

Capturing Twitter Streams for Opinion Mining on Airport Noise

Dimitris Kotzinos

Joint work with: Iheb Meddeb, Asma Gharbi, Catherine Lavandier

ETIS / ENSEA - Université de Cergy-Pontoise - CNRS
UMR 8051 F-95000 Cergy-Pontoise Cedex, France



Introduction and Problematic

Traditional Surveys methods : Phone calls, interviews, web surveys → Hard work + Time consuming



- Share news
- Participate in discussions
- Express emotions and opinions

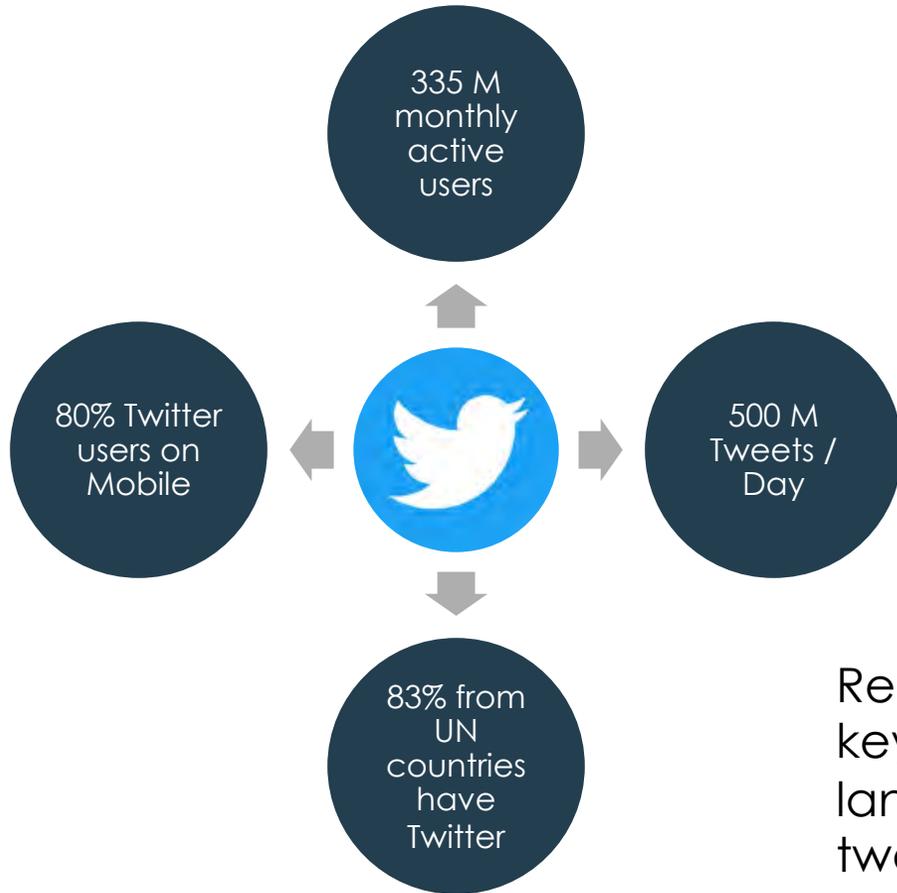
Social media and microblogging expansion → new source of information → Big Data

❑ Is it possible to offer an alternate survey method using this Data?

Problems emerging:

- How to deal with the big amount of data to extract only the relevant ones to the topic of interest?
- Which methods are used to classify data when the survey is based on opinions?
- Is there any bias in the data? How can we cope?

Twitter



Twitter properties:

- Short messages (280 maximum characters)
- Introducing hashtags (#) and mentioning users (@)
- Share tweets of others (retweet)
- Follow users / user followers

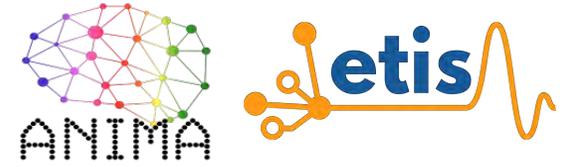
Twitter API:

- Integrate tweets into apps websites
- Advertise on twitter
- Access publicly available tweets (real-time stream or search)

Retrieve tweets by keywords, locations, languages, users, etc. → tweets presented as JSON objects

```
{
  "created_at" : "Thu May 10 15:24:15 +0000 2018" ,
  "id_str" : "850006245121695744" ,
  "text" : "Here is the Tweet message." ,
  "user" : {
  } ,
  "place" : {
  } ,
  "entities" : {
  } ,
  "extended_entities" : {
  }
}
```

Introduction and Problematic



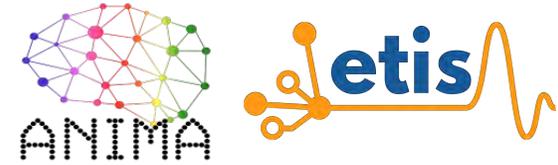
Pre-COVID19

- Working in real time, capturing and classifying tweets as they appear
- Work more interesting but more error-prone
- Data in the numbers of some hundred (relevant) tweets per day

During/After-COVID19

- Discussions mainly on COVID; lack of tweets about noise
- Access to millions of tweets (already collected more than 100 million tweets); starting from relative community accounts
- Richer dataset

Introduction and Problematic



Project goal: understand people's feeling against noise/quality of life around airports

Use case: the Heathrow Airport and the area around it

Problem 1: Relevance of tweets

- capture relevant twitter messages (tweets) about airport noise
Challenge: Filter out trivial and irrelevant messages from the stream

