

VaccinItaly: Monitoring Italian Conversations Around Vaccines On Twitter And Facebook

Francesco Pierri^{1,*}, Andrea Tocchetti¹, Lorenzo Corti¹, Marco Di Giovanni^{1,2},
Silvio Pavanetto¹, Marco Brambilla¹, Stefano Ceri¹

¹ Dipartimento di Elettronica, Informazione e Bioingegneria
Politecnico di Milano, Milano, Italy

² Università di Bologna, Bologna, Italy

*Corresponding author. E-mail: francesco.pierri@polimi.it

Abstract

We present VaccinItaly, a project which monitors Italian on-line conversations around vaccines, on Twitter and Facebook. We describe the ongoing data collection, which follows the SARS-CoV-2 vaccination campaign roll-out in Italy and we provide public access to the data collected. We show results from a preliminary analysis of the spread of low- and high-credibility news shared alongside vaccine-related conversations on both social media platforms. We also investigate the content of most popular YouTube videos and encounter several cases of harmful and misleading content about vaccines. Finally, we geolocate Twitter users who discuss vaccines and correlate their activity with open data statistics on vaccine uptake. We make up-to-date results available to the public through an interactive online dashboard associated with the project. The goal of our project is to gain further understanding of the interplay between the public discourse on online social media and the dynamics of vaccine uptake in the real world.

Introduction

On January 30th, 2020, the World Health Organization declared the outbreak of a novel coronavirus (SARS-CoV-2) a global pandemic¹. A year later, the spread of the virus has caused over 121 *M* confirmed cases and more than 2.5 *M* fatalities globally². Italy, in particular, has been one of the first European countries to be hit by the virus, with over 3.28 *M* confirmed cases and 100 *k* fatalities as of March 2021, and the first country outside China to implement national lockdown to circumvent its spreading with severe social and economic consequences (Bonaccorsi et al. 2020; Spelta et al. 2020). Despite the global crisis, we witnessed the most rapid vaccine development for a pandemic in history when the Pfizer-Biontech vaccine showed a 95% efficacy and was approved in several countries³ in late Fall,

2020. In the next few months, several other vaccines were going to be approved and made available to the public⁴. Italy, specifically, has started its vaccination campaign on December 27th, 2020, and reached over 6 *M* dispensed doses⁵ as of March 13th, 2021.

As COVID-19 was spreading around the world, online social networks experienced a so-called "infodemic", i.e. an over-abundance of information about the ongoing pandemic, which yield severe repercussions on public health and safety (Zarocostas 2020; Yang et al. 2021; Gallotti et al. 2020; Guarino et al. 2021). It is believed that low-credibility information might drive vaccine hesitancy and make it hard to reach herd immunity (Yang et al. 2021; Pierri et al. 2021). The European Social Observatory for Disinformation and Social Media Analysis has recently identified four macro-categories of unreliable information about COVID-19 vaccines⁶: (1) there haven't been enough tests on vaccines to guarantee their safety; (2) causal association for individuals who died after being vaccinated; (3) there are further medical complications due to vaccines; (4) vaccines can modify our DNA.

Since the 2016 US presidential elections, the research community has mostly focused its attention on political disinformation and election-related manipulation of online conversations (Lazer et al. 2018; Shao et al. 2018; Ferrara et al. 2016; Pierri and Ceri 2019). However, much concern has grown around health-related misinformation which became manifest during recent measles outbreaks (Filia et al. 2017) and other epidemics such as H1N1 and Ebola (Chew and Eysenbach 2010; Fung et al. 2014), eroding public trust in governments and institutions and undermining public countermeasures during such crises (Scotti et al. 2020; de Zarate et al. 2020).

In this paper, we describe VaccinItaly, a project to monitor Italian conversations around vaccines on multiple social media (Twitter, Facebook) with the aim of understanding the interplay between online public discourse and the vaccine roll-out campaign in Italy. Using a set of Italian vaccine-related

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹[https://www.who.int/publications/m/item/covid-19-public-health-emergency-of-international-concern-\(pheic\)-global-research-and-innovation-forum](https://www.who.int/publications/m/item/covid-19-public-health-emergency-of-international-concern-(pheic)-global-research-and-innovation-forum)

²<https://covid19.who.int>

³<https://www.pfizer.com/news/press-release/press-release-detail/pfizer-and-biontech-conclude-phase-3-study-covid-19-vaccine>

⁴<https://www.nytimes.com/interactive/2020/science/coronavirus-vaccine-tracker.html>

⁵<http://www.salute.gov.it/portale/nuovocoronavirus/dettaglioContenutiNuovoCoronavirus>

⁶<https://www.disinfobservatory.org/disinformation-about-covid-19-and-vaccines-a-journey-across-europe/>



Figure 1: Screenshot of the online dashboard associated to our project. Users can navigate through several sections, each providing different kind of analyses.

keywords, regularly updated to capture trending hashtags and relevant events, as of March 13th, 2021 we collected over 3 M tweets and 1 M Facebook posts published by public pages and groups (we started our collection on December 20th, 2020). We provide public access to the list of keywords and tweet IDs⁷, whereas access to Facebook data is granted by Crowdtangle (CrowdTangle Team 2020) to academics and researchers upon request⁸.

