# Comparing Emotion and Sentiment Analysis Tools on Italian anti-vaccination for COVID-19 posts

Elena Bellodi[1], Alessandro Bertagnon[2] and Marco Gavanelli[1]

[1]*Dipartimento di Ingegneria, University of Ferrara, Via Saragat 1, Ferrara, Italy*
[2]*Dipartimento in Scienze dell'Ambiente e della Prevenzione, University of Ferrara, C.so Ercole I D'Este, 32, Ferrara, Italy*

Abstract
Since the beginning of the vaccination campaign against Covid-19 in our country, resistance to vaccination has emerged on the part of a not negligible portion of the Italian population. Emotions (such as sadness, fear, etc.) and the polarity (positive / negative) of an opinion published on social media are essential for analyzing people's position towards a topic. For this reason, we applied two Natural Language Processing tools, FEEL-IT and SentIta, to a few thousands of social networks posts against the COVID-19 vaccine or specifically the booster shot. We find out some significant insights about the prevalent emotions among users and propose to combine the outputs of the tools in order to increase the classification performance of an opinion according to three possible sentiments (positive/neutral/negative).

**Keywords**
Emotion Recognition, Sentiment Analysis, COVID-19 vaccine

## 1. Introduction

Since the beginning of the vaccination campaign against Covid-19 in our country, resistance to vaccination has emerged on the part of a not negligible portion of the Italian population. A recent research in which the University of Ferrara participated showed that 60 % of the sample expressed some degree of hesitation towards coronavirus vaccines, which depends on the perception of the risk represented by the disease [2]. The Italian media - but not only - gave voice to the doubts of many citizens, but also of some experts, on the efficacy and safety of some vaccines, such as the one developed by the English AstraZeneca[1,2]. The phenomenon of "vaccine hesitancy" is well known to health institutions. The World Health Organization included this problem among the ten greatest threats to global health as early as 2019, in the pre-Covid era[3]. Various researches have shown that the group of opposers was formed only in small part by the so-called no-vax, while the majority were people hesitant for different and more nuanced reasons. Vaccine hesitancy does not depend exclusively on well-established and

*NL4AI 2022: Sixth Workshop on Natural Language for Artificial Intelligence, November 30, 2022, Udine, Italy [1]*
elena.bellodi@unife.it (E. Bellodi); alessandro.bertagnon@unife.it (A. Bertagnon); marco.gavanelli@unife.it (M. Gavanelli)
0000-0002-3717-3779 (E. Bellodi); 0000-0003-2390-0629 (A. Bertagnon); 0000-0001-7433-5899 (M. Gavanelli)

[1]https://www.who.int/news/item/17-03-2021-who-statement-on-astrazeneca-covid-19-vaccine-safety-signals
[2]https://wintoncentre.maths.cam.ac.uk/coronavirus/using-italian-data-illustrate-potential-harms-and-benefits-astrazeneca-vaccine
[3]https://www.who.int/news-room/spotlight/ten-threats-to-global-health-in-2019

clearly identifiable, albeit incorrect, reasons, as is the case for vaccine refusal, which is based on cognitive bias, disinformation and conspiracy theories. The phenomenon of hesitancy instead depends on more nuanced opinions and arguments that have a certain individual or group variability. Emotions (such as sadness, fear, etc.) and the polarity (positive / negative) of an opinion published on social media are essential for analyzing people's position towards a topic. The purpose of the article is to understand the emotions of that part of the Italian population that has resistance to vaccinating against the SARS-CoV-2 virus.

Despite the huge interest of the Natural Language Processing community, the majority of benchmark datasets have been proposed for English [3], [4], [5] showing a limited interest for other languages such as Italian [6], [7], [8], [9].

This paper applies the FEEL-IT [8] and SentIta [9] libraries to a few thousands anti-vaccination COVID-19 posts downloaded from Telegram, Facebook and Twitter between the end of 2021 and the beginning of 2022. The former performs an *emotion recognition* task on Italian texts by annotating every post with one out of four basic emotions: anger ('rabbia'), fear ('paura'), joy ('gioia'), sadness ('tristezza'). The latter performs sentiment analysis on Italian texts by applying a couple of polarity scores ranging between 0 and 1 to each post, indicating both positive and negative sentiment in the sentence. In order to test the performance of the tools we manually labelled a subset of the collected data and computed several machine learning performance metrics and statistics. Firstly, we show that by properly combining the output of the two tools we can get higher performance than using the systems alone, and a clearer and declarative indication of the polarity of an opinion, instead of relying on real-valued scores. Secondly, as regards the specific emotion in the collected opinions, results show that anger is the most spread emotion. Anger is mostly due to political aversion to deeds and decrees issued by the Italian government, especially related to the booster shot and the so-called "green pass"[4,5]. Following anger we find fear, to be understood as fear for adverse events caused both by the first shot and the booster shot.

