



Article

#lockdown: Network-Enhanced Emotional Profiling in the Time of COVID-19

Massimo Stella ^{1,*}  and Valerio Restocchi ² and Simon De Deyne ³

¹ Complex Science Consulting, 73100 Lecce, Italy

² School of Informatics, University of Edinburgh, Edinburgh EH8 9AB, UK; v.restocchi@ed.ac.uk

³ Human Complex Data Hub, School of Psychological Sciences, University of Melbourne, Melbourne, VIC 3010, Australia; simon.dedeyne@unimelb.edu.au

* Correspondence: massimo.stella@inbox.com

Received: 8 May 2020; Accepted: 6 June 2020; Published: 16 June 2020



Abstract: The COVID-19 pandemic forced countries all over the world to take unprecedented measures, like nationwide lockdowns. To adequately understand the emotional and social repercussions, a large-scale reconstruction of how people perceived these unexpected events is necessary but currently missing. We address this gap through social media by introducing MERCURIAL (Multi-layer Co-occurrence Networks for Emotional Profiling), a framework which exploits linguistic networks of words and hashtags to reconstruct social discourse describing real-world events. We use MERCURIAL to analyse 101,767 tweets from Italy, the first country to react to the COVID-19 threat with a nationwide lockdown. The data were collected between the 11th and 17th March, immediately after the announcement of the Italian lockdown and the WHO declaring COVID-19 a pandemic. Our analysis provides unique insights into the psychological burden of this crisis, focussing on—(i) the Italian official campaign for self-quarantine (*#iorestoacasa*), (ii) national lockdown (*#italylockdown*), and (iii) social denounce (*#sciacalli*). Our exploration unveils the emergence of complex emotional profiles, where anger and fear (towards political debates and socio-economic repercussions) coexisted with trust, solidarity, and hope (related to the institutions and local communities). We discuss our findings in relation to mental well-being issues and coping mechanisms, like instigation to violence, grieving, and solidarity. We argue that our framework represents an innovative thermometer of emotional status, a powerful tool for policy makers to quickly gauge feelings in massive audiences and devise appropriate responses based on cognitive data.

Keywords: COVID-19; social media; hashtag networks; emotional profiling; cognitive science; network science; sentiment analysis; computational social science

1. Introduction

The stunningly quick spread of the COVID-19 pandemic catalysed the attention of worldwide audiences, overwhelming individuals with a deluge of often contrasting content about the severity of the disease, the uncertainty of its transmission mechanisms, and the asperity of the measures taken by most countries to fight it [1–4]. Although these policies have been seen as necessary, they had a tremendous impact on the mental well-being of large populations [5] for a number of reasons. Due to lockdowns, many are facing financial uncertainty, having lost or being on the verge of losing their source of income. Moreover, there is much concern about the disease itself, and most people fear for their own health and that of their loved ones [6], further fueled by *infodemics* [1–3]. Finally, additional distress is caused by the inability of maintaining a normal life [7]. The extent of the impact of these factors is such that, in countries greatly struck by COVID-19 such as China, the population started to develop symptoms of post-traumatic stress disorder [8].

During this time more than ever, people have shared their emotions on social media. These platforms provide an excellent emotional thermometer of the population, and have been widely explored in previous studies investigating how online social dynamics promote or hamper content diffusion [1,2,9–11] and the adoption of specific positive/negative attitudes and behaviours [9,12,13].

1.1. Research Aim

Building on the above evidence, our goal is to draw a comprehensive quantitative picture of people's emotional profiles, emerging during the COVID-19 crisis, through a cognitive analysis of online social discourse. We achieve this by introducing MERCURIAL (Multi-layer Co-occurrence Networks for Emotional Profiling), a framework that combines cognitive network science [14–16] with computational social sciences [2,9,12,17,18]. Before outlining the methods and main contributions of our approach, we briefly review existing research on understanding emotions in social media.

