:

$$P(c_k|d) = \frac{P(c_k) \times \prod_{i=1}^m P(f_i|c_k)^{f_i(d)}}{P(d)} \quad (1)$$

where the objective is to assign a given feature vector for a document  $d$  consisting of  $m$  features to the highest probability class  $c_k$ .  $f_i(d)$  denotes the frequency count of feature  $i$  in document  $d$ . Typically the denominator  $P(d)$  is not computed explicitly as it remains constant for all  $c_k$ . In order to compute the highest

value numerator NB makes an assumption that features are conditionally independent given the set of classes. Right hand side values of the equation are estimates based on counts observed in the training corpus of classified Twitter messages. We used the freely available Rainbow toolkit<sup>3</sup> from CMU as the software package.

SVMs have been widely used in text classification achieving state of the art predictive accuracy. The major distinction between the two approaches are that whereas NB is a generative classifier which forms a statistical model of each class, SVM is a large-margin binary classifier. SVM operates as a two stage process. Firstly the feature vectors are projected into a high dimensional space using a kernel function. The second stage finds a maximum margin hyperplane within this space that separates the positive from the negative instances of the syndromic class. In practice it is not necessary to perfectly classify all instances with the level of tolerance for misclassification being controlled by the C parameter in the model. A series of binary classifiers were constructed (one for each syndrome) using the SVM<sup>Light</sup> software package<sup>4</sup>. We explored polynomial degree 1, 2, 3 and radial basis function kernels.

## 2.4 Temporal Model

In order to detect unexpected rises in the stream of messages for each syndrome we implemented a widely used change point detection algorithm called the Early Aberration and Reporting System (EARS) C2 [13]. C2 reports an alert when its test value  $S_t$  exceeds a number  $k$  of standard deviations above a historic mean:

$$S_t = \max(0, (C_t - (\mu_t + k\sigma_t))/\sigma_t) \quad (2)$$

where  $C_t$  is the count of classified tweets for the day,  $\mu_t$  and  $\sigma_t$  are the mean and standard deviation of the counts during the history period, set as the previous two weeks.  $k$  controls the number of standard deviations above the mean where an alert is triggered, set to 1 in our system. The output of C2 is a numeric score indicating the degree of abnormality but this by itself is not so meaningful to ordinary users. We constructed 5 banding groups for the score and showed this in the graphical user interface.

## 3 Results

### 3.1 Classifying Twitter Messages

Results for 10-fold cross validation experiments on the classification models are shown in Table 3. Overall the SVM with polynomial degree 1 kernel outperformed all other kernels with other kernels generally offering better precision at a higher cost to recall. Precision (Positive predictive) values ranged from 82.0 to

<sup>3</sup> <http://www.cs.cmu.edu/~mccallum/bow/rainbow/>

<sup>4</sup> <http://svmlight.joachims.org/>

93.8 for SVM (polynomial degree 1) and from 83.3 to 99.0 for NB. Recall (sensitivity) values ranged from 58.3 to 96.2 for SVM (polynomial degree 1) and from 74.7 to 90.3 for NB. SVM tended to offer a reduced level of precision but better recall. In the case of one syndrome (Hemorrhagic) we noticed an unusually low level of recall for SVM but not for NB.

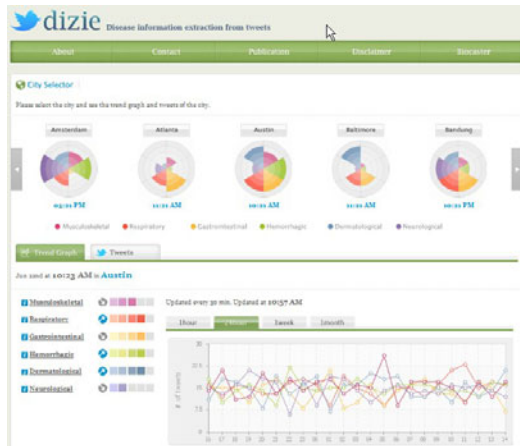
SVM's performance seemed moderately correlated to the positive/negative ratio in the training corpus and also showed weakness for the two classes (Hemorrhagic and Gastrointestinal) with the smallest positive counts. Naive Bayes performed robustly across classes with no obvious correlation either to positive/negative ratio or the volume of training data. Low performance was seen in both models for the gastrointestinal syndrome. This was probably due to the low number of training examples resulting from the low inter-annotator agreement on this class and the requirement for complete agreement between all three annotators.