The paper is organized as follows: Section 2 introduces related work, Section 3 explains the methodology adopted for data collection and manual annotation, Section 4 describes the application of FEEL-IT and SentIta to the data and Section 5 concludes the paper.

## 2. Related Work

FEEL-IT [8] is a benchmark corpus of Italian Twitter posts annotated with four basic emotions: anger, fear, joy, sadness. It was used to fine-tune the UmBERTo model (an Italian BERT model [10]) for the task of emotion recognition. UmBERTo[6] in turn was trained on the Commoncrawl ITA dataset[7], a corpus not related to social media data. This model was released as an open-source Python library called FEEL-IT, so that it is possible to use it for inferring emotions from Italian texts, as done in our experiments. For this reason, in the rest of the paper we will refer

[4]https://ec.europa.eu/info/live-work-travel-eu/coronavirus-response/safe-covid-19-vaccines-europeans/eu-digital-covid-certificate_en
[5]www.dgc.gov.it
[6]https://github.com/musixmatchresearch/umberto
[7]https://commoncrawl.org/

to the model applied to our data to infer emotions with the name of the corpus on which the model was trained.

SentIta[9] is a tool to perform sentiment analysis on Italian texts based on a Bidirectional LSTM-Convolutional Neural Network with two output signals ranging between 0 and 1, one for positive sentiment detection and one for negative sentiment detection. The model was trained and tested on Sentipolc2016 [11] and ABSITA2018 [12] datasets for a total of 15,000 positively and negatively labelled sentences plus 90,000 Wikipedia sentences automatically labelled as neutral. The two signals can be triggered both by the same input sentence if this contains both positive and negative sentiment (e.g. "The food is very good, but the location isn't nice"), and do not sum up necessarily to 1. As in the case of FEEL-IT, the model was released as an open-source Python library.

Considering works on the Italian language and about vaccination campaigns, Tavoschi et al. [13] developed an opinion mining system to monitor the Twitter posts; the system was targeted on vaccine hesitancy (in general) in Italy in the period from September 2016 to August 2017 (before the Covid pandemic). The authors manually labelled 693 training tweets into three categories (against vaccination, in favor or neutral) and trained several machine learning classification models; the best performing (according to a 10-fold cross validation analysis) was based on a Support Vector Machine and reached an average accuracy of 64.8%. Furini [14] performed a word frequency-based analysis to categorize posts by several dimensions (affective, biological, medical and social) distinguishing ProVax and NoVax posts from 2015 to 2017 (before the availability of Covid vaccines); concerning the affective class, he considers four categories, namely anxiety, anger, danger and rage. Gori et al. [15] labelled about 7000 tweets in Italian as pro-vax, no-vax or neutral to the Covid vaccines; their work could be used to develop sentiment analysis tools targeted to study the sentiment about the Covid vaccines.

## 3. Data Collection and Annotation

We retrieved opinions about COVID-19 anti-vaccination in general and against the vaccine third shot by monitoring Telegram, Facebook and Twitter social networks in different time intervals. The Telegram groups monitored were 'Io Non Mi Vaccino Chat' (I won't get vaccinated), 'Vittime vaccino Covid in Italia' (Italian Covid vaccine victims), 'Singles italiani NON vaccinati' (Italian not vaccinated singles), 'No Vaccini Covid sui Bambini' (No Covid vaccine for kids), 'COMBATTENTI NO BOOSTER - NO TERZA DOSE - NO VAC - NO GREEN PASS' (Fighters against the booster shot and green pass), 'Personale Scuola - No Green Pass - No Booster Vax' (School personnel against booster shot and green pass) between August 20th 2021 to February 27th 2022, a particularly relevant period in the evolution of the pandemic situation[8]. Posts were downloaded in JSON format, converted and grouped in a single CSV file. These groups collect opinions from people fully against the COVID-19 vaccine or against/hesitant to the vaccine booster shots.

From the Facebook group 'NON FARÒ LA TERZA DOSE' (I will not get the booster shot) posts were monitored and downloaded from December 8th 2021 to February 2nd 2022, as this group

[8]https://www.epicentro.iss.it/en/coronavirus/sars-cov-2-integrated-surveillance-data