A specific goal of our project is to investigate the spread of reliable and unreliable information related to vaccines. Following a huge corpus of literature (Lazer et al. 2018; Bovet and Makse 2019; Grinberg et al. 2019; DeVerna et al. 2021; Yang et al. 2021), we use a consolidated source-based approach to study how news articles, originated from low- and high-credibility websites (see next sections for more details), are shared alongside vaccine-related conversations on the two platforms. We also highlight YouTube as an additional potential source of misinformation about vaccines. Finally, we geolocate over 1 M users on Twitter and correlate their online activity with open data statistics about the Italian vaccine roll-out campaign⁹.

Up-to-date results from our ongoing analyses are also available to the public through an online dashboard accessible here: <http://genomic.elet.polimi.it/vaccinitaly/>. A pre-

⁷<https://github.com/frapijerri/VaccinItaly>

⁸Nevertheless, we provide a script to replicate our data collection using Crowdtangle keys.

⁹The data are available here: <https://github.com/italia/covid19-opendata-vaccini>

view of the dashboard is available in Figure 1. This is similar in spirit to CoVaxxy¹⁰ (DeVerna et al. 2021), a project based at the Observatory of Social Media (Indiana University) which aims to show the interplay between English-language online misinformation on Twitter and the US vaccine roll-out campaign. However, we focus on the Italian scenario and we also analyze Facebook data.

We believe that our project can contribute to a deeper understanding of the impact of online social networks in an unprecedented scenario where trust in science and governments will be critical to battle a global pandemic.

Related work

There is a huge corpus of literature around the diffusion of health-related (dis)information on online social networks. We describe a few contributions which are related to the Italian context and refer the reader to (Wang et al. 2019) for a deeper review of the existing literature on the subject.

Aquino et al. (2017) explored the relationships between Measles, mumps, and rubella (MMR) vaccination coverage in Italy and online search trends and social network activity from 2010 to 2015. Using a set of keywords related to the controversial link between MMR vaccines and autism, originated from a discredited 1998 paper, authors analyzed Google (search) Trends as well as the activity of Facebook pages and Twitter users on the same subject. They reported a significant negative correlation with the evolution of vaccination coverage in Italy (which decreased from 90% to 85% during the period of observation). They also identified real-world triggering events which most likely drove vaccine hesitancy, i.e. Court of Justice sentences that ruled in favor of a possible link between MMR vaccine and autism.

Donzelli et al. (2018) provide a quantitative analysis of the Italian videos published on YouTube, from 2007 to 2017, about the link between vaccines and autism or other serious side effects in children. They showed that videos with a negative tone were more prevalent and got more views than those with a positive attitude. However, they did not inspect how videos were treating the link between vaccines and autism.

(Righetti 2020) analyzed the Italian vaccine-related environment on Twitter in correspondence with the child vaccination mandatory law promulgated in 2017. Using a keyword-based data collection similar to ours, the author showed that the strong "politicization" of the debate was associated with an increase in the amount of problematic information, such as conspiracy theories, anti-vax narratives, and false news, shared by online users.

Cossard et al. (2020) also analyzed the debate about vaccinations in Italy on Twitter, following the mandatory law promulgated in 2017. They inspected the network of interactions between users, and they identified two main communities of people classified as "vaccine advocates" and "vaccine skeptics", in which they find evidence of echo chamber effects. Besides, they proposed a methodology to predicting the community in which a neutral user would fall, based on

¹⁰<https://osome.iu.edu/tools/covaxxy>

vaccini	vaccinarsi	vaccinerai
vaccino	vaccinare	vaccineremo
vaccinazioni	vacciniamoci	vaccinerete
iononmivaccino	vaccinareh24	iononmivaccinero
vaccinazione	vaccinerò	novaccinoainovax
vaccinocovid	vaccinoanticovid	iononsonounacavia

Table 1: List of keywords used to filter relevant tweets and Facebook posts. They all refer to vaccines and vaccination in general, and some indicate specific pro and anti-vax views (e.g. "iononmivaccino" means "I will not get vaccinated", "vaccinareh24" means "Vaccinate all day long").

a content-based analysis of the tweets shared by users in the two groups.

