Vaccine hesitancy in Italy: can we monitor it in real time?

L'esitazione sui vaccini in Italia: possiamo monitorarla in tempo reale?

C. Rizzo & S. Cheatham
Università di Pisa & Ospedale Pediatrico Bambino Gesù

Pisa, 8 Luglio 2022
Polo Didattico S. Rossore 1938 – Via Risorgimento 23

Vaccine hesitancy in Italy: can we monitor it in real time?

Social media is increasingly being used to express opinions and attitudes towards vaccines. Can we monitor in real-time the vaccine conversation using social media posts ?

This presentation describes our journey developing a platform to monitor public vaccine confidence through social media.

- Vaccine hesitancy in Italy
- EU-JAV platform for social media monitoring
- Development of a machine learning model for automatic classification of the stance towards vaccines expressed on Twitter
- Conclusions

Vaccine hesitancy

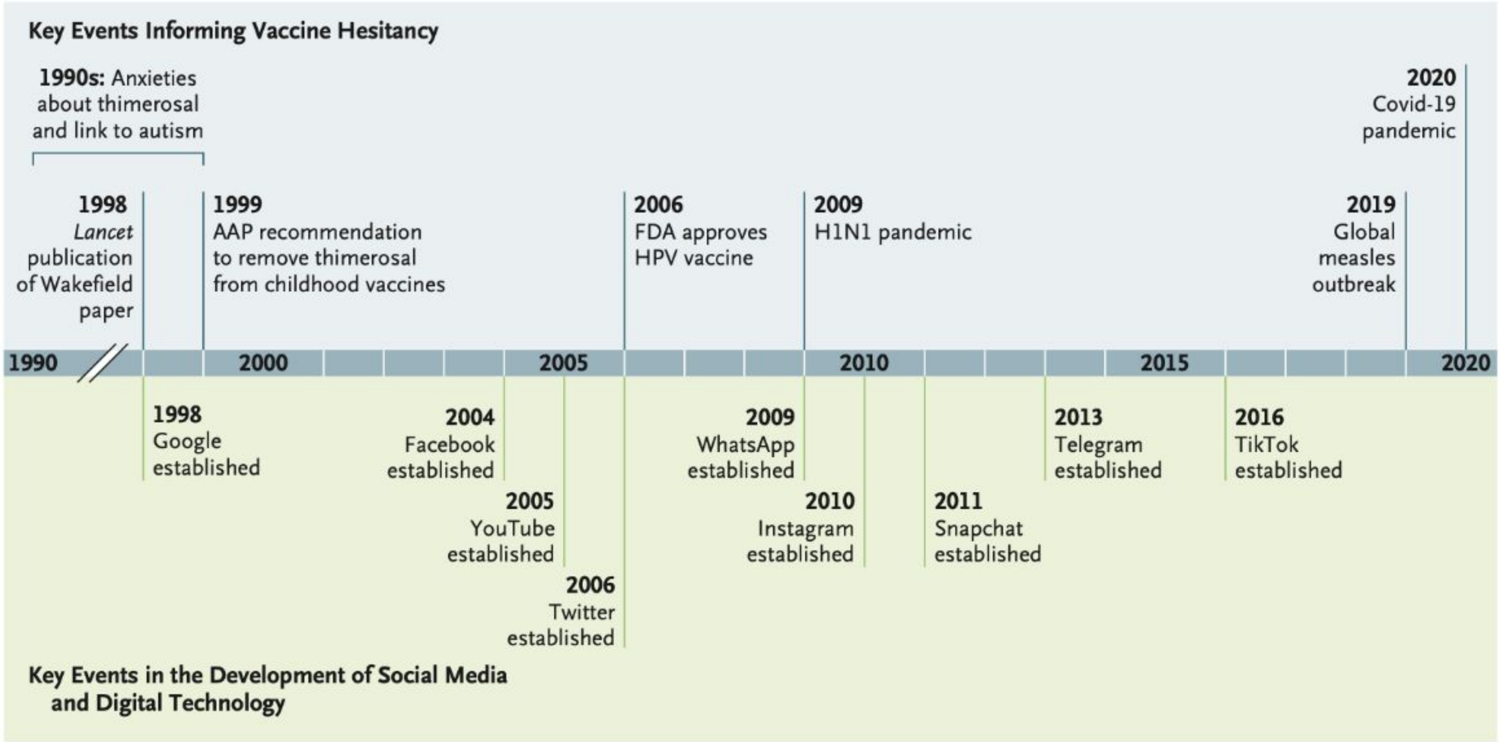

Figure 1. Timeline of Key Events Prompting Vaccine Hesitancy and Milestones in the Expansion of Social Media and Digital Technology.