1.2. Past Approaches Bridging Cognitive, Computer and Network Science

Much of the research on emotions in social media has been consolidated into two themes. On the one hand, there is the data science approach, which mostly focused over large-scale positive/negative sentiment detection [9]. A prominent example is the Hedonometer [19], a multi-language tool measuring over time the positive sentiment of Twitter discourse based on word frequency. Recent approaches identified the relevance of tracing more complex affect patterns for understanding social dynamics [16,20–22]. On the other hand, cognitive science research makes use of small-scale analysis tools, but explores the observed phenomena in much more detail in the light of its theoretical foundations [23–25]. Specifically, in cognitive science the massive spread of semantic and emotional information through verbal communication represent long-studied phenomena, known as *cognitive contagion* [25] and *emotional contagion* [24–26], respectively. This research suggests that ideas are composed of a cognitive component and an emotional content, much alike viruses containing the genomic information necessary for their replication [1]. Both these types of contagion happen when an individual is affected in their behaviour by an idea. Emotions elicited by ideas can influence users' behaviour without their awareness, resulting in the emergence of specific behavioural patterns such as implicit biases [24]. Unlike pathogen transmission, no direct contact is necessary for cognitive and emotional contagion to take place, since both are driven by information processing and diffusion, like it happens through social media [27,28]. In particular, during large-scale events, ripples of emotions can rapidly spread across information systems [28] and have dramatic effects, as it has recently been demonstrated in elections and social movements [12,26,29].

At the intersection of data- and cognitive science is emotional profiling, a set of techniques which enables the reconstruction of how concepts are emotionally perceived and assembled in user-generated content [9,15,17–20,30]. Emotional profiling conveys information about basic affective dimensions such how positive/negative or how arousing a message is, and also includes the analysis of more fine-grained emotions such as *fear* or *trust* that might be associated with the lockdown and people's hopes for the future [9,23,31].

Recently, an emerging important line of research has shown that reconstructing the knowledge embedded in messages through social and information network models [14,32,33] successfully highlight important phenomena in a number of contexts, ranging from the diffusion of hate speech during massive voting events [12] to reconstructing personality traits from social media [17]. Importantly, to reconstruct knowledge embedded in tweets, recent work has successfully merged data science and cognitive science, introducing linguistic networks of co-occurrence relationships between words in sentences [16,22,34] and between hashtags in tweets [12]. However, an important shortfall of these works is that these two types of networked knowledge representations were not merged together, thus missing on the important information revealed by studying their interdependence.

1.3. Main Contributions

We identify three important contributions that distinguish our paper from previous literature, and make a further step towards consolidating *cognitive network science* [14] as a paradigm suitable to analyse people's emotions.

First, we introduce a new framework exploiting the interdependence between hashtags and words, addressing the gap previously discussed. This framework, multi-layer co-occurrence networks for emotional profiling (MERCURIAL), combines both the semantic structure encoded through the co-occurrence of hashtags and the textual message to construct a multi-layer lexical network [35]. The multilayer network represented by MERCURIAL is a network of networks, where there is one hashtag co-occurrence network and many word co-occurrence networks. Each hashtag co-occurrence link identifies a word co-occurrence network, indicating language structure in tweets with certain hashtag combinations. This multi-layer network structure allows us to contextualise hashtags and, therefore, improve the analysis of their meaning. Importantly, these networks can be used to identify which concepts or words contribute to different emotions and how central they are.

Second, in contrast to previous work, which largely revolved around English tweets [4,21], the current study focusses on Italian Twitter messages. There are several reasons why the emotional response of Italians is particularly interesting. Specifically, (i) Italy was the first Western country to experience a vast number of COVID-19 clusters; (ii) the Italian government was the first to declare a national lockdown (cf. https://en.wikipedia.org/wiki/COVID-19_pandemic_lockdown_in_Italy, Last Access: 10 June 2020); (iii), the Italian lockdown was announced on 10th March, one day before the World Health Organization (WHO) declared the pandemic status of COVID-19. This enables us to address the urgent need of measuring the emotional perceptions and reactions to social distancing, lockdown, and, more generally, the COVID-19 pandemic.