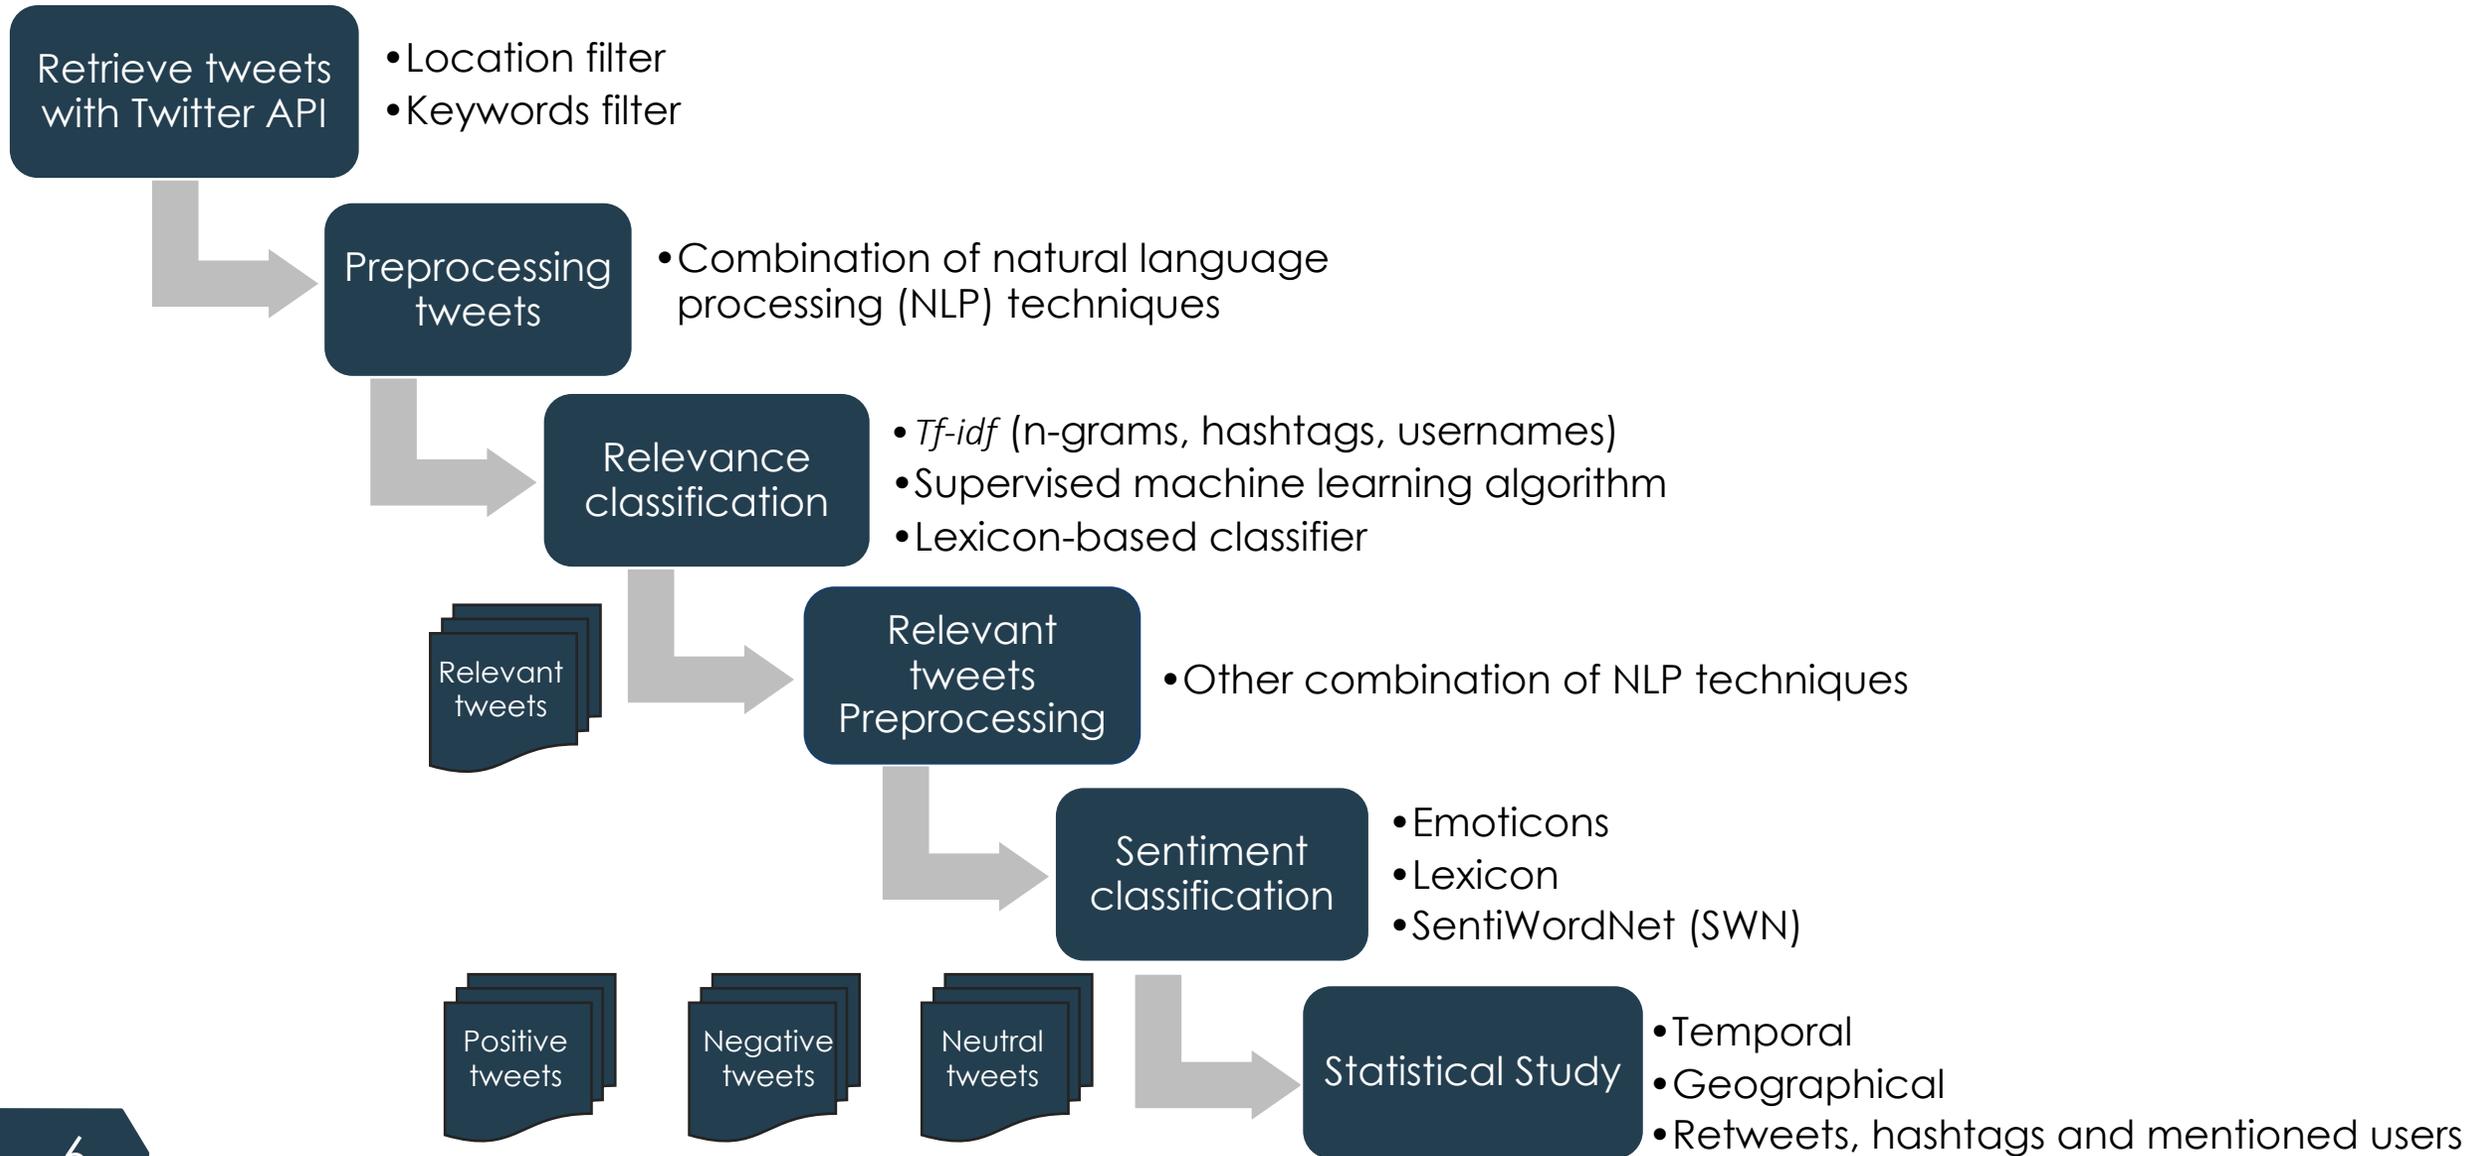
Problem 2: Sentiment Analysis

- Extract the sentiment of the user and classify tweets based on opinions (positive, negative or neutral)
Challenge: deal with Abbreviations, misspelling, incomplete sentences, negation, contrast, punctuations, irony, sarcasm
Challenge: Sentiment classification close to human judgements

Problem 3: Statistical analysis of the sentiments (post processing)

- Use classified tweets to:
 - pinpoint the areas mostly affected by the noise.
 - identify temporal bursts of tweets during the day and the week.
 - Trendy (hot) topics of discussion related to airport noise

Workflow: Main steps

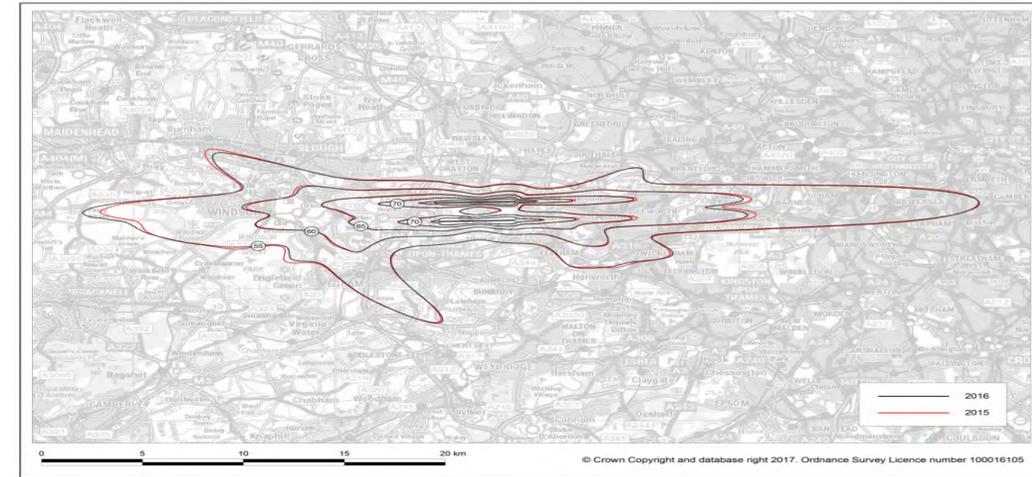


Workflow: Gathering tweets

□ Retrieve tweets with location filter

- Heathrow airport L_{den} noise contours [3] → set the minimum size of the area
- Bounding box of 167 km wide and 73 km long