EU-JAV platform monitoring social media

country selection

timeframe selection

http://www.opbg.cloud/#/

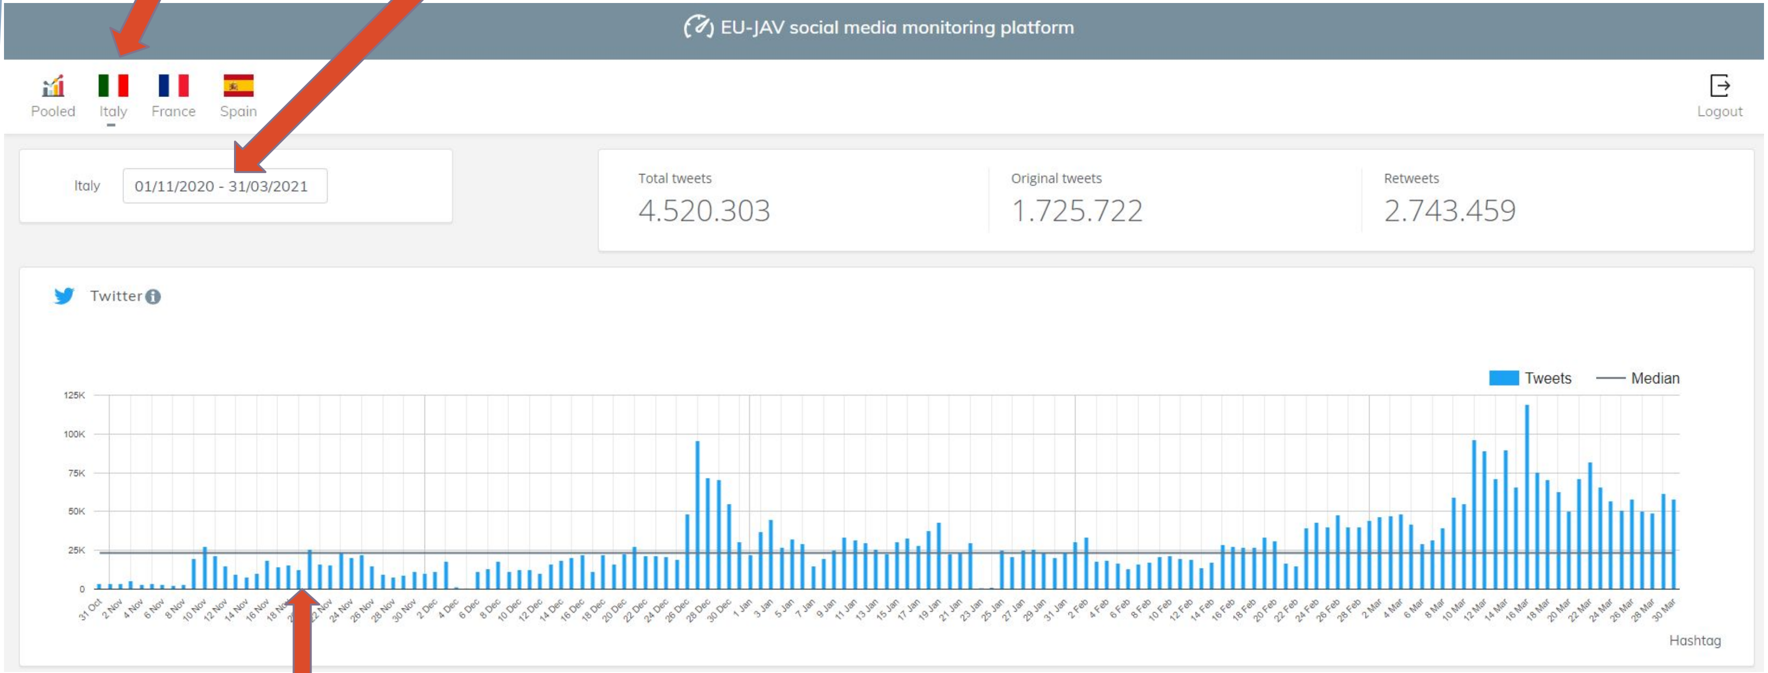

daily volume of vaccine-related tweets
based on validated keyword filter

EU-JAV platform monitoring social media
Event detection

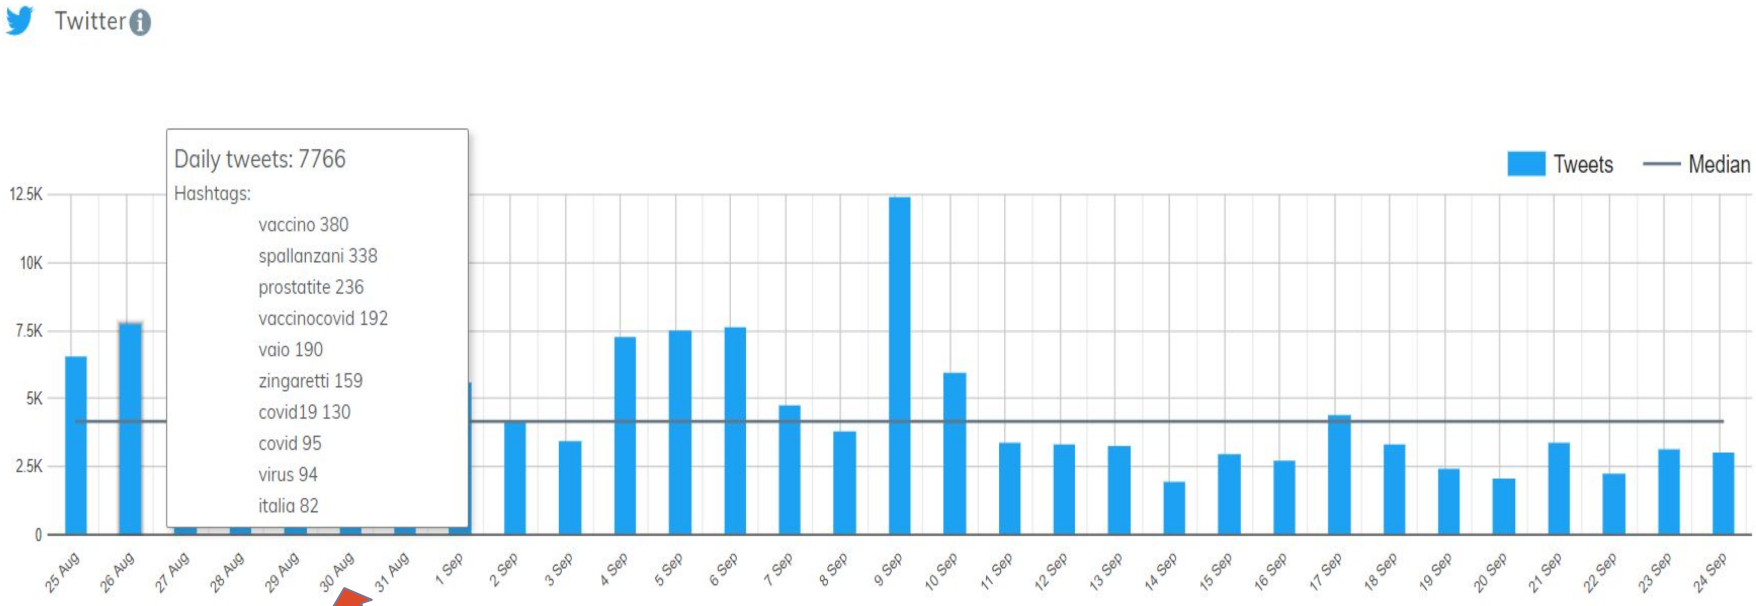

hovering the mouse over the column extra information are presented:
- precise number of daily tweets
- most popular hashtags and their daily frequency

EU-JAV platform monitoring social media
Most relevant users in the vaccine discourse on Twitter

Top influencers (Last month)

Ranking according to the mean between number of retweets done and received. Low ranking stands for many retweets (hence, high importance).

Screen Name	Times Retweeted	Number of Retweets	Relevance ↑	Most Central	Closest	Pagerank
Stefbazzi	885	17	17	19	2	1
PatriziaRametta	625	41	4	2	1	2
AntonioGrzt	259	2	205	28	5	3
MinervaMcGrani1	381	80	2	3	6	4
Cartabellotta	98	44	13	11	620	5
BarbaraRaval	277	32	9	34	10	6
piersar62	244	62	3	17	11	7
fbordo	40	8	82	13	305	8
valy_s	306	48	5	4	4	9

1 – 9 of 45

- **Relevance**: Top ranking = many retweets (**high level of activity**).
- **Most central**: Top ranking = the user frequently **acts as a bridge connecting other users**, and therefore can influence the **largest part of the community**.
- **Closest**: Top ranking = to users that are well positioned in the graph **to influence other users as fast as possible**.
- **Pagerank**: Top ranking = users who has the potential to **influence users that are not directly connected** with them.