Data collection

Twitter

Starting on December 20th, 2020, we use Twitter Filter¹¹ API to collect tweets matching the set of keywords in Table 1, in real-time. We routinely check for trending hashtags and relevant events to add new peculiar keywords, e.g. "#no-vaccinoainovax" and "#iononsonounacavia" were hashtags trending on specific days and consequently they were added to the list of keywords. The latter refers to vaccine advocates stating that no-vax should not be vaccinated, and the former indicates vaccine skeptics who "do not want to be guinea pigs for vaccines". The overall data up to March 13th, 2021 comprises approximately 3 *M* tweets shared by 258 *k* unique users.

Facebook

We used the *posts/search* endpoint of the CrowdTangle API (CrowdTangle Team 2020) to collect public posts shared by pages and groups which matched the list of keywords previously defined, resulting in over 10 *M* posts published by over 60 *k* public pages and groups, and re-shared over 100 *M* times, as of March 13th, 2021. In the following, we will use the number of shares to compare Facebook with Twitter.

A limitation to our collection of Facebook is the coverage of pages and groups, whose data can be retrieved using the API. The tool includes over 6M Facebook pages and groups: all those with at least 100k followers/members and a very small subset of verified profiles that can be followed like public pages. Besides, some pages and groups with fewer followers and members can be included by CrowdTangle upon request from users. This might bias the data as, for instance, researchers and journalists might be interested in monitoring pages and groups sharing low-credibility thus leading to an over-representation of such content.

Sources of low- and high-credibility information

We extract URLs contained in tweets and Facebook posts to understand the prevalence of low- and high-credibility information shared in vaccine-based conversations (Yang et al.

¹¹<https://developer.twitter.com/en/docs/twitter-api/v1/tweets/filter-realtime/api-reference/post-statuses-filter>

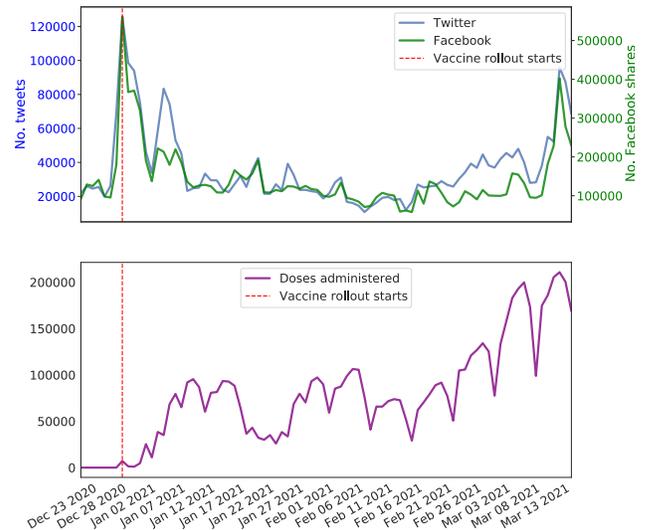


Figure 2: **Top.** Temporal evolution of the daily volume of vaccines-related posts shared on both Twitter and Facebook. We use a dashed red line to indicate the beginning of the Italian vaccination campaign (27th December, 2020). **Bottom.** Total number of vaccine doses administered over time.

2021). We use a consolidated source-based approach to label news articles (Lazer et al. 2018; Pierri, Artoni, and Ceri 2020; Gallotti et al. 2020; Shao et al. 2018; Pierri, Piccardi, and Ceri 2020a; Grinberg et al. 2019; DeVerna et al. 2021; Pierri, Piccardi, and Ceri 2020b; Pierri 2020) depending on the reliability of the source, referring to two lists of Italian low- and high credibility news websites. The former corresponds to websites flagged by Italian fact-checkers for publishing false news, hoaxes and conspiracy theories¹²; the latter corresponds to Italian traditional and most popular news websites (Vicario et al. 2019), and it is used as a reference to understand the prevalence of misleading and (potentially) harmful information. Lists are available in our repository¹³, and we plan to manually augment them during our analyses.

We are aware that this approach, widely adopted in the research community, is not 100% accurate, as cases of misinformation on mainstream websites are not rare and, similarly, low-credibility websites do not publish solely "fake news". However, to date, it is the most reliable and scalable way to study misleading and harmful information. Another limitation to our estimates is that our lists might not fully capture the amount of low- and high-credibility information circulating on Twitter. Besides, we do not consider different typologies of content such as photos, videos, memes, etc.