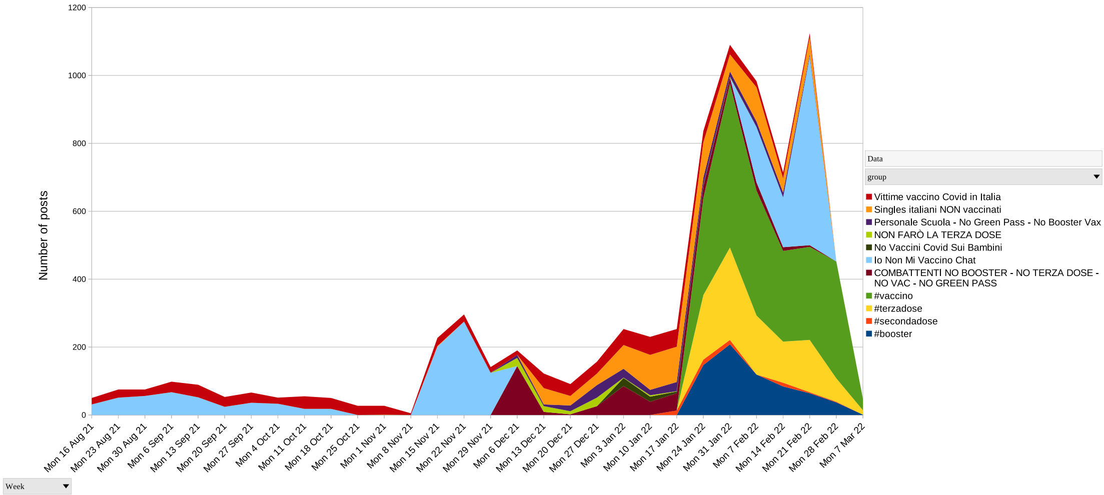


**Figure 1:** Number of posts per social network group per week (total 7928).

became active later than the Telegram groups. For the download the exportcomments.com tool was used, which allows one to export social media comments in CSV format.

Finally, we retrieved Twitter data using the Twitter API and restricting the search to the following Italian hashtags: #vaccino, #secondadose, #terzadose, #booster. Every day, twice a day, 100 tweets were downloaded between January 26th 2022 and March 7th 2022.

All texts were preprocessed by removing duplicates and retweets; moreover, Telegram and Facebook comments were checked to find if the initial substrings 'vac', 'terz', 'dos' were present, which are the roots of the Italian words *vaccino, vaccinazione, terza, dose* (terza dose' is the booster shot in English). This was done in order to check that they were indeed related to the vaccine, while for Twitter we trusted the presence of the hashtag.

Eventually, the number of collected Telegram posts was 4077, the number of Facebook posts was 84 and the number of tweets was 3767, for a total of 7928 texts so distributed in the groups/hashtags: #vaccino (2056), 'Io Non Mi Vaccino Chat' (1852), #terzadose (993), 'Vittime vaccino Covid in Italia' (780), 'Singles italiani NON vaccinati' (738), #booster (660), 'COMBATTENTI NO BOOSTER - NO TERZA DOSE - NO VAC - NO GREEN PASS' (449), 'Personale Scuola - No Green Pass - No Booster Vax' (211), 'NON FARÒ LA TERZA DOSE' (84), #secondadose (58), 'No Vaccini Covid sui Bambini' (47). Each of them is associated with its timestamp. Overall this search allowed us to collect opinions spanning a lot of weeks, from August 20th 2021 to the beginning of March 2022, and focusing both on the sentiment on the vaccine in general and on the booster shot. From the third week of August to the end of October 2021 we have between 30 and 100 opinions per week. From the second week of November 2021 more texts are available (from different social networks), ranging from 150 to 250 per week, with peaks between the third week of January and the third week of February 2022 (800-1100 posts per week). This is highlighted in Figure 1.

A subset of the posts, 1350, was manually labelled by all authors, who are native Italian

**Table 1**
Statistics of the subset of the dataset manually labelled with FEEL-IT emotions.

| *anger* | *fear* | *joy* | *sadness* | Total |
|---|---|---|---|---|
| 413 | 93 | 203 | 175 | 884 |

**Table 2**
Statistics of the subset of the dataset manually labelled with positive/neutral/negative sentiment.

| Pos. Sentiment | Neut. Sentiment | Neg. Sentiment | Total |
|---|---|---|---|
| 227 | 460 | 663 | 1350 |

**Table 3**
Statistics of the complete dataset labelled by FEEL-IT.

| *anger* | *fear* | *joy* | *sadness* | Total |
|---|---|---|---|---|
| 4795 | 1690 | 587 | 856 | 7928 |

speakers. Labelling was performed twice: 1) the first time we removed comments that did not contain any emotion, ending up with 884 posts: each of them was assigned one of the four emotions handled by FEEL-IT (joy, sadness, fear, anger); 2) the second time we kept all 1350 posts and each of them was assigned one out of three classes (positive, negative, neutral), meaning that a post has either a predominant positive sentiment, or a predominant negative sentiment, or does not express any emotion according to the reader, in order to evaluate SentIta performance (see Section 4.2). Table 1 shows the label distribution after the manual annotation for task 1). Table 2 shows the result of the manual annotation for task 2).