EU-JAV platform monitoring social media
Google Trends

Google Trends

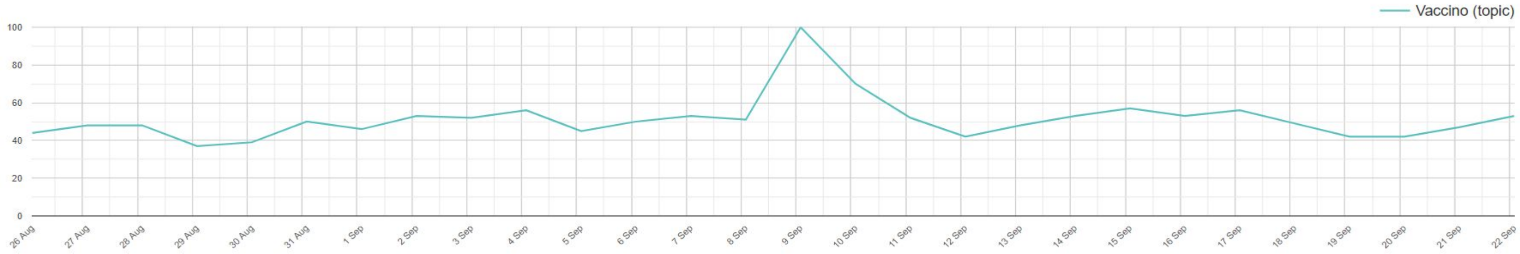

Google Trends GeoChart

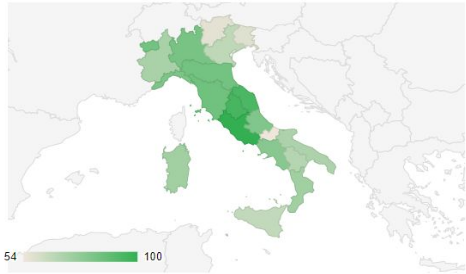

Related Queries

- coronavirus
- coronavirus vaccino
- vaccino covid
- covid
- vaccino antinfluenzale
- vaccino antinfluenzale 2020
- vaccino oxford

Vaccine mentioned in a random sample of 3000 tweets

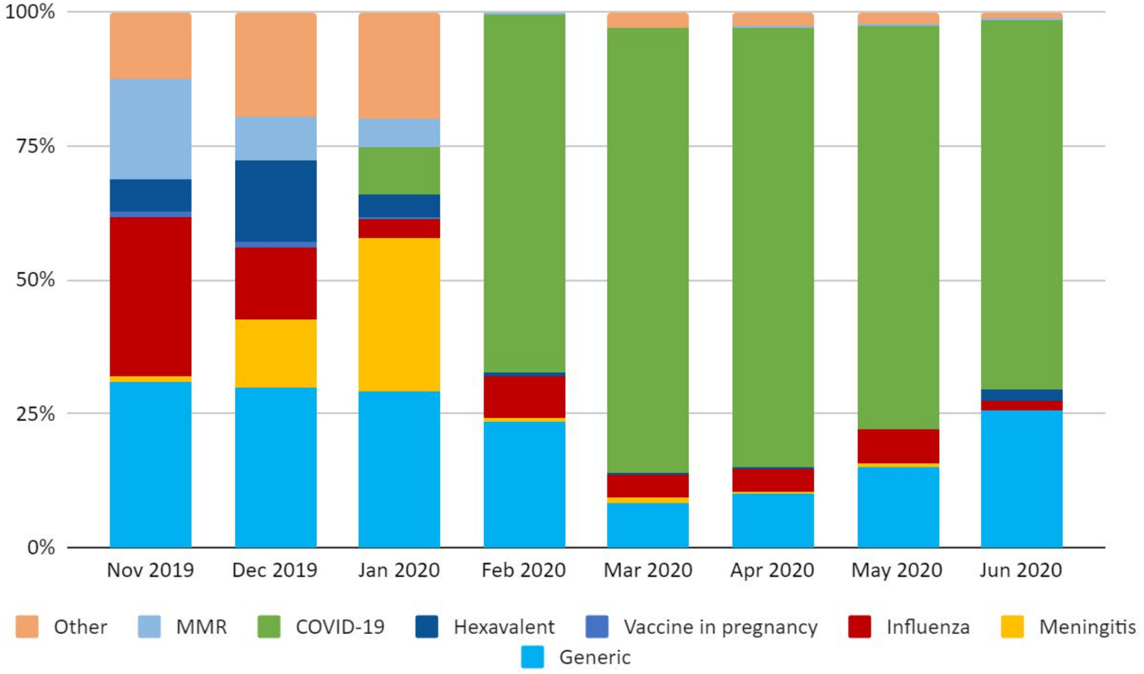

Since February 2020 the discussion about vaccines has been dominated by covid-19 mentions

How can we better understand vaccine hesitancy? Monitor stance?

Contents lists available at ScienceDirect

Vaccine

journal homepage: www.elsevier.com/locate/vaccine

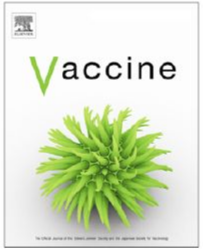

“Vaccines for pregnant women…?! Absurd” – Mapping maternal vaccination discourse and stance on social media over six months

Sam Martin [a,*], Eliz Kilich [a], Sara Dada [a], Per Egil Kummervold [c], Chermain Denny [b], Pauline Paterson [a], Heidi J Larson [a,d]

[a] *Faculty of Epidemiology and Population Health, London School of Hygiene & Tropical Medicine, London, UK*
[b] *Faculty of Science, Vrije University, Amsterdam, Netherlands*
[c] *FISABIO-Public Health – Vaccine Research Department, Valencia, Spain*
[d] *Institute of Health Metrics and Evaluation, University of Washington, Seattle, USA*

What is stance? Sentiment vs stance

Sentiment analysis determines the overall tone or language of some text, whether it is positive, negative or neutral. It does not consider the meaning or message of the text.