¹²See www.pagellapolitica.it, www.facta.news and www.butac.it

¹³<https://github.com/frapietri/VaccinItaly>

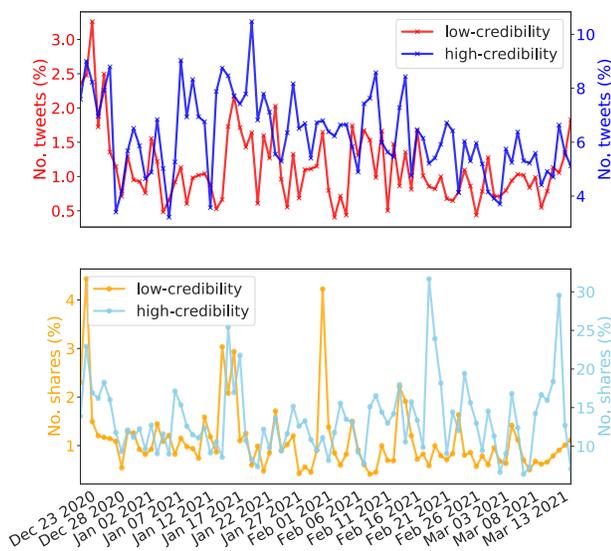


Figure 3: Daily fraction of high-credibility and low-credibility content, compared to online conversations altogether, for Twitter (Top) and Facebook (Bottom).

Online conversations and vaccine roll-out campaign

As previously mentioned, we started our collection on December 20th, 2020, in order to capture the beginning of the Italian vaccination campaign. A symbolic start took place on December 27th 2020, when a few thousand doses of Pfizer–BioNTech COVID-19 vaccine were used to vaccinate part of the medical and health personnel of hospitals, while a few days after 2021 New Year’s eve over 300 *k* doses were delivered to Italy. In this ongoing phase, the priority is given first to health medical and administrative personnel, together with the guests and personnel of nursing homes, and then to elderly people and public service personnel.

Accordingly, we notice a huge spike in both Twitter and Facebook volumes following the symbolic start (over 120 *k* tweets and 500 *k* Facebook posts shared in a single day), and a slightly smaller spike after the actual beginning of the campaign (a peak of 80 *k* tweets and 400 *k* Facebook shares), as shown in Figure 2. As the number of doses administered increases to a steady level, we notice that public attention slowly decreases. However, we notice a second surge of online conversations in March, in correspondence with the suspension of Astrazeneca vaccine in several European countries following an investigation of the European Medicines Agency about unusual blood disorders¹⁴.

Overall we notice that the volume of vaccine-related conversations on Facebook is much higher than on Twitter, and this is probably due to the different size of their user base¹⁵.

¹⁴<https://www.ema.europa.eu/en/news/covid-19-vaccine-astrazeneca-prac-preliminary-view-suggests-no-specific-issue-batch-used-austria>

¹⁵<https://www.statista.com/statistics/787390/main-social-networks-users-italy/>

Prevalence of low- and high-credibility information

A key focus of our research is to analyze the spread of low-credibility information on social media, using high-credibility information as a reference. Overall, we report over 30 *k* tweets and 130 *k* Facebook shares linking to low-credibility news, and over 188 *k* tweets and 1.6 *M* Facebook shares linking to high-credibility news.

In Figure 3, we plot the fraction of tweets and Facebook posts shared that contain a link to either low- or high-credibility information. We note that the amount of low-credibility articles shared on both social media is much smaller compared to high-credibility, on both platforms. Relatively, the mean daily amount of low-credibility information is similar on the two platforms (1.13% on Twitter, 1.10% on Facebook), whereas, interestingly, the mean daily amount of high-credibility circulating on Facebook is higher compared to Twitter (6.27% on Twitter, 13.41% on Facebook). Given the limitations of our analysis, we might not simply state that the information spreading on Facebook is more reliable than on Twitter. Besides, the amount of low-credibility information is non-negligible on both platforms, and it might play a relevant role in shaping the public discourse and opinion around vaccines.

In Figure 4, we show a leaderboard of the top-20 news sources shared on Facebook and Twitter, considering both low- and high-credibility information. We also add the totality of low-credibility information. As previously noted, Facebook shares are an order of magnitude larger than Twitter. Besides, high-credibility domains are shared more than low-credibility websites on both platforms. Except for “liberoquotidiano.it”, a right-wing news website which notably publishes misleading information, we notice the same two most shared low-credibility domains in the leaderboard, namely “imolaoggi.it” and “byoblu.it”. The former is a well-known far-right-wing website that regularly publishes false news with nationalist and anti-immigration views, the latter is a blog that has been repetitiously flagged for sharing hoaxes about health-related subjects, including the COVID-19 pandemic.