## 4. Experiments

In our experimental evaluation we (i) perform emotion recognition with FEEL-IT and sentiment classification with SentIta, (ii) we produce some statistics from the results, (iii) we compare the automatic labelling with our manual labelling in order to test the performance of these tools. Both FEEL-IT[9] and SentIta[10] are open-source Python libraries. The SentIta model is written in Python 3.6 and is implemented in Keras 2.2.4 with Tensorflow 1.11 backend. The Sentita package contains the model and the necessary pre-processing functions.

### 4.1. Emotion recognition

We first experimented with emotion recognition with the FEEL-IT library applied over the dataset of 7928 posts. Results are shown in Table 3.

On the subset of opinions manually labelled we tested the performance of FEEL-IT using the open source ML library scikit-learn[11]. We considered the 884 posts expressing an emotion and computed the confusion matrix, accuracy, precision, recall and F1-score for multi class

[9]https://github.com/MilaNLProc/feel-it
[10]https://nicgian.github.io/Sentita/
[11]https://scikit-learn.org/

**Table 4**
Confusion matrix, Precision, Recall and F1-score for every FEEL-IT class.

| | | Predicted Class | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | | Joy | Anger | Fear | Sadness | Total | Precision | Recall | F1-score |
| Actual Class | Joy | 46 | 86 | 45 | 26 | 203 | 0.742 | 0.227 | 0.347 |
| | Anger | 10 | 344 | 34 | 25 | 413 | 0.604 | 0.833 | 0.700 |
| | Fear | 3 | 37 | 35 | 18 | 93 | 0.269 | 0.376 | 0.314 |
| | Sadness | 3 | 102 | 16 | 54 | 175 | 0.439 | 0.309 | 0.362 |
| | Total | 62 | 569 | 130 | 123 | 884 | | | |

classification, with the number of classes equal to 4. We obtained an accuracy of 0.542, other results are shown in Table 4.

Results show that joy is the most mistaken emotion, and FEEL-IT performance varies depending on the class. It is quite precise in recognizing true joy and anger emotions in texts, but not the other two. It is able to find correctly positive comments based on anger (recall 83.3%), but not based on the other emotions. Anger is the most recognized emotion in the data both by humans and the tool. The varying results should take into account that the FEEL-IT corpus was from a very different context than the COVID-19 texts, and [8] themselves show a reduction in performance (precision (P) = recall (R) = F1-score (F1) $= 0.56$, accuracy $= 0.69$) when the model is applied to different contexts than the Commoncrawl ITA dataset ($P = 0.72, R = 0.73, F1 = 0.71$, accuracy= 0.82): one of these contexts is precisely represented by 662 tweets about COVID-19. Higher performance reduction in our case could be due to the fact that we consider more testing data (almost 900 opinions).

## 4.2. Sentiment analysis

Secondly, we experimented with sentiment analysis with the SentIta library over the dataset of 7928 posts. As SentIta provides two scores, each ranging in the [0,1] interval, one measuring the positive content in the message and the other the negative, we plotted the distribution of the scores in Figure 2. Table 5 shows some examples of SentIta output.

Most of the posts had very low scores in both dimensions: the highest density is around the origin, and the median was 0.09 for the positive score and 0.16 for the negative. This shows that in the considered dataset many posts were classified as neutral. The average is 0.15 and 0.26, respectively for positive and negative, with higher values for negative, as many posts were anti-vaccination.

In order to evaluate the performance of SentIta on the subset of opinions manually labelled with 3 classes (1350 posts), we decided to discretize the two scores applied by SentIta to each text into 3 distinct classes: *positive*, *neutral* and *negative*. The conversion of the scores into classes was based on the fact that we know, as said above, that a low value of both scores indicates a *neutral* sentiment, a high value of the positive score and a low value of the negative one means a *positive* polarity, finally a high value of the negative score and a low value of the positive one indicates a *negative* polarity in an opinion. This excludes the case in which the scores are similar and high, but here we decided to consider a *neutral* sentiment again, since the positive and

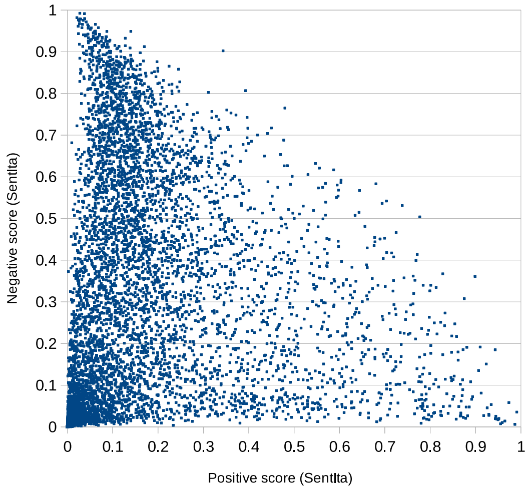


**Figure 2:** Output of SentIta over the complete dataset.

**Table 5**
Examples of the output of SentIta for 3 opinions from our dataset, representative of positive, negative and neutral polarity respectively.