 Gets only tweets with enabled location → small proportion

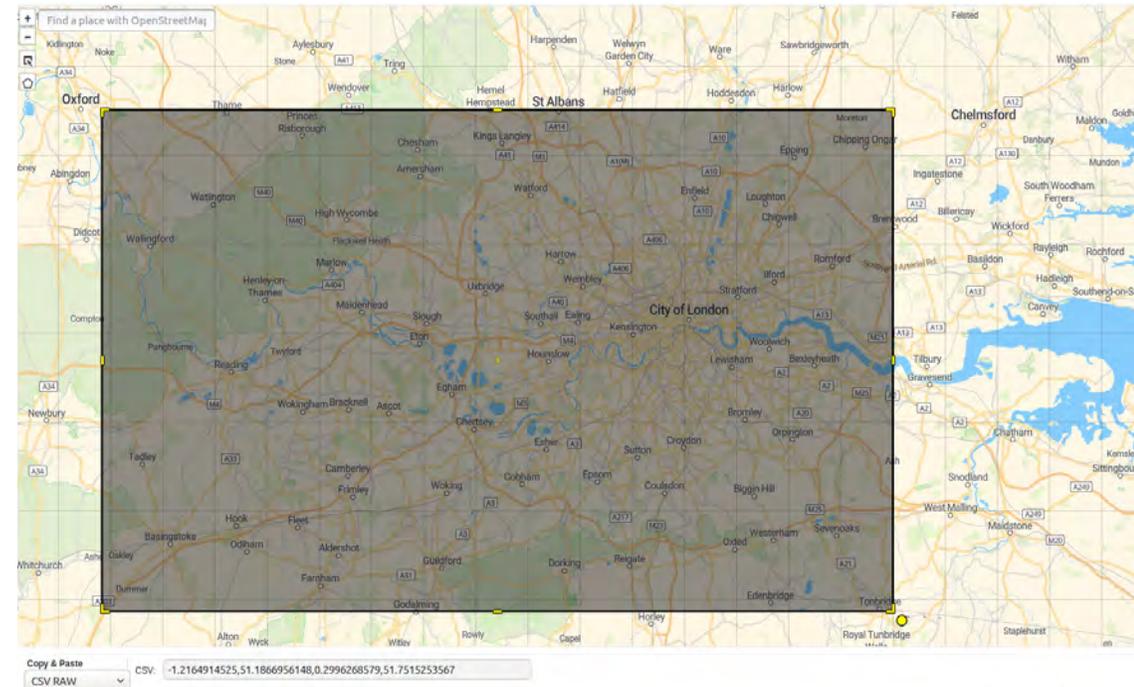


□ Retrieve tweets with keywords filter

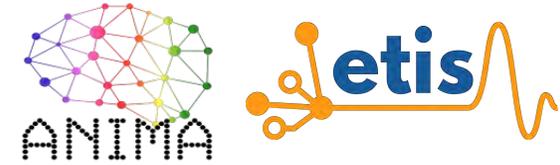
- “Heathrow” and “noise” or “LHR” and “noise” to filter tweets

Both retrieval methods are configured to get only English language tweets.

→ Two datasets: tweets with location query (TWLQ) and tweets with keywords query (TWKQ)

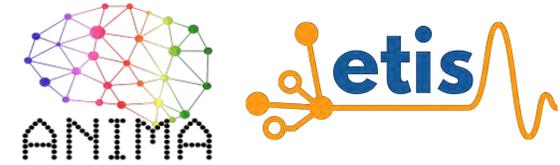


Workflow: Preprocessing tweets



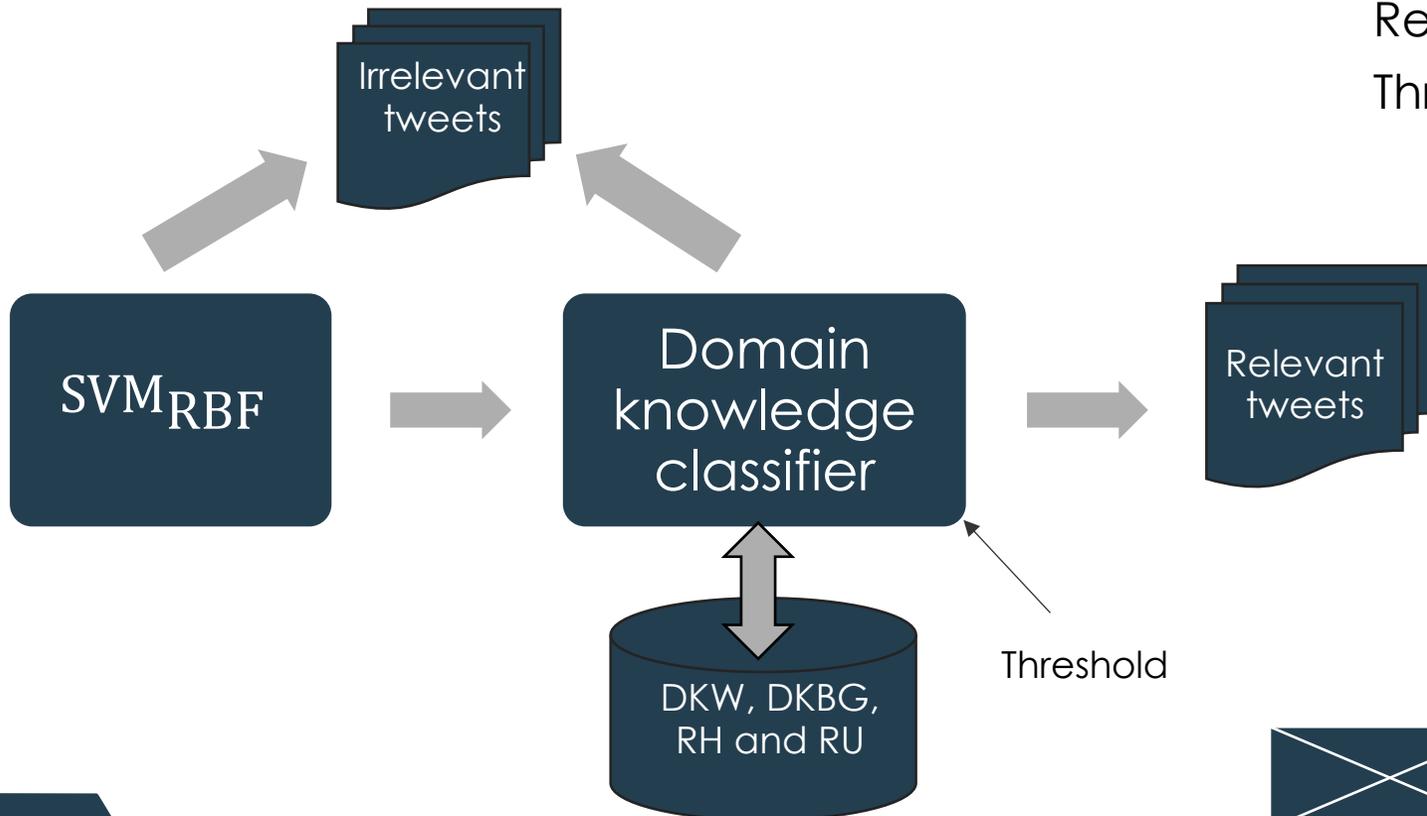
Preprocessing tweets		
For both Sentiment and relevance classification	Specific to relevance classification	Specific to sentiment classification
URL and punctuation removal	Numbers removal	Username removal
Tokenization (split words into tokens)	Stop words removal (a, the, this, can, etc.)	Emoticons extraction
Part-of-speech (POS) tagging	Lemmatization (bring words into their respective root words)	Intensifier detection: character repetition (e.g. "happyyy") and all caps (e.g. "NOISE")
Spelling correction (character repetition case)	-	Detect negation words and the respective scope
-	-	Detect contrast words

Workflow: Relevance classification



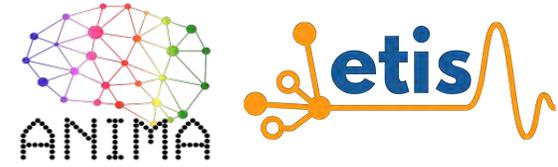
Tweets fed as a *tf-idf* matrix (t_i^T is tweet i in corpus T , d_j is term j): $t_i^T \rightarrow \begin{pmatrix} x_{1,1} & \dots & x_{1,n} \\ \vdots & \ddots & \vdots \\ x_{m,1} & \dots & x_{m,n} \end{pmatrix}$ } Relevant