By looking at the overall amount of low-credibility news shared on both platforms, we notice interestingly that on Twitter this is larger than any individual high-credibility source. We do not observe the same for Facebook, where still total low-credibility is comparable to the top-3 high-credibility domains. This shows that even though trustworthy information is more prevalent on both social media, the amount of disinformation content shared is still remarkable.

Finally, we investigate the relative popularity of low-credibility news websites on the two platforms, by computing Spearman’s correlation coefficient of the websites ranked by their volumes. We find a significant positive correlation for low-credibility websites ($R= 0.65$, $p\text{-value}= 1.14e-05$), indicating that the majority of unreliable sources is popular on both platforms.

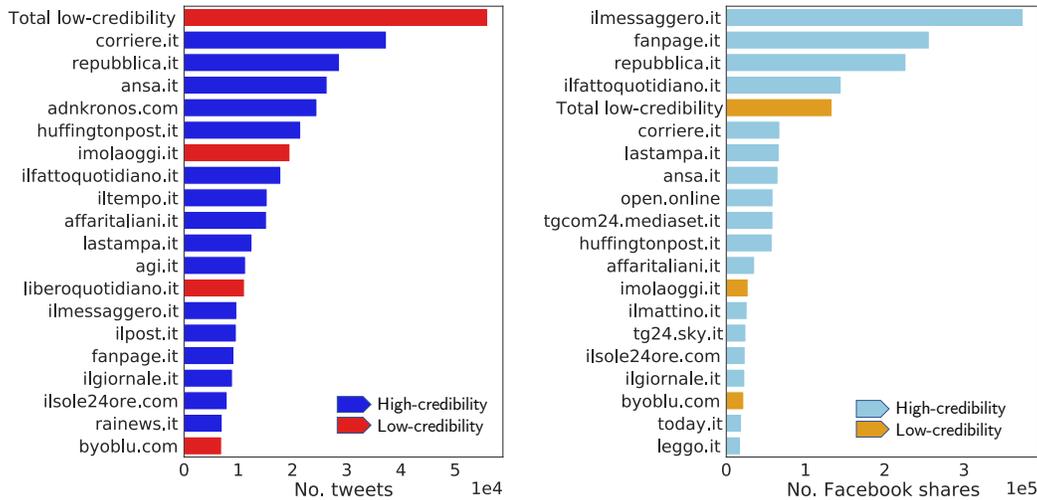


Figure 4: Top-20 news w.r.t the overall number of Twitter (left) and Facebook (right) shares. We also indicate the total amount of low-credibility content and compare it with individual sources of high-credibility information.

YouTube as a potential source of misinformation

As an additional source of information about vaccines, we consider links to YouTube videos shared alongside Facebook and Twitter posts. Previous work (Donzelli et al. 2018) has shown that YouTube is used by both vaccine advocates and skeptics, and we aim to investigate the quality of videos shared on the two platforms. Overall, our dataset contains over 6 *k* links to YouTube shared 21,407 times on Twitter and 132,553 on Facebook.

After extracting URLs pointing to YouTube from tweets and Facebook posts, we use their IDs to query the Youtube API and collect metadata available for such videos. We collected data for approximately 3 *k* videos (published by 1.6 *k* unique channels) shared on Twitter, and 3.2 *k* videos (published by 1.5 *k* unique channels) shared on Facebook. For approximately 300 videos (50 of which were present on both platforms) the API did not return any results, meaning these videos were removed from YouTube due to copyright or policy infringement. Such videos were shared over 800 times on Twitter and 6.5 *k* times on Facebook. Following (Yang et al. 2021), we argue that these videos might have contained suspicious and harmful content. However, we cannot confirm this hypothesis as they were deleted and are no longer available.

We manually inspected the top-20 videos based on the number of tweets and Facebook shares. On Twitter, these videos achieved a total of 5,511 retweets and 5,770,308 YouTube visualizations, while on Facebook they were shared 61,154 times and reached 11,521,158 visualizations. The number of YouTube views was extracted on March 18th. Interestingly, we find several popular videos, on both platforms, which are associated with anti-vax views and misleading information.

A relevant case is the 1st most shared video on Twit-

ter (and 4th on Facebook) with the title "IL PARERE DEL PREMIO NOBEL LUC MONTAGNIER SULLA VACCINAZIONE ANTI-COVID [VIDEO IN ITALIANO]"¹⁶, with over 700 retweets, 4k Facebook shares, and 450k YouTube views. In this video, the Nobel prize winner Luc Montagnier refers to Moderna company as "sorcerer apprentices" stating that they only tested the vaccine on animals, and it's thus not possible to foresee the effects of the vaccine on humans. He also proposes alternative natural treatments against COVID-19 and states that vaccinating the whole population is not the solution.