| Text | negative score (oneg) | positive score (opos) |
|---|---|---|
| Ragazzi io vi capisco! Vi voglio bene! Sono con voi! Io niente dosi zero! Vi appoggio e vi capisco!!! Spero veramente vi unirete a me. *Guys I understand you! I love you! I'm with you! Zero shots! I support you and I understand you !!! I truly hope you will join me.* | 0.072429597 | 0.893467069 |
| Non lo voglio fare, nemmeno una di dose, ma mi stanno costringendo. *I don't want to do it, not even a shot, but they are forcing me.* | 0.898216963 | 0.068787813 |
| "Questo studio ha mostrato che l'impatto della vaccinazione sulla trasmissione nella comunità delle varianti circolanti di SARS-CoV-2 non pare essere significativamente diverso dall'impatto tra le persone non vaccinate." *"This study showed that the impact of vaccination on community transmission of circulating variants of SARS-CoV-2 does not appear to be significantly different from the impact among unvaccinated people."* | 0.000597686 | 0.00036788 |

negative sentiment content is somewhat equivalent. More specifically, we generated the rules listed in Equation 1 where the numerical coefficients were chosen to implement the partition as specified previously. A graphical representation of the rules is shown in Figure 3.

$$\begin{cases} \textit{positive} \longleftarrow \texttt{negative\_score} \leq (1-0.05)*\texttt{positive\_score} - 0.1 \\ \textit{negative} \longleftarrow \texttt{negative\_score} \geq (1+0.05)*\texttt{positive\_score} + 0.1 \\ \textit{neutral} \longleftarrow \text{otherwise} \end{cases} \tag{1}$$

After this conversion we could compute accuracy, precision (P), recall (R) and F1-score for multi class classification, with the number of classes equal to 3.

We obtained an accuracy of 0.487, other results are shown in Table 6.

The best sentiment recognition is done about neutral opinions, while positive opinions are often misclassified. In order to improve SentIta and FEEL-IT performance on the manually

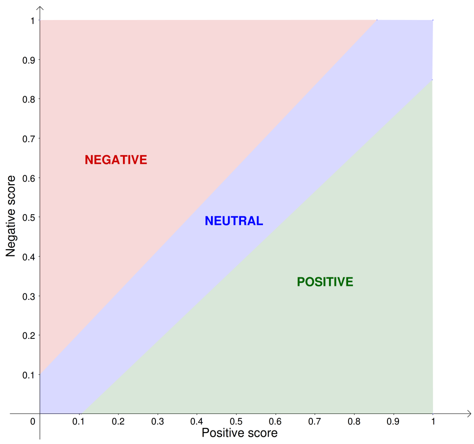


**Figure 3:** Graphical representation of how the scores provided by SentIta are used to obtain 3 discrete classes.

**Table 6**
Confusion matrix, Precision, Recall and F1-score for every SentIta discrete class.

| | | Predicted Class | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | | Pos | Neutral | Neg | Total | Precision | Recall | F1-score |
| Actual Class | Pos | 58 | 79 | 90 | 227 | 0.331 | 0.256 | 0.289 |
| | Neutral | 44 | 285 | 131 | 460 | 0.446 | 0.620 | 0.519 |
| | Neg | 73 | 275 | 315 | 663 | 0.588 | 0.475 | 0.525 |
| | Total | 175 | 639 | 536 | 1350 | | | |

labelled dataset, we tried to combine the two as described in the next section.

## 4.3. Combination of Emotion Recognition and Sentiment Analysis

To compare the results of the application of the two systems on the same data, we performed several tests over the complete dataset (7928 posts).

Firstly, we colored the scatter plot according to the emotion recognition by FEEL-IT; the results are in Figure 4.

This plot confirms that anger, fear and sadness (yellow, red and green points resp.) characterize opinions to which SentIta assigns lower values for the positive score and higher values for the negative one, concentrating near the Y-axis; joy (blue) receives higher values for positive score. Red and blue points (fear and joy) concentrating near zero should represent neutral opinions which received two low scores by SentIta but could not be labelled with a proper emotion by FEEL-IT (as a neutral opinion does not contain an emotion).