Third, thanks to MERCURIAL, we obtain richer and more complex emotional profiles that we analyse through the lens of established psychological theories of emotion. This is a fundamental step in going beyond positive/neutral/negative sentiment and to provide accurate insights on the mental well-being of a population. To this end, we take into account three of the most trending hashtags, *#iorestoacasa* (English: "I stay at home"), *#sciacalli* (English: "jackals"), and *#italylockdown*, as representative of positive, negative, and neutral social discourse, respectively. We use these hashtags as a starting point to build multi-layer networks of word and hashtag co-occurrence, from which we derive our profiles. Our results depict a complex map of emotions, suggesting that there is co-existence and polarisation of conflicting emotional states, importantly fear and trust towards the lockdown and social distancing. The combination of these emotions, further explored through semantic network analysis, indicates mournful submission and acceptance towards the lockdown, perceived as a measure for preventing contagion but with negative implications over economy. As further evidence of the complexity of the emotional response to the crisis, we also find strong signals of hope and social bonding, mainly in relation to social flash mobs, and interpreted here as psychological responses to deal with the distress caused by the threat of the pandemic.

1.4. Manuscript Outline

The paper is organised as follows. In the Methods section we describe the data we used to perform our analysis, and describe MERCURIAL in detail. In Section 3 we present the emotional profiles obtained from our data, which are then discussed in more detail in the section Section 4. Finally, the last section highlights the psychological implications of our exploratory investigation and its potential for follow-up monitoring of COVID-19 perceptions in synergy with other datasets/approaches.

We argue that our findings represent an important first step towards monitoring both mental well-being and emotional responses in real time, offering policy-makers a framework to make timely data-informed decisions.

2. Methods

In this section we describe the methodology employed to collect our data and perform the emotional profiling analysis. First, we describe the dataset and how it was retrieved. Then, we introduce co-occurrence networks, and specifically our novel method that combines hashtag co-occurrence with word co-occurrence on multi-layer networks. Finally, we describe the cognitive science framework we used to perform the emotional profiling analysis on the so-obtained networks.

2.1. Data

We gathered 101,767 tweets in Italian to monitor how online users perceived the COVID-19 pandemic and its repercussions in Italy. These tweets were gathered by crawling messages containing three trending hashtags of relevance for the COVID-19 outbreak in Italy and expressing three different sentiment polarities:

- *#iorestoacasa* (English: “I stay at home”), a positive-sentiment hashtag introduced by the Italian Government in order to promote a responsible attitude during the lockdown;
- *#sciacalli*, (English: “jackals”), a negative sentiment hashtag used by online users in order to address unfair behaviour rising during the health emergency;
- *#italylockdown*, a neutral sentiment hashtag indicating the application of lockdown measures all over Italy.

We refer to *#iorestoacasa*, *#sciacalli* and *#italylockdown* as *focal hashtags* to distinguish them from other hashtags. We collected the tweets through *Complex Science Consulting (@ComplexConsult)*, which was authorised by Twitter, and used the *ServiceConnect* crawler implemented in Mathematica 11.3. The collection of tweets comprises 39,943 tweets for *#iorestoacasa*, 26,999 for *#sciacalli* and 34,825 for *#italylockdown*. Retweets of the same text message were not considered. For each tweet, the language was detected. Pictures, links, and non-Italian content was discarded and stop-words (i.e., words without intrinsic meaning such as “di” (English: “of”) and “ma” (English: “but”) removed. Other interesting datasets with tweets about COVID-19 are available in References [21,36].

2.2. Multi-Layer Co-Occurrence Networks

Word co-occurrence networks have been successfully used to characterise a wide variety of phenomena related to language acquisition and processing [16,37,38]. Recently, researchers have also used hashtags to investigate various aspects of social discourse. For instance, Stella et al. [12] showed that hashtag co-occurrence networks were able to characterise important differences in the social discourses promoted by opposing social groups during the Catalan referendum. In this work we introduce MERCURIAL (Multi-layer Co-occurrence Networks for Emotional Profiling), a framework combining:

- Hashtag co-occurrence networks (or hashtag networks) [12]. Nodes represent hashtags and links indicate the co-occurrence of any two nodes in the same tweet.
- Word co-occurrence networks (or word networks) [16]. Nodes represent words and links represent the co-occurrence of any two words one after the other in a tweet without stop-words (i.e., words without an intrinsic meaning).