We omit other examples for reasons of space, but we report that several other popular videos mention conspiracy theories behind the origin of the virus and/or the effects of vaccines as well as proposing alternative therapies and suggesting the audience not to get vaccinated.

The fact that through a simple manual evaluation we encounter almost a dozen of suspicious and, in some cases, explicitly harmful videos among most popular videos indicates that further investigation is required. Indeed, it appears that YouTube is a potential source of online misinformation about vaccines.

Geolocating Twitter conversations

A goal of our project is to link online conversations with geographical details on the ongoing vaccination campaign, e.g. the number of doses administered in each Italian region.

To this aim, we attempt to geolocate Twitter users by using a naive string matching algorithm, i.e. checking whether they have a "location" field disclosed in their profile and matching it against a list of Italian municipalities, provinces,

¹⁶Translation: "The opinion of Nobel Prize Winner Luc Montagnier on COVID-19 vaccination". Available at www.youtube.com/watch?v=kHGtn_vnrJ8.

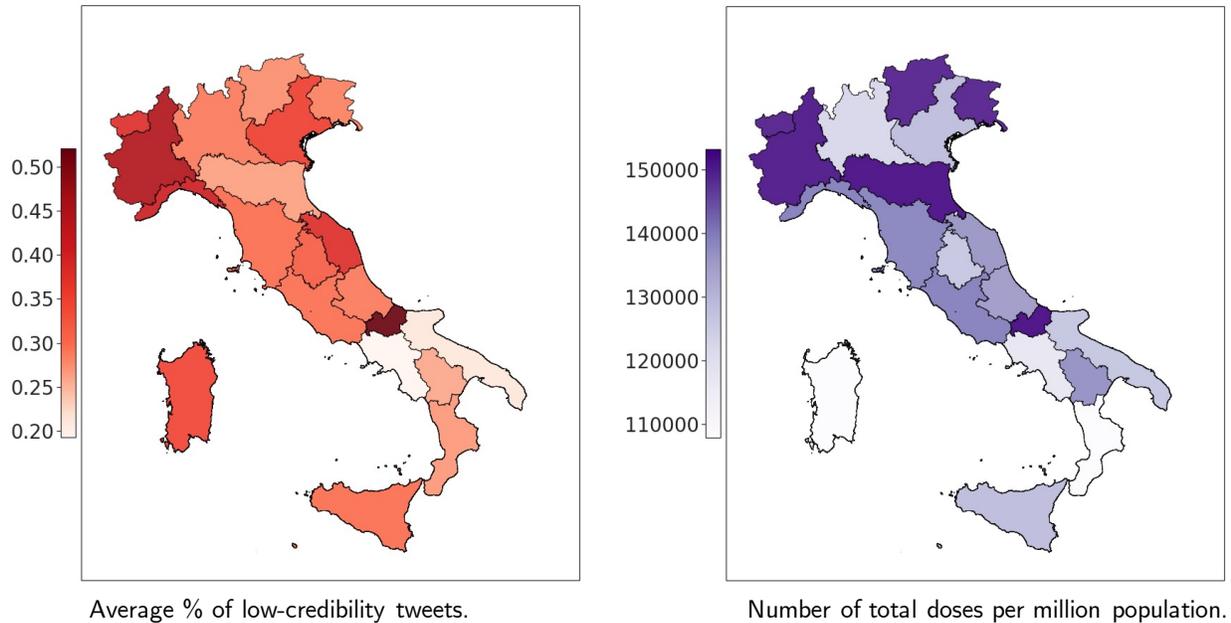


Figure 5: **Left.** Average fraction of low-credibility tweets shared by users, for each Italian region. **Right.** Total amount of vaccine doses administered, per million population, in each Italian region.

and regions¹⁷. In the case of multiple matches, we retain the longest one. We matched circa 16 *k* unique locations and, among over 135 *k* users putting a "location" in their profile, we accordingly geolocated 73 *k* users to either an Italian municipality or region. These shared over 1.3 *M* tweets. The number of accounts mapped to each Italian region is significantly positively correlated with the actual population (Pearson $R=0.89$, $PVAL < 0.001$). However, it is known that the Twitter sample of users might not be fully representative of the Italian population, and this is a limitation to analyses that infer demographics from Twitter (Alessandra, Gentile, and Bianco 2017).