Relevance score = $\frac{4}{4} = 1$
 Threshold = 0,5



	MNB	SVM _{Linear}	SVM _{RBF}	RF
F-measure	90.74%	95.53%	95.64%	94.46%

Workflow: Sentiment classification



Sentiment classification: Emoticons, lexicon and SWN

Emoticon (Em) score calculation:

- uses datasets of positive and negative emoticons (PE and NE)

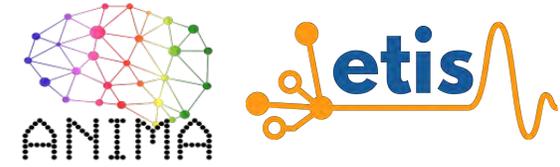
$$emscore(e_j) = \begin{cases} 1, (e_j \in PE) \\ -1, (e_j \in NE) \\ 0, (e_j \notin PE) \wedge (e_j \notin NE) \end{cases} \quad (1)$$

Lexicon polarity (LP) score calculation:

- uses datasets of positive and negative words (PW and NW)
- Datasets are created from multiple existing lexicon (MPQA **[13]**, Bing Liu **[1]** and Bill McDonald **[11]** → 10529 words).
- Subjective intensity of words: strong subjectivity, weak subjectivity and unknown subjectivity (weights)
- Intensifiers (ACI set of all caps intensifier score and CRI set of character repetition intensifier score of a tweet).

$$swscore_{LP}(w_j) = \begin{cases} 1 \times weight \times aci_j \times cri_j, (w_j \in PW) \wedge (aci_j \in ACI) \wedge (cri_j \in CRI) \\ (-1) \times weight \times aci_j \times cri_j, (w_j \in NW) \wedge (aci_j \in ACI) \wedge (cri_j \in CRI) \\ 0, (w_j \notin PW) \wedge (w_j \notin NW) \end{cases} \quad (2)$$

Workflow: Sentiment classification



- Negation with dynamic scope: stops when sentences ends (e.g. “,”, “.”, “-”, “!”, “?”) or when conjunction or contrast words are found (e.g. “and”, “where”, “which”, “but”). → opposite scores of words in the scope.
- Contrast effect: take the opposite scores of all words before the contrast word. Is not operating at sentence level.

Tweet: “@HeathrowAirport Stop these noise sewers, my kids are not sleeping. scum”

Negation scope 1

Negation scope 2

$LPscore = 0$

$j = 9$

$$swscore_{LP}(sleeping) = (1) * (1) * (1) * (1) = 1$$

$$swscore_{LP}(sleeping) = (-1) * swscore_{LP}(sleeping) = (-1) * (1) = -1$$

$$LPscore = 0 + (-1)$$

$j = 10$

$$swscore_{LP}(scum) = (-1) * (1) * (1) * (1) = -1$$

$$LPscore = (-1) + (-1) = -2$$

$$score_{LP}(\text{Tweet}) = \frac{(-2)}{10} = -0.2$$

Workflow: Sentiment classification



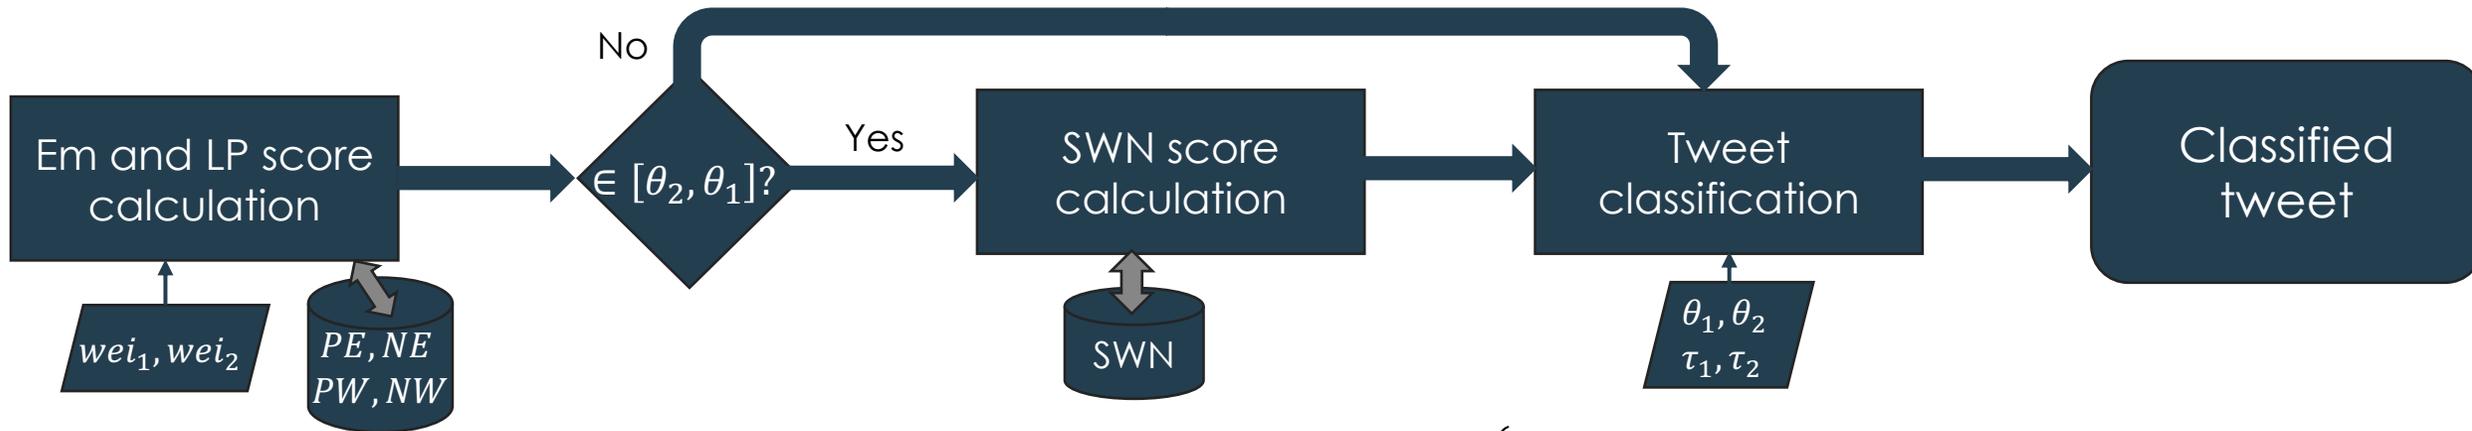
SWN score calculation: uses SWN dataset to get positive/negative score of word w_j and its synsets sy_j in SY_{w_j}

$$score_{SWN}(sy_j) = posscore_{SWN}(sy_j) - negscore_{SWN}(sy_j), (sy_j \in SY_{w_j}) \quad (3)$$

Proposed sentiment classifier (PC) score calculation and classification:

- Hierarchical use of Em, LP and SWN scores by weights and priority steps. First, use Em and LP score of the tweet to classify

$$score_{Em+LP}(\text{Tweet}) = wei_1 \times score_{Em}(\text{Tweet}) + wei_2 \times score_{LP}(\text{Tweet}) = 0,7 \times 0 + 0,3 \times (-0,2) = -0,06 \quad (4)$$



1st classification case: thresholds θ_1 and θ_2

$$sclass_{Em+LP}(\text{Tweet}) = \begin{cases} \text{positive}, & score_{Em+LP}(\text{Tweet}) > \theta_1 \\ \text{negative}, & score_{Em+LP}(\text{Tweet}) < \theta_2 \\ score_{SWN}(rt_j), & score_{Em+LP}(\text{Tweet}) \in [\theta_2, \theta_1] \end{cases} \quad (5)$$

2nd classification case: thresholds τ_1 and τ_2

$$sclass_{SWN}(\text{Tweet}) = \begin{cases} \text{positive}, & score_{SWN}(\text{Tweet}) > \tau_1 \\ \text{negative}, & score_{SWN}(\text{Tweet}) < \tau_2 \\ \text{neutral}, & score_{SWN}(\text{Tweet}) \in [\tau_2, \tau_1] \end{cases} \quad (6)$$