Stance analysis determines favorability towards a chosen topic of interest.
E.g. A tweet criticising the anti-vax movement can have a negative sentiment, despite expressing a favourable opinion towards vaccines

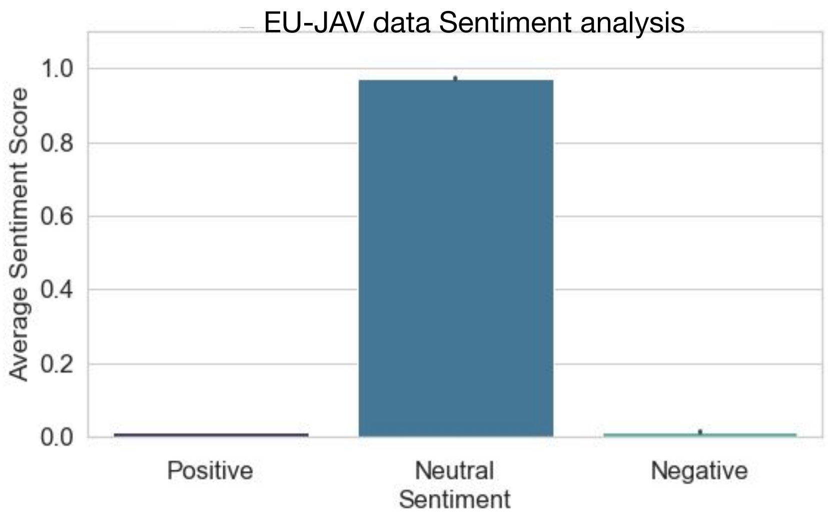

Using Python Natural Language Toolkit SentimentIntensityAnalyzer, the majority of tweets have neutral sentiment.

This demonstrates that a more nuanced study of the data is required to extract meaning.

EU-JAV Data: Stance Labels

Promotional: communicate public health benefits or safety of vaccination, encourage vaccination. Contains negative attitude against not vaccinating. Describe risks of not vaccinating

Neutral: no elements of uncertainty, promotional or negative content. Often statements. Includes factual recommendations or reporting information

Discouraging: contains negative attitude/arguments against vaccines. Contains questions regards effectiveness/safety or possibility of adverse reactions that may or may not be proven

Ambiguous: Content contains indecision/uncertainty on the risks or benefits of vaccination

Other: often in a language other than Italian, or the word vaccine used in a figurative way

Well labelled data is key to the performance of the machine learning model.
Labelling the data takes a lot of time and effort.

Example tweets and their stance

Text	Stance
My children are vaccinated against #chickenpox. There's been an epidemic in their school and they've been fine, happy, did not have any symptom. Davide was born in 2012, we payed for the vaccine because it was not free for his age. Renounce everything but #vaccinate, for your children and for the others.	Promotional
Coronavirus, WHO: For the vaccine we'll have to wait at least for 18 months.	Neutral
#vaccines DISABLED for a #fluvaccine. Ministry of Health condemned to pay. Vaccines are DRUGS. Be wary of all those that say that vaccines are safe and do not have #sideffects. You can't go back.	Discouraging
I still have to decide and I'm scared…if I get the vaccine, I don't know side effects…but what if I don't get the vaccine and I get covid? for how long will I be in danger? If I get vaccinated, how long is the vaccine going to work? I'm scared…any advice?	Ambiguous

Vaccine stance in a random sample of 3000 tweets

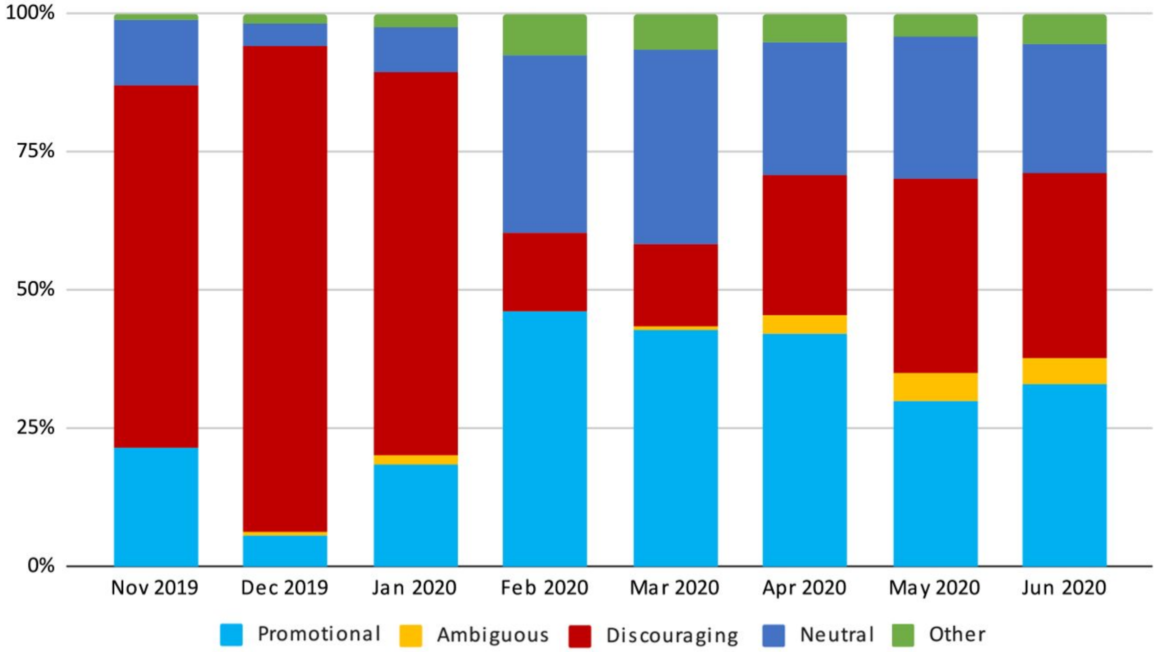

Strong representation of discouraging tweets in the Italian vaccine discourse.