**Table 3.** Evaluation of automated syndrome classification using naive Bayes and Support Vector Machine models on 10-fold cross validation. P - Precision, R - Recall, F1 - F1 score. <sup>1</sup> SVM using a linear kernel, <sup>2</sup> SVM using a polynomial kernel degree 2, <sup>3</sup> SVM using a polynomial kernel degree 3, <sup>R</sup> SVM using a radial basis function kernel.


	Naive Bayes			SVM <sup>1</sup>			SVM <sup>2</sup>			SVM <sup>3</sup>			SVM <sup>R</sup>			
	Synd.	P	R	F1	P	R	F1	P	R	F1	P	R	F1	P	R	F1
Resp.	90.3	82.4	86.2	85.4	82.5	83.8	83.0	71.0	76.5	86.4	61.3	71.7	66.7	3.2	6.2	
Gast.	83.3	75.5	79.2	85.9	78.4	81.8	92.7	79.2	85.4	91.4	66.7	77.1	73.1	39.6	51.3	
Neur.	98.2	74.7	84.8	83.2	95.0	88.6	77.9	98.2	86.9	62.4	98.2	76.3	90.0	63.0	74.1	
Rash	94.5	76.1	84.3	82.0	90.6	86.0	76.9	91.2	83.4	67.7	94.5	78.9	60.7	100.0	75.5	
Hem.	89.4	90.3	89.9	93.8	58.3	71.7	100.0	50.0	66.7	100.0	50	66.7	87.5	43.8	58.3	
Con.	99.0	79.8	88.4	83.6	96.2	89.3	83.6	93.3	88.2	78.6	99.0	87.7	76.5	100	86.7	

### 3.2 Technology Dissemination

An experimental service for syndromic surveillance called DIZIE has been implemented based on the best of our classifier models and we are now observing its performance. The service is freely available from an online portal at <http://born.nii.ac.jp/dizie>. As shown in Figure 3.2 the graphical user interface (GUI) for DIZIE shows a series of radial charts for each major world city with each band of the chart indicating the current level of alert for one of the six syndromes. Alerting level scores are calculated using the Temporal Model presented above. Each band is colour coded for easy recognition. Alerting levels are calculated on the classified twitter messages using the EARS C2 algorithm described above. Data selection is by city and time with drill down to a selection of user messages that contributed to the current level. Trend bars show the level of alert and whether the trend is upwards, downwards or sideways. Charting is also provided over an hourly, daily, weekly and monthly period. The number of positively classified messages by city is indicated in Figure 3.2 for a selection of cities.



**Fig. 1.** Radial graphs showing syndromic alert levels for major world cities. Colour coding on the radial segments indicates the alerting degree automatically assigned to a syndrome in a city based on the previous hour's Twitter counts and the previous 2 weeks as a baseline. The page is updated every hour. Clicking on the graph for a city displays the frequency graph and also the matching tweets for the current hour.



**Fig. 2.** Number of Tweets by a sample of major world cities classified by DIZIE during the period 2nd March 2011 to 31st August 2011

Navigation links are provided to and from BioCaster, a news event alerting system, and we expect in the future to integrate the two systems more closely to promote greater situation awareness across media sources. Access to the GUI is via regular Web browser or mobile device with the page adjusting automatically to fit smaller screens.



## 4 Conclusion

Twitter offers unique challenges and opportunities for syndromic surveillance. Approaches based on machine learning need to be able (a) to handle biased data, and (b) to adjust to the rapidly changing vocabulary to prevent a flood of false positives when new topics trend. Future work will compare keyword classifiers against more conceptual approaches such as [12] and also compare the performance characteristics of change point detection algorithms.