As an illustrative example, we show in Figure 5 statistics on the amount of low-credibility information circulating on Twitter, and the status of the vaccination campaign. Specifically, in the left panel, we show the average fraction of low-credibility tweets shared by users geolocated in each region; darker colors correspond to higher values. We note that on average, Italian users share low-credibility information around 0.20-0.50% of the time. In the right panel, we show the total number of doses administered per million population, in each region. We can note that Lombardy is performing worse than most regions, even though it was the region most struck by the pandemic during the first wave.

These results are still preliminary, as the methodology presents several limitations and needs further assessment, e.g., how to handle multiple locations appearing in the "location" field of user profiles or when false places match with Italian municipalities with misleading names (e.g. "Paese"

which translates as "village").

Conclusions

We present an ongoing project which monitors online conversations of Italian users around vaccines on Twitter and Facebook. We give full access to the data we are collecting, and we provide up-to-date results in an online interactive dashboard. Preliminary analyses show that there is a non-negligible amount of low-credibility information circulating on both platforms, and they indicate YouTube as a potential source of misinformation about vaccines. Our final goal is to understand the interplay between the public discourse on online social media and the vaccine roll-out campaign. In particular, we aim to investigate the impact of online sentiment (e.g. communities of pro and anti-vax) and misinformation about vaccines on vaccine uptake in Italy. We also aim to assess whether there are geographical socio-economic differences that shape both online conversations and the vaccination campaign.

Acknowledgments

This work has been partially supported by the PRIN grant HOPE (FP6, Italian Ministry of Education), and the EU H2020 research and innovation programme, COVID-19 call, under grant agreement No. 101016233 "PERISCOPE" (<https://periscopeproject.eu/>).

References

Alessandra, R.; Gentile, M. M.; and Bianco, D. M. 2017. Who Tweets in Italian? Demographic Characteristics of

¹⁷Taken from the Italian National Institute of Statistics and available at <https://www.istat.it>.