Secondly, we interpreted the positive and negative scores provided by SentIta as the coordinates of a point in the plane (in a square of size $1 \times 1$). Note that all the points stand in a circle of radius 1 from the origin, so it might be interesting to identify each point by its polar coordinates $(\rho, \theta)$; in such a representation the angle $\theta \in \left[0, \frac{\pi}{2}\right]$ is a measure of the negativity of the post,

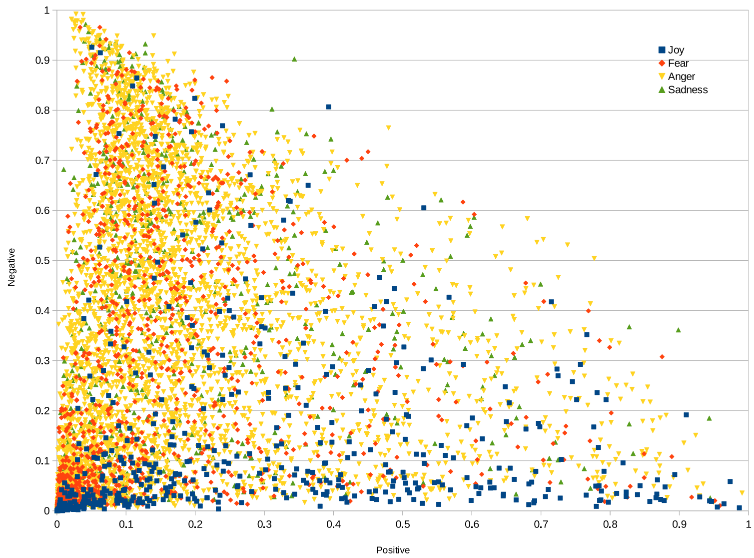


**Figure 4:** Output of SentIta highlighting the emotion of each post.

while the radius $\rho$ can represent the strength with which the argument is pushed forward. Figure 5 is a bubble plot that shows in a synoptic way the evolution of the sentiments for each emotion during time, and can be seen as a comparison of the outputs of the two considered systems. For each week in the considered time range (plotted in the $x$-axis) and for each emotion (each emotion is plotted with a different color) there is a bubble having the $y$-coordinate of the center equal to the average angle $\theta$ that was obtained in that week and for that emotion; the radius of the bubble is proportional to the average strength $\rho$ of the posts. Posts having strength less than 0.1 were removed from the computation of the average angle. The graph shows that joy, the only positive emotion provided by FEEL-IT, is associated to lower angles in SentIta, corresponding to points having a larger positive score and a smaller negative score; the three negative emotions (fear, anger and sadness), are associated to angles closer to $\frac{\pi}{2}$, i.e. closer to the $y$-axis or having a negative score larger than the positive score in SentIta. The plot also highlights that negative emotions (fear, anger and sadness) prevail with respect to joy, confirming that people posting in the monitored groups are against COVID-19 vaccine or the booster shot.

Thirdly, we tried to combine the outputs provided by the two systems by formulating the rule given in Equation 2: $positive_{SentIta}$ and $negative_{SentIta}$ is the conversion of SentIta's scores into discrete classes as described in Section 4.2 and $joy_{\text{FEEL-IT}}$ indicates that FEEL-IT predicted 'joy' as an emotion. These rules maintain the principle of discretization in 3 classes, positive, neutral, and negative.

$$\begin{cases} positive \longleftarrow positive_{SentIta} \text{ and } joy_{\text{FEEL-IT}} \\ negative \longleftarrow negative_{SentIta} \text{ and } \textbf{not } joy_{\text{FEEL-IT}} \\ neutral \longleftarrow \text{otherwise} \end{cases} \tag{2}$$

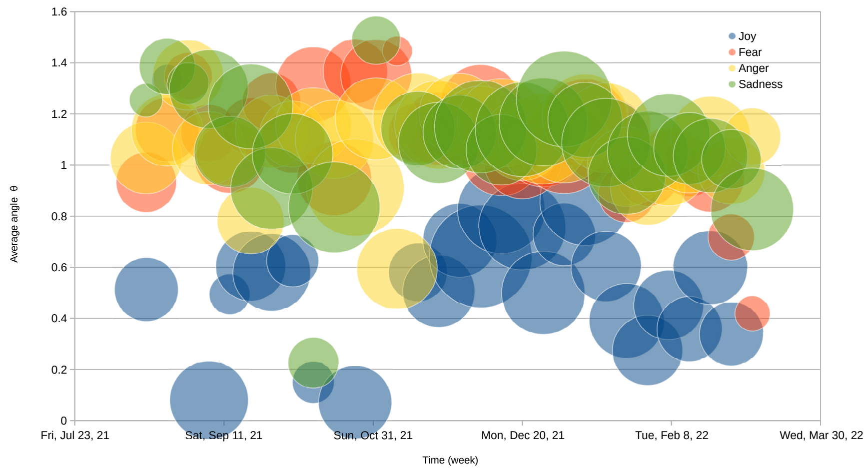


**Figure 5:** Bubble plot of the average strength of the 4 emotions during the period of time monitored.

**Table 7**
Confusion matrix, Precision, Recall and F1-score for every SentIta discrete class.

| | | Predicted Class | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | | Pos | Neutral | Neg | Total | Precision | Recall | F1-score |
| Actual Class | Pos | 57 | 148 | 4 | 209 | 0.781 | 0.272 | 0.404 |
| | Neutral | 16 | 354 | 96 | 466 | 0.415 | 0.760 | 0.536 |
| | Neg | 0 | 352 | 323 | 675 | 0.764 | 0.479 | 0.588 |
| | Total | 73 | 854 | 423 | 1350 | | | |