Based on the experiments reported here we have built an experimental application called DIZIE that samples Twitter messages originating in major world cities and automatically classifies them according to syndromes. Access to the system is openly available. Based on the outcome of our follow up study we intend to integrate DIZIE's output with our event-based surveillance system BioCaster which is currently used by the international public health community.

**Acknowledgements.** This work was in part supported by grant in aid support from the National Institute of Informatics' Grand Challenge Project (PI:NC). We are grateful to Reiko Matsuda Goodwin for commenting on the user interface in the early stages of this study and helping in data collection for the final system.

## References

1. Java, A., Song, X., Finin, T., Tseng, B.: Why we twitter: Understanding microblogging usage and communities. In: Proc. 9th WebKDD and 1st SNA-KDD Workshop on Web Mining and Social Network Analysis, August 12. ACM (2007)
2. Collier, N., Nguyen, S.T., Nguyen, M.T.N.: OMG U got flu? analysis of shared health messages for bio-surveillance. *Biomedical Semantics* 2(suppl. 5), S10 (2011)
3. Earle, P.: Earthquake twitter. *Nature Geoscience* 3(4), 221–222 (2010), doi:10.1038/ngeo832
4. Sakaki, T., Okazaki, M., Matsuo, Y.: Earthquake shakes twitter users: real-time event detection by social sensors. In: Proc. of the 19th International World Wide Web Conference, Raleigh, NC, USA, pp. 851–860 (2010)
5. Hartley, D., Nelson, N., Walters, R., Arthur, R., Yangarber, R., Madoff, L., Linge, J., Mawudeku, A., Collier, N., Brownstein, J., Thinus, G., Lightfoot, N.: The landscape of international biosurveillance. *Emerging Health Threats J.* 3(e3) (January 2010), doi:10.1093/bioinformatics/btn534
6. Szomszor, M., Kostkova, P., De Quincey, E.: swineflu : Twitter predicts swine flu outbreak in 2009 (December 2009)
7. Lampos, V., De Bie, T., Cristianini, N.: Flu Detector - Tracking Epidemics on Twitter. In: Balcázar, J.L., Bonchi, F., Gionis, A., Sebag, M. (eds.) ECML PKDD 2010. LNCS, vol. 6323, pp. 599–602. Springer, Heidelberg (2010)
8. Signorini, A., Segre, A.M., Polgreen, P.M.: The use of twitter to track levels of disease activity and public concern in the U.S. during the influenza a h1n1 pandemic. *PLoS One* 6(5), e19467 (2011)
9. Wagner, M.M., Espino, J., Tsui, F.C., Gesteland, P., Chapman, W., Ivanov, W., Moore, A., Wong, W., Dowling, J., Hutman, J.: Syndrome and outbreak detection using chief-complaint data - experience of the real-time outbreak and disease surveillance project. *Morbidity and Mortality Weekly Report (MMWR)* 53(suppl.), 28–31 (2004)

10. Collier, N., Doan, S., Kawazoe, A., Matsuda Goodwin, R., Conway, M., Tateno, Y., Ngo, Q., Dien, D., Kawtrakul, A., Takeuchi, K., Shigematsu, M., Taniguchi, K.: BioCaster: detecting public health rumors with a web-based text mining system. *Bioinformatics* 24(24), 2940–2941 (2008), doi:10.1093/bioinformatics/btn534
11. Joachims, T.: Text categorization with support vector machines: Learning with many relevant features. In: Nédellec, C., Rouveirol, C. (eds.) *ECML 1998. LNCS*, vol. 1398, pp. 137–142. Springer, Heidelberg (1998)
12. Christensen, L.M., Haug, P.J., Fiszmann, M.: Mplus: A probabilistic medical language understanding model. In: *Proceedings of the Workshop on Natural Language Processing in the Biomedical Domain, Philadelphia, USA (July 2002)*
13. Hutwagner, L., Thompson, W., Seeman, M.G., Treadwell, T.: The bioterrorism preparedness and response early aberration reporting system (EARS). *J. Urban Health* 80(2), i89–i96 (2003)