Post processing



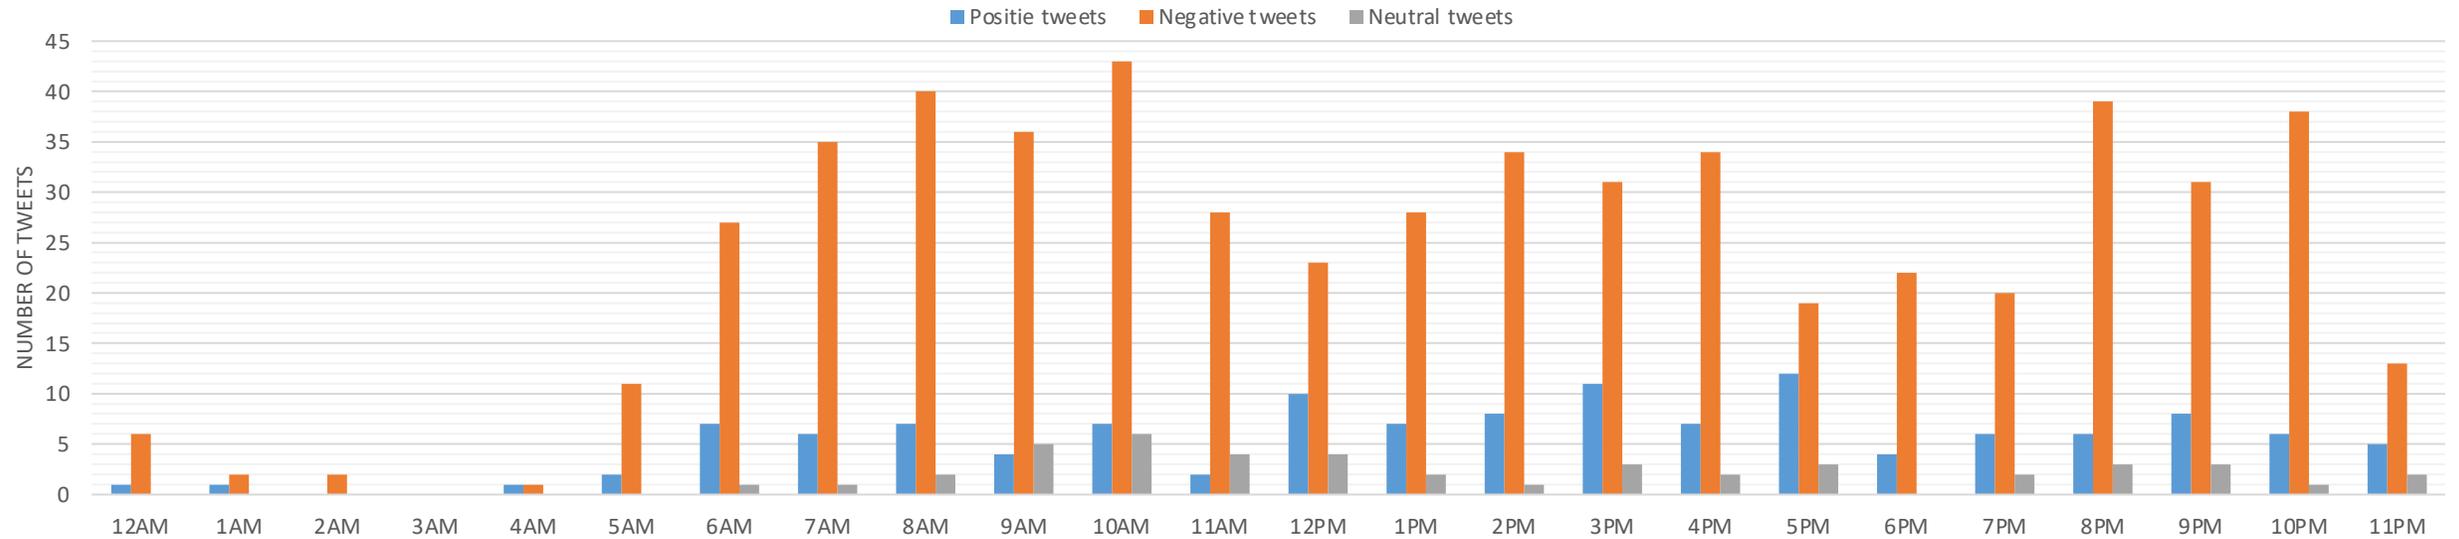
D_2	Relevant tweets	Total	Tweets with location	Positive tweets	Negative tweets	Neutral tweets
TWLQ	265	742	250	128	569	45
TWKQ	477					

D_2	Most mentioned user	Most used hashtags	Most retweeted tweets
Positive tweets	@HeathrowNoise: 17 @BBCParliament: 12 @StopHeathrowExp: 11	#Heathrow: 8 #HeathrowExpansion: 2 #HeathrowNoise: 2	“RT .@BBCParliament MPs who vote in favour of a third runway at Heathrow will be condemning yet more people to the torture of aircraft noise. Which ones think it's a good idea to fly another 260,000 flights pa over this densely-populated area?": 11
			“RT @HeathrowNoise See? It's reversible, as we always knew: "The changes were made under an agreement to provide aircraft noise relief to residents of historic neighborhoods": 4
Negative tweets	@HeathrowNoise: 164 @TeddingtonTAG: 45 @NeilSpurrier: 33	#Heathrow: 52 #heathrow: 28 #care2: 23	“RT Demand an end to noise sewers... #care2 #heathrow": 21
			“RT Here's an online petition to end the 'noise sewers' caused by more concentrated flight paths out of #Heathrow... #twickenham #teddington #whitton #etc": 20
Neutral tweets	@HeathrowNoise: 29 @NeilSpurrier: 21 @bakerairlondon: 9	#Heathrow: 1 #care2: 1 #NoiseActionWeek: 1	“RT @NeilSpurrier @HeathrowNoise JHK said planes at 1000ft by airport boundary, so can Matt or @Heathrownoise clarify why a procedure to 1000ft that only benefits the airport is in a community noise action plan?: 9
			“RT One must question whether @yourHeathrow have the slightest intention of being truthful with the communities surrounding the airport or whether their draft Noise Action Plan bears even a passing resemblance to reality. @HeathrowCEB": 7