The idea is to classify as neutral those posts in which the two systems do not agree, i.e. when 'joy' is identified as an emotion (by FEEL-IT) but a negative polarity is found by SentIta, or vice versa when a positive polarity is identified together with one of the "negative emotions" (sadness, anger, fear).

After the conversion of the scores we computed accuracy, precision, recall and F1-score for multi class classification, with the number of classes equal to 3. We obtained an accuracy of 0.544, other results are shown in Table 7.

With respect to the performance obtained at the end of Section 4.2, precision increases significantly; a minor improvement can be seen also in recall. About 50% of the negatives are classified as neutral, and a lot of positives are classified as neutral (70.8%). In general, the combined system misclassifies with the "adjacent class" instead of the opposite class. This third experiment demonstrates that a simple yet effective combination of the outputs of the two NLP systems allows to increase the classification performance of SentIta alone, and to improve the understanding of its results by the user, by means of the discretization of the scores (real values) in 3 more intuitive classes by means of a simple conversion.

## 5. Conclusions and Future Work

We applied two pre-trained neural network-based models for natural language processing, FEEL-IT and SentIta, to COVID-19 anti-vaccination posts downloaded from the major social networks between the end of 2021 and the beginning of 2022. We evaluated the performance of the two models separately on a subset of manually labelled data and then of the two models jointly by combining their outputs. We extracted both an insight of the most prevalent emotions during several months of the pandemic and proposed a method for increasing the performance of the two systems alone through a combination of their output.

In the future this work could benefit both from improvements to the experimental activity and from the application of new techniques, by:

- deepening the analysis of the training sets of the two systems used, making it clear the differences with the texts analyzed for this work. For example, by computing the different average length of sentences, the percentage of lexical overlap, etc.
- when considering the combination of the systems, analyzing in what percentage the two systems are in agreement on the test set (both in cases of correct and incorrect classification) and how many times they are not (comparing the accuracy of the two systems in cases of disagreement);
- applying a different interpretation to SentIta scores when both similar and high: it would be interesting to see in how many of those cases the systems guess or fail with respect to the manual labels;
- applying Deep Learning techniques to the complete dataset, for instance for identifying classes of opinions by unsupervised learning in order to avoid manual labelling.

## Acknowledgments

This work was partially supported by the project "NO Esitazione ('NOE'). Per una comunicazione efficace della vaccinazione anti COVID-19" funded by Fondo per l'Incentivazione alla Ricerca 2021 (FIR) of Ferrara University. Special thanks to Dr. Davide Civolani for his help in manually labelling the dataset.

## References

[1] D. Nozza, L. Passaro, M. Polignano, Preface to the Sixth Workshop on Natural Language for Artificial Intelligence (NL4AI), in: D. Nozza, L. C. Passaro, M. Polignano (Eds.), Proceedings of the Sixth Workshop on Natural Language for Artificial Intelligence (NL4AI 2022) co-located with 21th International Conference of the Italian Association for Artificial Intelligence (AI*IA 2022), November 30, 2022, CEUR-WS.org, 2022.

[2] M. Caserotti, P. Girardi, E. Rubaltelli, A. Tasso, L. Lotto, T. Gavaruzzi, Associations of COVID-19 risk perception with vaccine hesitancy over time for Italian residents, Social Science and Medicine 272 (2021) 113688. URL: https://www.sciencedirect.com/science/article/pii/S0277953621000204. doi:https://doi.org/10.1016/j.socscimed.2021.113688.