- Twitter Users. In *Convegno della Società Italiana di Statistica*, 329–344. Springer.
- Aquino, F.; Donzelli, G.; De Franco, E.; Privitera, G.; Lopalco, P. L.; and Carducci, A. 2017. The web and public confidence in MMR vaccination in Italy. *Vaccine* 35(35): 4494–4498.
- Bonaccorsi, G.; Pierri, F.; Cinelli, M.; Porcelli, F.; Galeazzi, A.; Flori, A.; Schmidth, A. L.; Valensise, C. M.; Scala, A.; Quattrociochi, W.; and Pammolli, F. 2020. Economic and Social Consequences of Human Mobility Restrictions Under COVID-19. *Proceedings of the National Academy of Sciences*.
- Bovet, A.; and Makse, H. A. 2019. Influence of fake news in Twitter during the 2016 US presidential election. *Nature Communications* 10(1): 7. ISSN 2041-1723. URL <https://doi.org/10.1038/s41467-018-07761-2>.
- Chew, C.; and Eysenbach, G. 2010. Pandemics in the age of Twitter: content analysis of Tweets during the 2009 H1N1 outbreak. *PloS one* 5(11): e14118.
- Cossard, A.; Morales, G. D. F.; Kalimeri, K.; Mejova, Y.; Paolotti, D.; and Starnini, M. 2020. Falling into the echo chamber: the Italian vaccination debate on Twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 14, 130–140.
- CrowdTangle Team. 2020. CrowdTangle. Menlo Park, CA: Facebook. Accessed December 2020.
- de Zarate, J. M. O.; Di Giovanni, M.; Feuerstein, E. Z.; and Brambilla, M. 2020. Measuring Controversy in Social Networks Through NLP. In Boucher, C.; and Thankachan, S. V., eds., *String Processing and Information Retrieval*, 194–209. Cham: Springer International Publishing. ISBN 978-3-030-59212-7.
- DeVerna, M.; Pierri, F.; Truong, B.; Bollenbacher, J.; Axelrod, D.; Loynes, N.; Torres-Lugo, C.; Yang, K.-C.; Menczer, F.; and Bryden, J. 2021. CoVaxxy: A global collection of English Twitter posts about COVID-19 vaccines. *Proceedings of the International AAAI Conference on Web and Social Media*.
- Donzelli, G.; Palomba, G.; Federigi, I.; Aquino, F.; Cioni, L.; Verani, M.; Carducci, A.; and Lopalco, P. 2018. Misinformation on vaccination: a quantitative analysis of YouTube videos. *Human vaccines & immunotherapeutics* 14(7): 1654–1659.
- Ferrara, E.; Varol, O.; Davis, C.; Menczer, F.; and Flammini, A. 2016. The rise of social bots. *Communications of the ACM* 59(7): 96–104.
- Filia, A.; Bella, A.; Del Manso, M.; Baggieri, M.; Magurano, F.; and Rota, M. C. 2017. Ongoing outbreak with well over 4,000 measles cases in Italy from January to end August 2017- what is making elimination so difficult? *Euro-surveillance* 22(37): 30614.
- Fung, I. C.-H.; Tse, Z. T. H.; Cheung, C.-N.; Miu, A. S.; and Fu, K.-W. 2014. Ebola and the social media.
- Gallotti, R.; Valle, F.; Castaldo, N.; Sacco, P.; and De Domenico, M. 2020. Assessing the risks of ‘infodemics’ in response to COVID-19 epidemics. *Nature Human Behaviour* 4: 1285–1293. doi:10.1038/s41562-020-00994-6.
- Grinberg, N.; Joseph, K.; Friedland, L.; Swire-Thompson, B.; and Lazer, D. 2019. Fake news on Twitter during the 2016 U.S. presidential election. *Science* 363(6425): 374–378. ISSN 0036-8075. URL <http://science.sciencemag.org/content/363/6425/374>.
- Guarino, S.; Pierri, F.; Di Giovanni, M.; and Celestini, A. 2021. Information disorders during the COVID-19 infodemic: The case of Italian Facebook. *Online Social Networks and Media* 22: 100124.
- Lazer, D.; Baum, M.; Benkler, Y.; Berinsky, A.; Greenhill, K.; et al. 2018. The science of fake news. *Science* 359(6380): 1094–1096. doi:10.1126/science.aao2998.
- Pierri, F. 2020. The diffusion of mainstream and disinformation news on Twitter: the case of Italy and France. *Companion Proceedings of the Web Conference 2020 (WWW '20 Companion)*.
- Pierri, F.; Artoni, A.; and Ceri, S. 2020. Investigating Italian disinformation spreading on Twitter in the context of 2019 European elections. *PloS one* 15(1): e0227821.
- Pierri, F.; and Ceri, S. 2019. False news on social media: a data-driven survey. *ACM Sigmod Record* 48(2).
- Pierri, F.; Perry, B.; DeVerna, M. R.; Yang, K.-C.; Flammini, A.; Menczer, F.; and Bryden, J. 2021. The impact of online misinformation on US COVID-19 vaccinations. *arXiv preprint arXiv:2104.10635*.
- Pierri, F.; Piccardi, C.; and Ceri, S. 2020a. A multi-layer approach to disinformation detection in US and Italian news spreading on Twitter. *EPJ Data Science* 9(35).
- Pierri, F.; Piccardi, C.; and Ceri, S. 2020b. Topology comparison of Twitter diffusion networks effectively reveals misleading news. *Scientific Reports* 10: 1372.
- Righetti, N. 2020. Health Politicization and Misinformation on Twitter. A Study of the Italian Twittersphere from Before, During and After the Law on Mandatory Vaccinations. doi:10.31219/osf.io/6r95n. URL osf.io/6r95n.
- Scotti, F.; Magnanimi, D.; Urbano, V. M.; and Pierri, F. 2020. Online feelings and sentiments across Italy during pandemic: investigating the influence of socio-economic and epidemiological variables. In *2020 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 453–459. IEEE.
- Shao, C.; Ciampaglia, G. L.; Varol, O.; Yang, K.-C.; Flammini, A.; and Menczer, F. 2018. The spread of low-credibility content by social bots. *Nature Communications* 9: 4787. doi:10.1038/s41467-018-06930-7.
- Spelta, A.; Flori, A.; Pierri, F.; Bonaccorsi, G.; and Pammolli, F. 2020. After the lockdown: simulating mobility, public health and economic recovery scenarios. *Scientific reports* 10(1): 1–13.
- Vicario, M. D.; Quattrociochi, W.; Scala, A.; and Zollo, F. 2019. Polarization and fake news: Early warning of Potential misinformation targets. *ACM Transactions on the Web (TWEB)* 13(2): 10.

Wang, Y.; McKee, M.; Torbica, A.; and Stuckler, D. 2019. Systematic literature review on the spread of health-related misinformation on social media. *Social Science & Medicine* 240: 112552.

Yang, K.-C.; Pierri, F.; Hui, P.-M.; Axelrod, D.; Torres-Lugo, C.; Bryden, J.; and Menczer, F. 2021. The COVID-19 Infodemic: Twitter versus Facebook. *Big Data & Society* Forthcoming in the special issue "Studying the COVID-19 Infodemic at Scale".

Zarocostas, J. 2020. How to fight an infodemic. *The Lancet* 395(10225): 676.