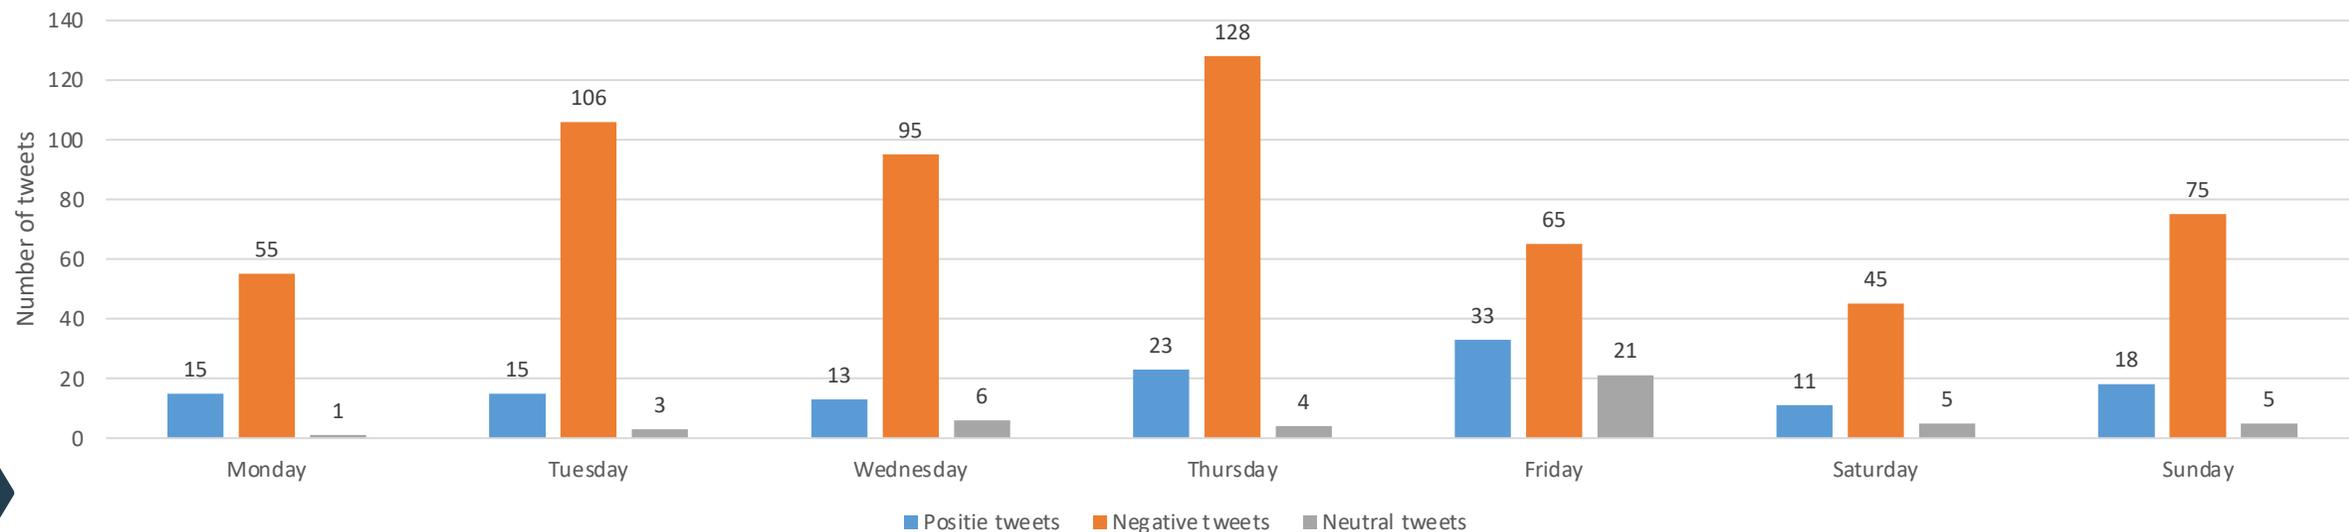
Post processing: Temporal analysis



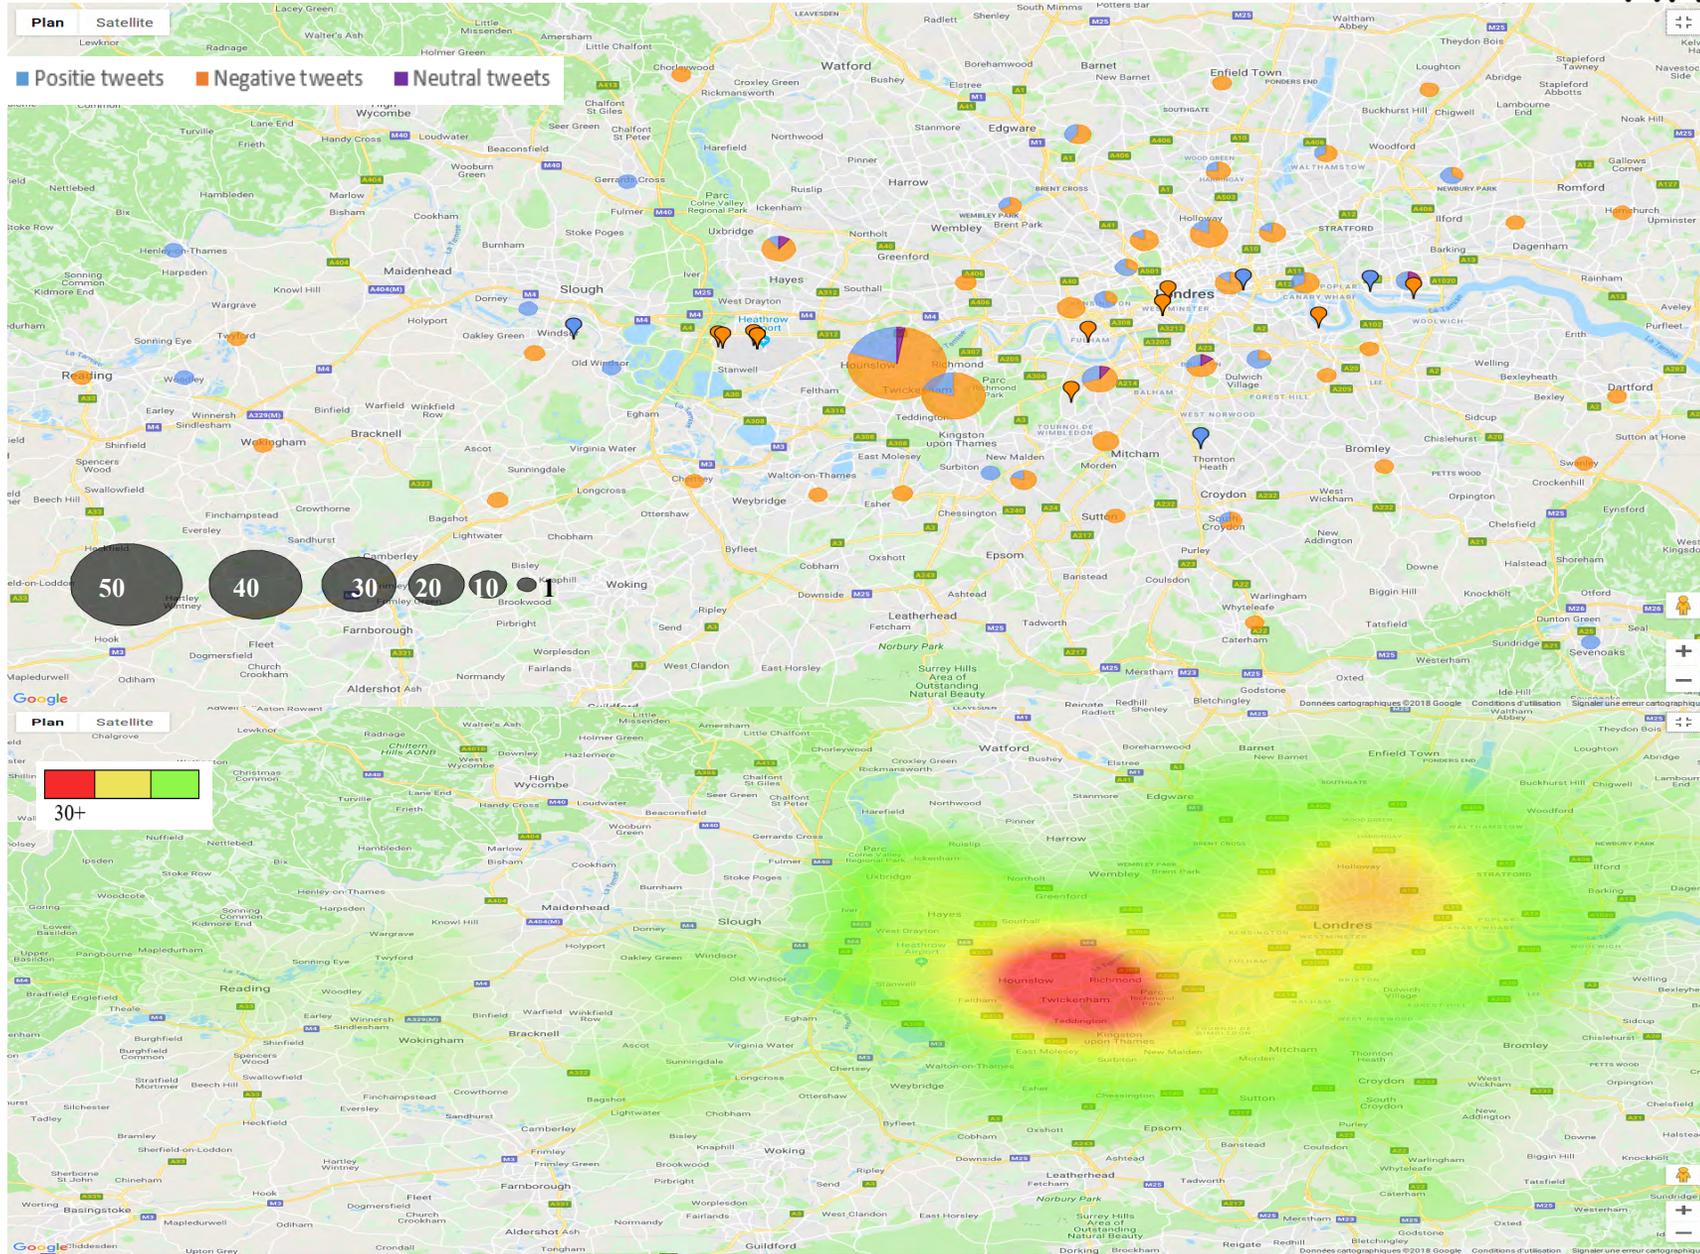
Daily temporal distribution of tweets in D2



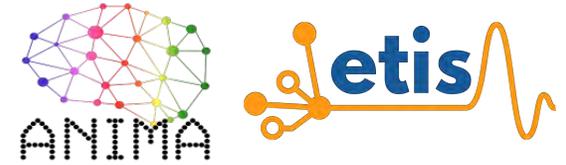
Weekly temporal distribution of tweets in D2