We combine these two types of networks in a multi-layer network to exploit the interdependence between hashtags and words. This new, resulting network enables us to contextualise hashtags, and capture their real meaning through context, thereby enhancing the accuracy of the emerging emotional profile. To build the multi-layer network, we first build the single hashtag and word layers.

For sake of simplicity, word networks are unweighted and undirected. Robustness checks were performed for various cut-offs in co-occurrence frequency, for example, pruning co-occurrence networks according to co-occurrence weights did not change the most central hashtags or global emotional profiles. Note that the hashtag network was kept at a distinct level from word networks, for example, common words were not explicitly linked with hashtags. As reported in Figure 1,

each co-occurrence link between any two hashtags A and B (*#coronavirus* and *#restiamoacasa* in the figure) is relative to a word network, including all words co-occurring in all tweets featuring hashtags A and B.

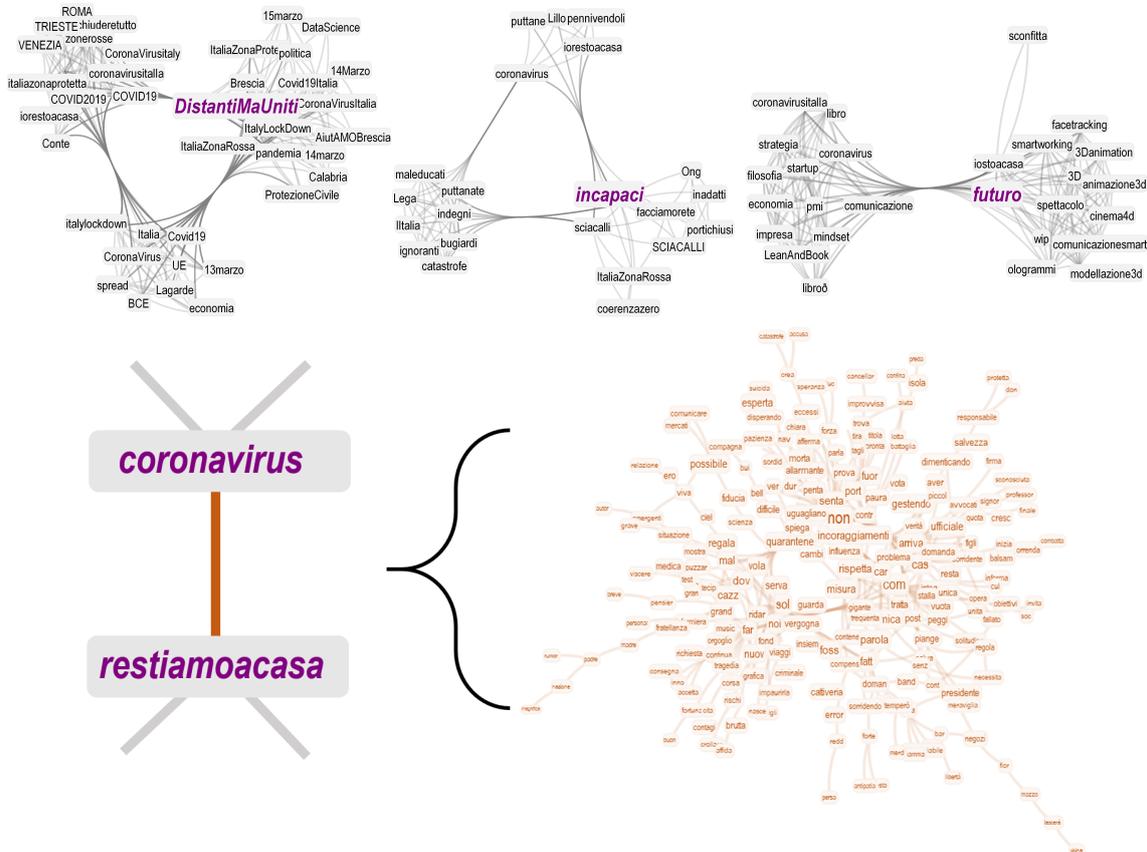


Figure 1. Top: Example of co-occurrence networks for different hashtags: *#distantimauniti* (English: distant but united) in *#iorestoacasa* on the left, *#incapaci* (English: inept) in *#sciacalli* in the middle, and *#futuro* (English: future) in *#italylockdown* on the right. Clusters of co-occurring hashtags were obtained through spectral clustering [39]. These clusters highlight the co-occurrence of European-focused content, featuring hashtags like *#BCE* (i.e., European Central Bank), *#Lagarde* and *#spread* (i.e., spread between Italian and German bonds) together with social distance practices related to *#iorestoacasa*. **Bottom:** In MERCURIAL, any link in a co-occurrence network of hashtags (left) corresponds to a collection of tweets whose words co-occur according to a word network (right). Larger words have a higher closeness centrality.

The hashtag and word networks capture the co-occurrence of lexical entities within the structured online social discourse. Words possess meaning in language [32] and their network assembly is evidently a linguistic network. Similar to words in natural language, hashtags possess linguistic features that express a specific meaning and convey rich affect patterns [12]. The resulting networks capture the meaning of a collection of tweets by identifying which words/hashtags co-occurred together.

This knowledge embedded in hashtag networks was used in order to identify the most relevant or central terms associated within a given collection of thematic tweets. Rather than using frequency to indicate centrality, which makes it difficult to compare hashtags that do not co-occur in the same message, the current work relies on distance-based measures to detect how central a hashtag is in the network.

The first measure that implements this notion is closeness centrality. Closeness $c(i)$ identifies how many links connect i to all its N neighbours and is formalised as follows:

$$c(i) = \frac{N}{\sum_{j=1}^N d_{ij}}, \quad (1)$$

where d_{ij} is the network distance between i and j , that is, the smallest amount of links connecting nodes i and j . In co-occurrence networks, nodes (i.e., hashtags and words) with a higher closeness tend to co-occur more often with each other or with other relevant nodes at short network distance. We expect that rankings of closeness centrality will reveal the most central hashtags in the networks for #iorestocasa, #sciacalli and #italylockdown, in line with previous work in which closeness centrality was used to measure language acquisition and processing [14,33,40].

Importantly, closeness is a more comprehensive approach compared to the simpler frequency analysis. Imagine a collection of hashtags A, B, C, D, \dots . Computing the frequency of hashtag A co-occurring with hashtag B is informative about the frequency of the so-called 2-grams “AB” or “BA” but it does not consider how those hashtags co-occur with C, D , and so forth. In other words, a 2-gram captures the co-occurrence of two specific hashtags within tweets but does not provide the simultaneous structure of co-occurrences of all hashtags across tweets, for which a network of pairwise co-occurrences is required. On such a network, closeness can then highlight hashtags at short distance from all others, that is, co-occurring in a number of contexts in the featured discourse.

In addition to closeness, we also use *graph distance entropy* to measure centrality. This centrality measure captures which hashtags are uniformly closer to all other hashtags in a connected network. Combining closeness with graph distance entropy led to successfully identifying words of relevance in conceptual networks with a few hundreds of nodes [33].

The main idea behind graph distance entropy is that it provides info about the spread of the distribution of network distances between nodes (i.e., shortest path), a statistical quantity that cannot be extracted from closeness (which is, conversely, a mean inverse distance). Considering the set $\mathbf{d}^{(i)} \equiv (d_{i1}, \dots, d_{ij}, \dots, d_{iN})$ of distances between i and any other node j connected to it ($1 \leq j \leq N$) and $M_i = \text{Max}(\mathbf{d}^{(i)})$, then graph distance entropy is defined as:

$$h(i) = -\frac{1}{\log(M_i - 1)} \sum_{k=1}^{M_i-1} p_k^{(i)} \log p_k^{(i)}, \quad (2)$$

where p_k is the probability of finding a distance equal to k . Therefore, $h(i)$ is a Shannon entropy of distances and it ranges between 0 and 1. In general, the lower the entropy, the more a node resembles a star centre [35] and is at equal distances from all other nodes. Thus, nodes with a lower $h(i)$ and a higher closeness are more uniformly close to all other connected nodes in a network. Words with *simultaneously* low graph distance entropy and high closeness were found to be prominent words for early word learning [35] and mindset characterisation [33].