With the pandemic, increased vocality of neutral and promotional tweets.

Natural Language Processing:
How do humans understand the meaning of a sentence?

Language contains words with one or more meanings and grammar rules which determine sentence structure.

The meaning of a sentence is interpreted by the separate words, and by the positions and presence of the other words.

Cognitive attention allows humans to concentrate on selected stimuli. It is not just about centering focus on one particular thing; it also involves ignoring competing information and stimuli.

Huge quantities of text allow the machine learning algorithms to work out the rules.

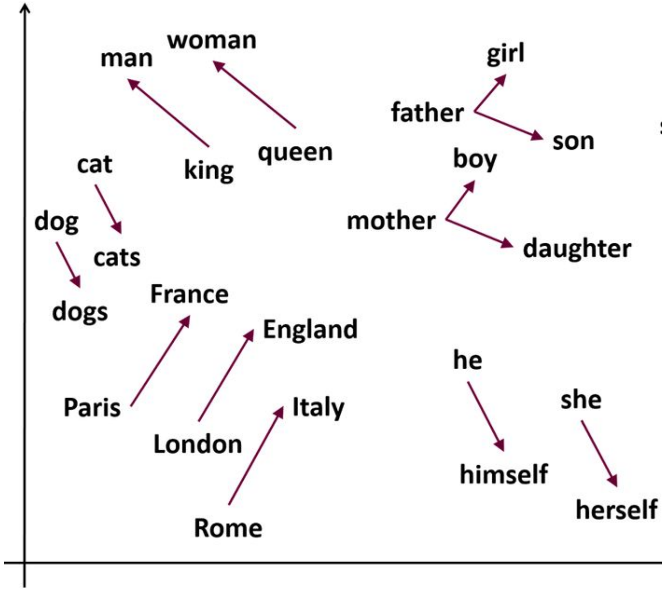

Attention is a machine learning technique that mimics cognitive attention. Which part of the data is more important than others depends on the context and is learned through training

Machine Learning

In natural language processing, models can be developed and **pre-trained** using huge quantities of unlabelled text. Pre-training imitates the way human beings process new knowledge.

- **Unsupervised learning**: no labels are given to the learning algorithm
- **Supervised learning**: the model is fed with example data and labels

For model **fine-tuning** labelled data is split into 3 independent sets

- **Training**: used for to construct a predictive relationship
- **Validation**: used to build the model, to assess the performance of the model during training
- **Testing**: the model labels the data, then uses the human labels to assess the accuracy of the model