Post processing: Geospatial analysis



Conclusion and perspectives



Conclusion 

- Extract sentiments from twitter messages on airport noise
- Capture relevant tweets from a stream: SVM + Lexicon-based classifier
- Sentiment analysis: emoticons, subjective intensity of words, intensifiers, negation with dynamic scope, contrast and SWN
- Use hierarchy of emoticons, lexicon polarity and SWN scores to classify tweets
- Classified tweets are used to understand causes, time, areas and to extract related topics

perspectives 

- Improve relevance classifier: apply more domain knowledge features (e.g. presence of related expressions)
- Improve sentiment classifier: expand lexicon dataset, more types of spelling errors to correct, contrast effect at sentence level
- Post processing: compare affected areas and flight paths, extract related events not only with hashtags and create users links graph

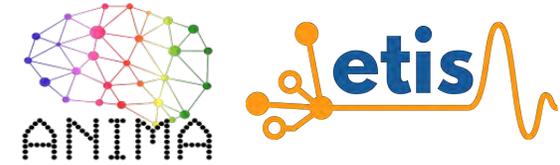
Perspectives – Phase 2, “Historical” Data



- Getting tweets from followers of community/relevant accounts
- Going back in time as we like
- More relevant tweets
- Ability to compare with other datasets, currently into the pipeline:
 - Tripadvisor
 - AirBnB
 - Land use/values
- Provide a toolkit applicable in different airports/areas
- Compare results with more “traditional” surveys
- Investigate bias in the data and the process

perspectives

Experimental results: Classifiers comparison



$wei_1 = 0.7$ and $wei_2 = 0.3$

$weight$ {
 strong subjectivity: 1
 weak subjectivity: 0,75
 unknown subjectivity: 0,75

	Positive tweets	Negative tweets	Neutral tweets	Total	Tweets with emoticons
D_1	26	601	26	653	18

D_1		Confusion matrix			Metrics				
		Positive	Negative	Neutral	<i>precision</i>	<i>recall</i>	<i>F – measure</i>	<i>accuracy</i>	
EmC	Thresholds	Positive	0	0	26	0 %	0 %	-	4.90 %
	$\tau_1 = \tau_2 = 0$	Negative	3	6	592	100%	1%	1,98 %	
		Neutral	0	0	26	4.04 %	100%	7.76%	
LPC	Thresholds	Positive	10	7	9	12.50 %	38.46 %	18.86 %	62.17 %
	$\tau_1 = 0.027$ $\tau_2 = -0.001$	Negative	66	380	155	96.69 %	63.23 %	76.46 %	
		Neutral	4	6	16	8.89 %	61.54 %	15.53 %	
SWNC	Thresholds	Positive	4	15	7	4.44 %	15.38 %	6.90 %	69.37%
	$\tau_1 = 0.015$ $\tau_2 = 0.005$	Negative	79	443	79	95.05 %	73.71 %	82.64 %	
		Neutral	7	13	6	6.52 %	23.07 %	10.17%	
PC	Thresholds	Positive	12	10	4	12.12%	46.15%	19.20 %	77.79 %
	$\theta_1 = 0.01$ $\theta_2 = -0.001$ $\tau_1 = 0.015$ $\tau_2 = 0.004$	Negative	80	491	30	95.34 %	81.70 %	87.99 %	
		Neutral	7	14	5	12.82 %	19.23 %	15.38 %	

→ Detects only tweets with emoticons

→ Many false neutral tweets

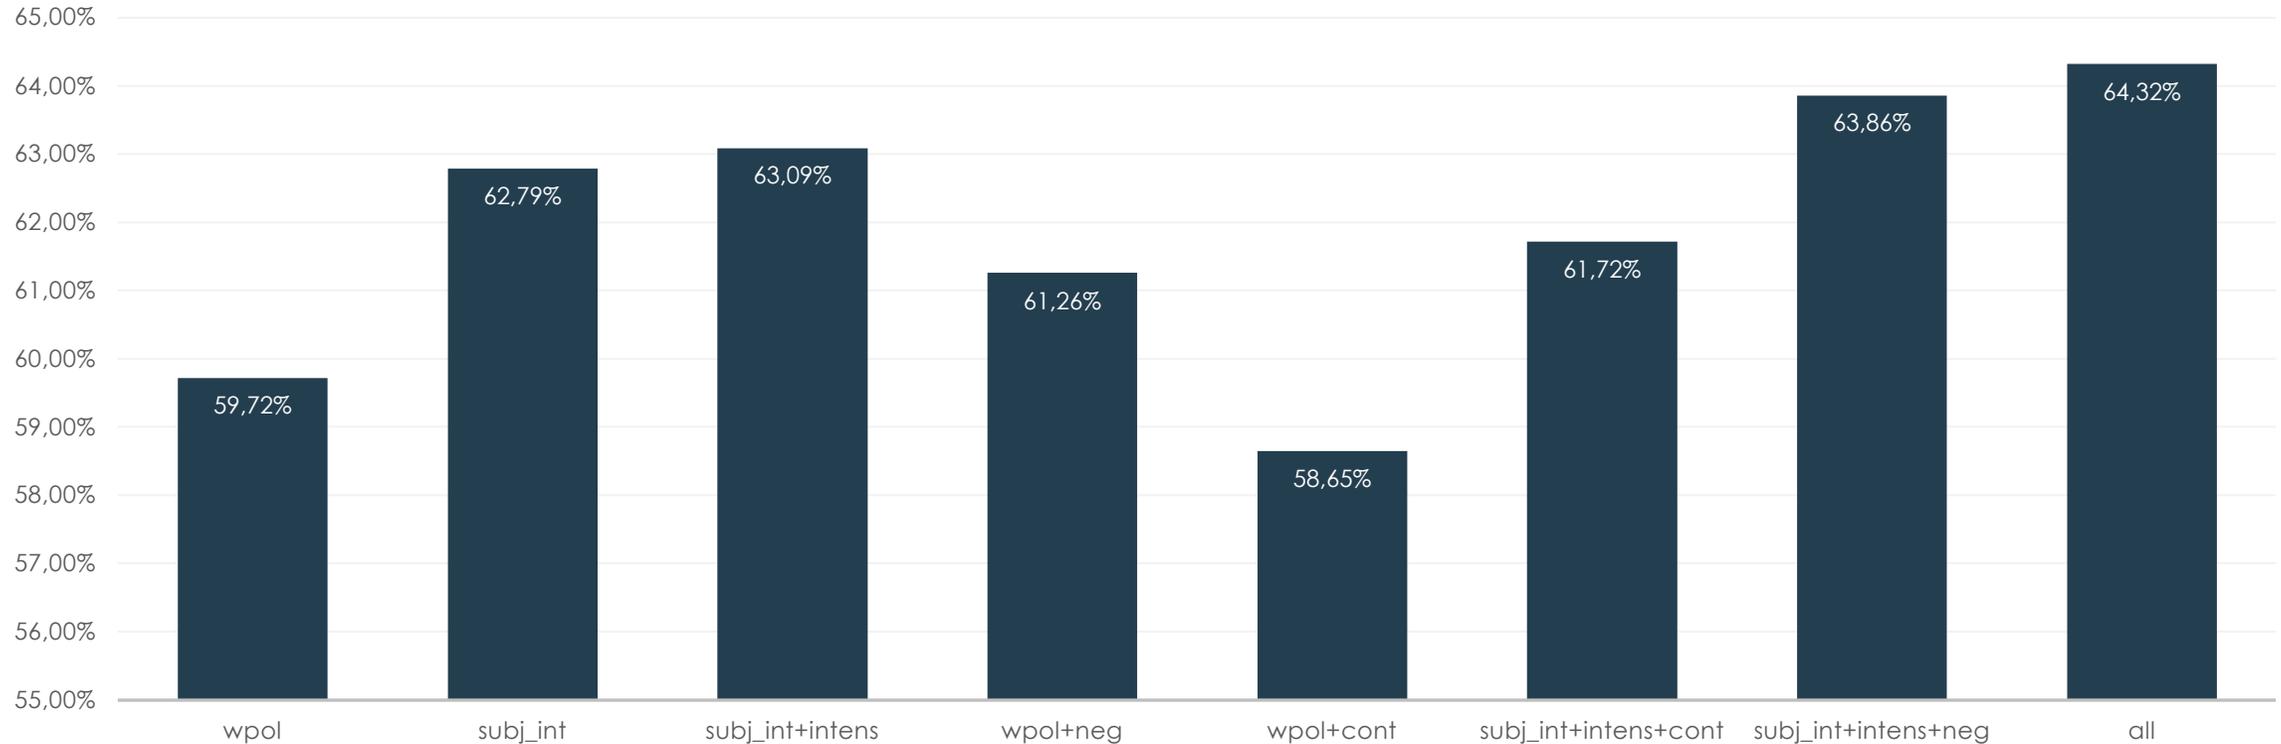
→ Few true positive and true neutral tweets