2.3. Attributing Meaning and Emotions to Focal Hashtags by Using Word Networks

In addition to hashtag networks, we also build word networks obtained from a collection of tweets containing any combination of the focal hashtags #iorestocasa or #sciacalli and #coronavirus. For all tweets containing a given set of hashtags, we performed the following:

1. Subdivide the tweet in sentences and delete all stop-words from each sentence, preserving the original ordering of words;
2. Stem all the remaining words, that is, identify the root or stem composing a given word. In a language such as Italian, in which there is a number of ways of adding suffixes to words, word stemming is essential in order to recognise the same word even when it is inflected for different gender, number or as a verb tense. For instance, *abbandoneremo* (we will abandon) and *abbandono* (abandon, abandonment) both represent the same stem *abband*;
3. Draw links between a stemmed word and its subsequent one. Store the resulting edge list of word co-occurrences.

4. Sentences containing a negation (i.e., “not”) underwent an additional step parsing their syntactic structure. This was done in order to identify the target of negation (e.g., in “this is not peace”, the negation refers to “peace”). Syntactic dependencies were not used for network construction but intervened in emotional profiling, instead (see below).

The resulting word network also captures syntactic dependencies between words [16] related by online users to a specific hashtag or combination of hashtag. We used closeness centrality to detect the relevance of words for a given hashtag. Text pre-processing such as word stemming was performed with the R package Snowball C while syntactic dependencies was performed using Mathematica 11.3, which was also used to extract networks and compute network metrics.

The presence of hashtags in word networks provided a way of linking words, which express common language, with hashtags, which express content but also summarise the topic of a tweet. Consequently, by using this new approach, the meaning attributed by users to hashtags can be inferred not only from hashtag co-occurrence but also from word networks. An example of MERCURIAL, featuring hashtag-hashtag and word-word co-occurrences, is reported in Figure 1 (bottom). In this example, hashtags #coronavirus and #restiamoacasa co-occurred together (left) in tweets featuring many co-occurring words (right). The resulting word network shows relevant concepts such as “incoraggiamenti” (English: encouragement) and “problemi” (English: problems), highlighting a positive attitude towards facing problems related to the pandemic. More in general, the attribution and reconstruction of such meaning was explored by considering conceptual relevance and emotional profiling in one or several word networks related to a given region of a hashtag co-occurrence network.

2.4. Emotional Profiling

As a first data source for emotional profiling, this work also used valence and arousal data from Warriner and colleagues [41], whose combination can reconstruct emotional states according to the well-studied circumplex model of affect [31,42]. In psycholinguistics, word valence expresses how positively/negatively a concept is perceived (equivalently to sentiment in computer science). The second dimension, arousal, indicates the alertness or lethargy inspired by a concept. Having a high arousal and valence indicates excitement and joy, whereas a negative valence combined with a high arousal can result in anxiety and alarm [31]. Finally, some studies also include dominance or potency as a measure of the degree of control experienced [41]. However, for reasons of conciseness, we focus on the two primary dimensions of affect: valence and arousal.

Going beyond the standard positive/negative/neutral sentiment intensity is of utmost importance for characterising the overall online perception of massive events [9]. Beyond the primary affective dimension of sentiment, the affect associated with current events [12] can also be described in terms of arousal [43] and of basic emotions such as fear, disgust, anger, trust, joy, surprise, sadness, and anticipation. These emotions represent basic building blocks of many complex emotional states [24], and they are all self-explanatory except for anticipation, which indicates a projection into future events [18]. Whereas fear, disgust, and anger (trust and joy) elicit negative (positive) feedback, surprise, sadness and anticipation have been recently evaluated as neutral emotions, including both positive and negative feedback reactions to events in the external world [44].