EU-JAV Twitter data

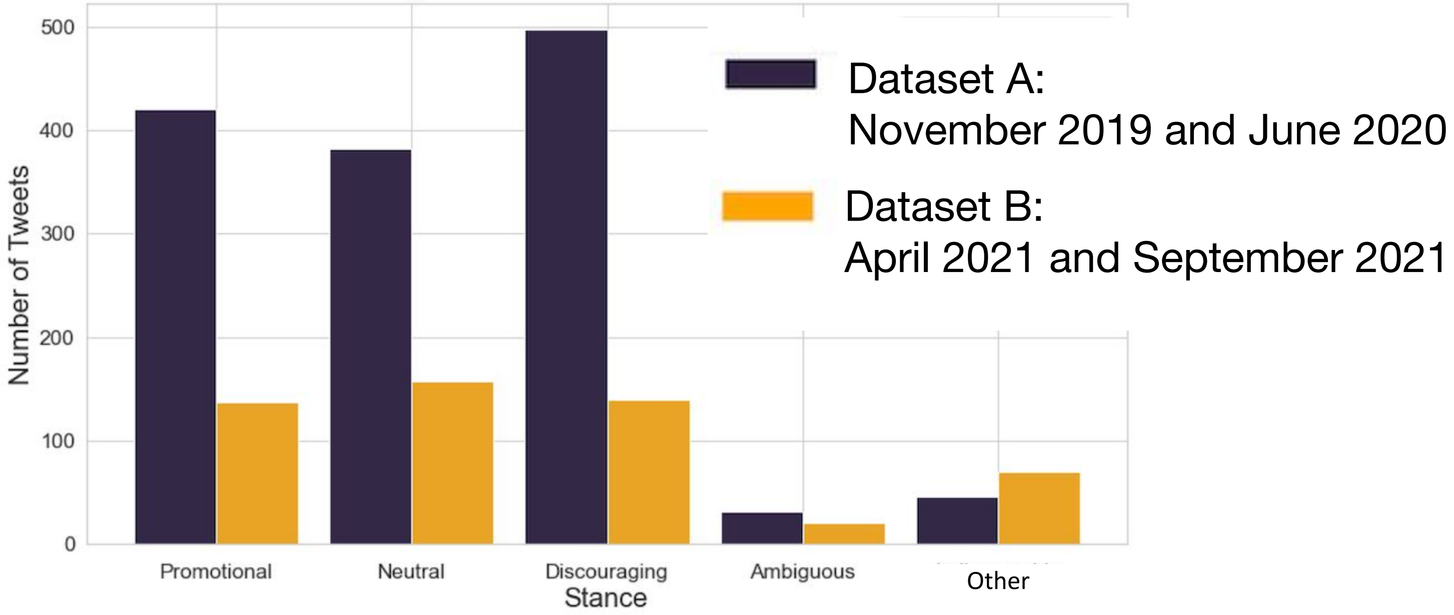

Data from Twitter collected using the EU-JAV platform for social media monitoring.
Two distinct datasets: A and B, 10 months apart.

Given the very low proportion of ambiguous and other tweets, and the inconsistent nature of the texts, the data was reduced to 3 categories: promotional, neutral and discouraging.

Machine learning models

A Transformer is a deep learning model that adopts the mechanism of self-attention, differentially weighting the significance of each part of the input data.

Transformers were introduced in 2018 by a team at Google and are now ubiquitous in natural language processing problems.

Bidirectional Encoder Representations from Transformers (BERT) models come pre-trained on huge unlabelled datasets, which can then be fine-tuned with a much smaller task-specific labelled dataset.

Multilingual models are language agnostic, encoding text from different languages into a shared embedding space.

Development of a machine learning model to monitor stance: methodology

- Logistic Regression used to understand what accuracy can be expected from a machine learning model. Logistic Regression is a well-established classification model which uses maximum likelihood estimation.
- Machine learning models identified that might be relevant to our study:
  - The vast majority of machine learning models are developed for English text, we wanted to keep the data in Italian.
  - On Hugging Face, an open source library of NLP models, there are around 15 000 mono-lingual models in English and just 10 000 mono-lingual models in all other languages.
- Hyperparameters (eg learning rate, warm-up, epochs) tuned for best performance of each model.
- Accuracy of the models assessed to select best model for our task.

Stance analysis results

	Accuracy (%)	F-score
Annotator average	**83.7**	**0.817**
Annotator 1	83.9	0.814
Annotator 2	88.1	0.865
Annotator 3	79.2	0.772
Logistic Regression	63.2	0.625
XLM-RoBERTa-large	**72.7**	**0.720**

The XLM-RoBERTa-large model achieved an accuracy 72.7% (F-score= 0.720) compared to the agreed score between the three annotators. The average accuracy annotators was 83,7% (F-score= 0.817)

This shows that machine learning models can achieve close to the same accuracy as individual annotators at labelling such data. The huge advantage of machine learning models is that they can classify tweets in near real-time.

Social media language evolves: need to regularly retrain social media monitoring algorithms

Train dataset	Test dataset	Accuracy (%)	F-score
dataset A	dataset A	72.8	0.713
dataset A	dataset B	62.1	0.617
dataset A	dataset A+B	70.1	0.689
dataset A+B	dataset A	72.8	0.721
dataset A+B	dataset B	71.3	0.713
dataset A+B	**dataset A+B**	**72.4**	**0.720**

Training on dataset A and testing on dataset A results in an accuracy of 72.8%.
Testing the same model on dataset B the accuracy drops to 62.1%. A drop of about 10%.

Dataset B is too small to be used for independent training but **if the model is retrained** with the full dataset (A+B), then the accuracy of the predictions for dataset A remains the same and the accuracy for dataset B **increases by 9%** (from 62.1% to 71.3%).

Retraining a model on more recent data improves the model's performance on that data.