→ Decrease false positive and false neutral and increase true positive

Experimental results: LPC features comparison



LPC accuracy



wpol: word polarity

subj_int: subjective intensity of words

subj_int+intens: subjective intensity of words and intensifiers

wpol+neg: word polarity and negation

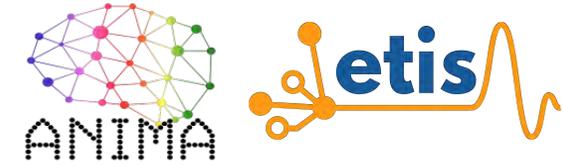
wpol+cont: word polarity and contrast

subj_int+intens+cont: subjective intensity, intensifiers and contrast

subj_int+intens+neg: subjective intensity, intensifiers and negation

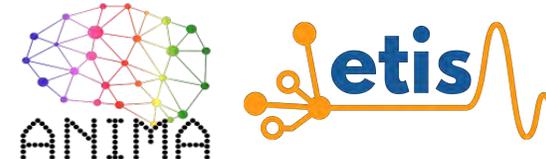
all: subjective intensity, intensifiers, negation and contrast

State of the art



Problem	Methods	Advantages	Disadvantages
Relevance	Keyword query	simple and fast	Captures also many irrelevant tweets [8]
	Feature-based techniques [12] / Rule-based techniques [10]	Offer more accuracy	Specific to the topic and its characteristics
	Document-pivot techniques (<i>tf-idf</i> [4] , <i>named entity</i>)	Based on similarity of texts	Less suitable for tweets
Sentiment Analysis	Emoticons	Text and topic independent [6]	Effective only for tweets with emoticons
	POS features	Separate objective and subjective sentences	Tagger errors due to tweets
	Convolutional neural nets [14]	Good to learn subjective expressions (n-grams)	Time consuming
	Lexicon-based	Detect the sentiment of a tweet based the prior polarity of its words	Needs to have big lexicon dataset / Performs better with other features (negation, contrast) [9]

Other examples



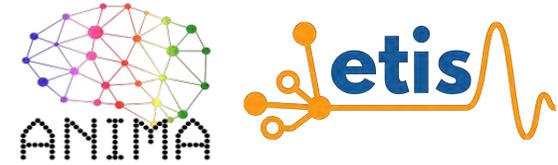
Tweet: BBC News - #Heathrow: Cabinet set for new runway decision five and a half hours sleep is not enough sleep for healthy living #Nothirdrunway #HeathrowExpansion #NoisePollution #Noise #AirQuality #Asthma → LP negative

Tweet: RT It's the wind that has been a lot of easterly recently. However if this garbage airport did something about the noise we wouldn't have to suffer so much. → Unclassified by LP → SWN classified negative

Tweet: This man lives closer to Heathrow Airport than anyone else and claims noise is not a problem, but that's only if planes are landing and not departing over his house, his windows remain closed and he's not out in his garden, confused? I am. → LP classified negative

Tweet: Interesting to see how people are so badly affected by noise. I live under the Heathrow flight path and right next to the piccadilly line and have never been kept awake at night by the noise → LP positive

Misclassifications



Tweet: RT Detailed evidence from community groups @TeddingtonTAG @HACAN1 with expert science clearly shows Heathrow needs to be "better not bigger" addressing noise air pollution for the sake of hundreds of thousands of Londoners. Londoners and London MPs cannot support this. #No3rdRunway → LP positive

Tweet: "@Outofthecontour So agree! All life is governed by LHR flight paths and schedules now." → Unclassified by LP → SWN classified neutral

Tweet: People can complain about church bells & cowbells & sue against noise caused by neighbors including the washing up & damage to hearing from playing in an orchestra but disgusting that the government wants to subject hundreds of thousands more to aircraft noise due to lobbying → LP positive



Tweet: RT Want to sleep with the window open? Would be lovely on a hot evening, but
#planes #Heathrow #noise. Please stand against night time take-offs

$$LPscore = 0$$

$$j = 3$$
$$swscore_{LP}(sleep) = (1) * (1) * (1) * (1) = 1$$

$$LPscore = 0 + 1$$

$$j = 7$$
$$swscore_{LP}(open) = (1) * (0,75) * (1) * (1) = 0,75$$

$$LPscore = 1 + 0,75$$

$$j = 13$$
$$swscore_{LP}(hot) = (-1) * (0,75) * (1) * (1) = -0,75$$

$$LPscore = 1,75 + (-0,75) = 1$$

$$j = 15$$

$$LPscore = (-1) \times LPscore = -1$$

$$j = 17$$
$$swscore_{LP}(stand) = (1) * (0,75) * (1) * (1) = 0,75$$

$$LPscore = -1 + 0,75 = -0,25$$

$$score_{LP}(\text{Tweet}) = \frac{(-0,25)}{25} = -0.1$$

$$score_{Em}(\text{Tweet}) = 0$$

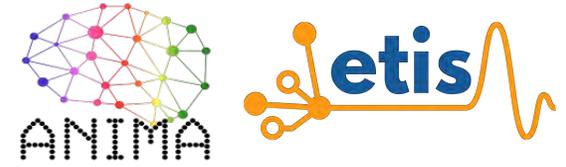
$$score_{Em+LP}(\text{Tweet}) = 0,7 \times 0 + 0,3 \times (-0,1)$$
$$= -0,03$$

$$-0,03 < \theta_2 = -0,01 \rightarrow \text{Negative tweet}$$

- ▶ Example Tweet: “@HeathrowAirport Stop these noise sewers, my kids are not sleeping. scum” → negative

- [1] B. Liu, M. Hu and J. Cheng, "Opinion Observer: Analyzing and Comparing Opinions on the Web.," in *Proceedings of the 14th International World Wide Web conference (WWW-2005)*, Chiba, Japan, 2005.
- [2] C. Cortes and V. Vapnik, "Support-vector networks," *Machine Learning*, pp. 273-297, 1995
- [3] Civil Aviation Authority, "Heathrow Airport 2016 Summer Noise Contours and Noise Action Plan Contours," London, 2017.
- [4] G. Salton, *Automatic Text Processing: The Transformation, Analysis and Retrieval of information by Computer*, Boston, Massachusetts, USA: Addison-Wesley, 1988
- [5] H. Zhang, "The Optimality of Naive Bayes," in *Proceedings of Florida Artificial Intelligence Research Society Conference (FLAIRS)*, 2004.
- [6] J. Read, "Using emoticons to reduce dependency in machine learning techniques for sentiment classification," in *ACL the Association for Computer Linguistics*, 2005
- [7] L. Breiman, "Random Forests," *Machine Learning*, vol. 45, no. 1, pp. 5-32, 2001.
- [8] M. Krieck, J. Dreesman, L. Otrusina and K. Denecke, "A New Age of Public Health: Identifying Disease Outbreaks by Analyzing Tweets," in *Proceedings of Health WebScience Workshop, ACM Web Science Conference*, 2011.
- [9] M. Z. Asghar, A. Khan, S. Ahmad, M. Qasim and I. A. Khan, "Lexicon-enhanced sentiment analysis framework using rule-based classification scheme," *PLoS ONE*, 2017.
- [10] O. Dan, J. Feng and B. D. Davison, "A Bootstrapping Approach to Identifying Relevant Tweets for Social TV," in *Proceedings of the Fifth International AAAI Conference on Web and Social Media*, North America, 2011.

Bibliography



- [11] T. L. B. McDonald, "When is a Liability not a Liability? Textual Analysis, Dictionaries, and 10-Ks,," *Journal of Finance*, vol. 66, no. 1, pp. 35-65, 2011
- [12] T. Sakaki, M. Okazaki and Y. Matsuo, "Earthquake Shakes Twitter Users: Real-Time Event Detection by Social Sensors," in *International World Wide Web Conference*, Raleigh, North Carolina, USA, 2010.
- [13] T. Wilson, J. Wiebe and P. Hoffmann, " Recognizing Contextual Polarity in Phrase-Level Sentiment Analysis," in *Proceedings of HLT-EMNLP-2005*, Vancouver, Canada, 2005
- [14] Y. Ren, Y. Zhang, M. Zhang and D. Ji, "Context-Sensitive Twitter Sentiment Classification Using Neural Network," in *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, Phoenix, Arizona, USA, 2016