[3] F. Calefato, F. Lanubile, N. Novielli, Emotxt: A toolkit for emotion recognition from text, in: 2017 Seventh International Conference on Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW), 2017, pp. 79–80. doi:10.1109/ACIIW.2017.8272591.
[4] M. Abdul-Mageed, L. Ungar, EmoNet: Fine-grained emotion detection with gated recurrent neural networks, in: Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Vancouver, Canada, 2017, pp. 718–728. URL: https://aclanthology.org/P17-1067. doi:10.18653/v1/P17-1067.
[5] M. S. Akhtar, D. Chauhan, D. Ghosal, S. Poria, A. Ekbal, P. Bhattacharyya, Multi-task learning for multi-modal emotion recognition and sentiment analysis, in: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), Association for Computational Linguistics, Minneapolis, Minnesota, 2019, pp. 370–379. URL: https://aclanthology.org/N19-1034. doi:10.18653/v1/N19-1034.
[6] F. Barbieri, V. Basile, D. Croce, M. Nissim, N. Novielli, V. Patti, Overview of the evalita 2016 sentiment polarity classification task, in: CLiC-it/EVALITA, 2016.
[7] R. Sprugnoli, MultiEmotions-it: A new dataset for opinion polarity and emotion analysis for Italian, in: Proceedings of the Seventh Italian Conference on Computational Linguistics (CLiC-it 2020), 2020, pp. 402–408. 7th Italian Conference on Computational Linguistics, CLiC-it 2020 ; Conference date: 01-03-2021 Through 03-03-2021.
[8] F. Bianchi, D. Nozza, D. Hovy, FEEL-IT: emotion and sentiment classification for the italian language, in: O. D. Clercq, A. Balahur, J. Sedoc, V. Barrière, S. Tafreshi, S. Buechel, V. Hoste (Eds.), Proceedings of the Eleventh Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis, WASSA@EACL 2021, Online, April 19, 2021, Association for Computational Linguistics, 2021, pp. 76–83. URL: https://aclanthology.org/2021.wassa-1.8/.
[9] G. Nicola, Bidirectional attentional LSTM for aspect based sentiment analysis on italian, in: T. Caselli, N. Novielli, V. Patti, P. Rosso (Eds.), Proceedings of the Sixth Evaluation Campaign of Natural Language Processing and Speech Tools for Italian. Final Workshop (EVALITA 2018) co-located with the Fifth Italian Conference on Computational Linguistics (CLiC-it 2018), Turin, Italy, December 12-13, 2018, volume 2263 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2018. URL: http://ceur-ws.org/Vol-2263/paper016.pdf.
[10] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, BERT: Pre-training of deep bidirectional transformers for language understanding, in: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), Association for Computational Linguistics, Minneapolis, Minnesota, 2019, pp. 4171–4186. URL: https://aclanthology.org/N19-1423. doi:10.18653/v1/N19-1423.
[11] F. Barbieri, V. Basile, D. Croce, M. Nissim, N. Novielli, V. Patti, Overview of the evalita 2016 sentiment polarity classification task, in: P. Basile, A. Corazza, F. Cutugno, S. Montemagni, M. Nissim, V. Patti, G. Semeraro, R. Sprugnoli (Eds.), Proceedings of Third Italian Conference on Computational Linguistics (CLiC-it 2016) & Fifth Evaluation Campaign of Natural Language Processing and Speech Tools for Italian. Final Workshop (EVALITA 2016), Napoli, Italy, December 5-7, 2016, volume 1749 of *CEUR Workshop Proceedings*,

CEUR-WS.org, 2016. URL: http://ceur-ws.org/Vol-1749/paper_026.pdf.

[12] P. Basile, V. Basile, D. Croce, M. Polignano, Overview of the EVALITA 2018 aspect-based sentiment analysis task (ABSITA), in: T. Caselli, N. Novielli, V. Patti, P. Rosso (Eds.), Proceedings of the Sixth Evaluation Campaign of Natural Language Processing and Speech Tools for Italian. Final Workshop (EVALITA 2018) co-located with the Fifth Italian Conference on Computational Linguistics (CLiC-it 2018), Turin, Italy, December 12-13, 2018, volume 2263 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2018. URL: http://ceur-ws.org/Vol-2263/paper003.pdf.

[13] L. Tavoschi, F. Quattrone, E. D'Andrea, P. Ducange, M. Vabanesi, F. Marcelloni, P. L. Lopalco, Twitter as a sentinel tool to monitor public opinion on vaccination: an opinion mining analysis from september 2016 to august 2017 in Italy, Human Vaccines & Immunotherapeutics 16 (2020) 1062–1069. URL: https://doi.org/10.1080/21645515.2020.1714311. doi:10.1080/21645515.2020.1714311. arXiv:https://doi.org/10.1080/21645515.2020.1714311, pMID: 32118519.

[14] M. Furini, Identifying the features of ProVax and NoVax groups from social media conversations, Comput. Hum. Behav. 120 (2021) 106751. URL: https://doi.org/10.1016/j.chb.2021.106751. doi:10.1016/j.chb.2021.106751.

[15] D. Gori, F. Durazzi, M. Montalti, Z. Di Valerio, C. Reno, M. P. Fantini, D. Remondini, Mis-tweeting communication: a vaccine hesitancy analysis among twitter users in italy, Acta Bio Medica: Atenei Parmensis 92 (2021).