EU-JAV platform monitoring social media
Stance

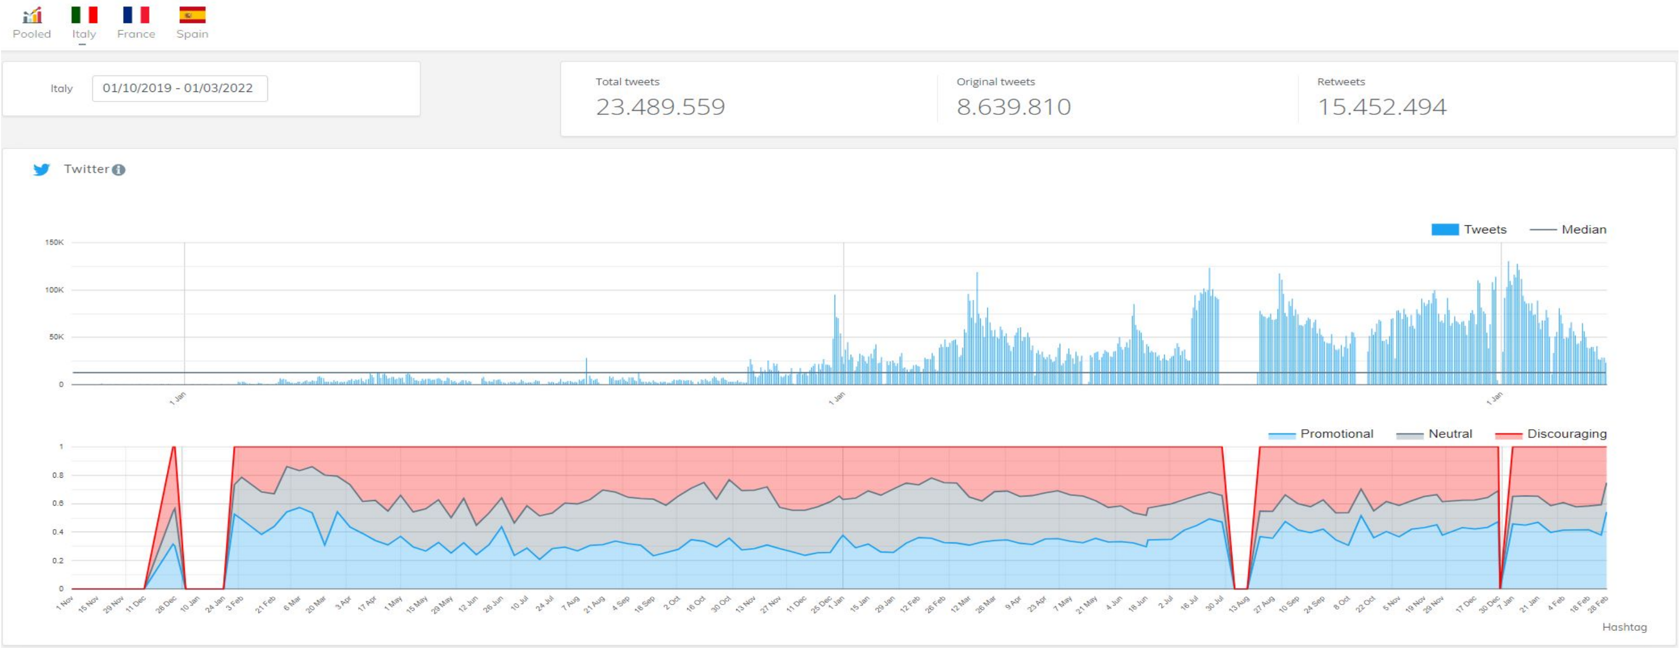

The fine-tuned model is used on the EU-JAV monitoring website to classify tweets and present the current stance of the social media discourse about vaccines on Twitter

Monitoring social media using ML: Limitations

Italian Twitter language is intellectual, sophisticated and rhetorical, making tweets difficult to classify.

The selected machine learning Transformer-based model, XLM-RoBERTa-large, was published in November 2019. It was pre-trained on 2.5TB of data containing 100 languages.

- SARS-CoV-2 was first confirmed to have spread to Italy on 31 January 2020, thus the model had not seen any COVID-19 related conversations or COVID-19 specific language. E.g. words such as "mask", "confinement" and "vaccine" are not present in the XLM-RoBERTa vocabulary. Pre-training the XLM-RoBERTa-large model on recent unlabelled Twitter data may well increase its performance on this task.

Kummervold et al. studied the stance of tweets on maternal vaccination using Transformer-based machine learning models. They achieved an accuracy of 81.8% (F-score= 0.776).

- The lower accuracy of our model might be in part due to the reduced dataset size (1736 cf 2722 tweets), or to the drift in stance over the extended time period of our data collection (Kummervold's data collected over 6 month period).
- The Kummervold study was carried out in English (non-English tweets were translated into English using a Google Translate script), whereas our study was carried out in Italian. Work on machine learning models has predominantly been done in English.

Conclusion
Vaccine hesitancy in Italy: can we monitor it in real time?

- The EU-JAV platform http://www.opbg.cloud/#/ monitors social media in near real time: daily volume of vaccine-related tweets, most popular hashtags, most relevant Twitter users, the stance of the Twitter vaccine discourse and Google trends
- A machine learning model was selected and fine-tuned on Italian Twitter data to help understand the public stance regards vaccines
- We demonstrated that machine learning models achieve close to the same accuracy in categorising tweets as could be expected by a single human annotator.
- Italian institutions should improve their presence online, establish media monitoring activities and engage in conversations in a timely manner
- Conversation on routine vaccinations should be maintained also during emergencies
- Systematic media monitoring + periodic traditional surveys could improve vaccine acceptance through data-based communication

LE SFIDE DELLE TECNOLOGIE DIGITALI PER LA SALUTE DEL FUTURO

http://www.opbg.cloud/#/

Special thanks to the JAV group and to **Dr Francesco Gesualdo**