To attribute emotions to individual words, we use the National Research Council of Canada (NRC) lexicon [18] and the circumplex model [31]. These two approaches allow us to quantify the *emotional profile* of a set of words related to hashtags or combinations of hashtags. The NRC lexicon enlists words eliciting a given emotion. The circumplex model attributes valence and arousal scores to words, which in turn determine their closest emotional states. Because datasets of similar size were not available for Italian, the data from the NRC lexicon and the Warriner norms were translated from English to Italian using a forward consensus translation of Google Translate, Microsoft Bing and DeepL translator, which was successfully used in previous investigations with Italian [45]. Although the valence of some concepts might change across languages [41], word stemming related several scores to the same stem, for example, scores for “studio” (English: “study”) and “studiare”

(English: “to study”) were averaged together and the average attributed to the stem root “stud”. In this way, even if non-systematic cross-language valence shifting introduced inaccuracy in the score for one word (e.g., “studiare”), averaging over other words relative to the same stem reduced the influence of such inaccuracy. No statistically significant difference ($\alpha = 0.05$) was found between the emotional profiles of 200 Italian tweets, including 896 different stems, and their automatic translations in English, holding for each dimension separately (z-scores < 1.96).

Then, we build emotional profiles by considering the distribution of words eliciting a given emotion/valence/arousal and associated to specific hashtags in tweets. Assertive tweets with no negation were evaluated directly through a bag of words model, that is, by directly considering the words composing them. Tweets including negations underwent an additional intermediate step where words syntactically linked to the negation were substituted with their antonyms [46] and then evaluated. Source-target syntactic dependencies were computed in Mathematica 11.3 and all words targeted by a negation word (i.e., *no*, *non* and *nessuno* in Italian) underwent the substitution with their antonyms.

To determine whether the observed emotional intensity $r(i)$ of a given emotion in a set S of words was compatible with random expectation, we perform a statistical test (Z-test) using the NRC dataset. Remember that emotional intensity here was measured in terms of richness or count of words eliciting a given emotion in a given network. As a null model, we use random samples as follows: let us denote by m the number of words stemmed from S that are also in the NRC dataset. Then, m words from the NRC lexicon are sampled uniformly at random and their emotional profile is compared against that of the empirical sample. We repeated this random sampling 1000 times for each single empirical observed emotional profile $\{r(i)\}_i$. To ensure the resulting profiles are indeed compatible with a Gaussian distribution, we performed a Kolmogorov-Smirnov test ($\alpha = 0.05$). All the tests we performed gave random distributions of emotional intensities compatible with a Gaussian distribution, characterised by a mean random intensity for emotion i , $r^*(i)$ and a standard deviation $\sigma^*(i)$. For each emotion, a z-score was computed:

$$z_i = \frac{r(i) - r^*(i)}{\sigma^*(i)}. \quad (3)$$

In the remainder of the manuscript, every emotional profile incompatible with random expectation was highlighted in black or marked with a check. Since we used a two-tailed Z-test (with a significance level of 0.05), this means that an emotional richness can either be higher or lower than random expectation. Notice that emotional richness and random expectations are different across different emotions, since in the NRC lexicon there are more words eliciting fear and trust and fewer words eliciting other emotions like joy or surprise. For this reason, it is important to consider z-scores for each individual emotion.

3. Results

The investigated corpus of tweets represents a complex multilevel system, where conceptual knowledge and emotional perceptions are entwined on a number of levels. Tweets are made of text and include words, which convey meaning [32]. From the analysis of word networks, we can obtain information on the organisation of knowledge proper of social media users, which is embedded in their generated content [16]. However, tweets also convey meaning through the use of hashtags, which can either refer to specific words or point to the overall topic of the whole tweet. Both words and hashtags can evoke emotions in different contexts, thus giving rise to complex patterns [17]. Similar to words in natural language, the same hashtags can be perceived and used in language differently by different users, according to the context.

The simultaneous presence of word- and hashtag-occurrences in tweets is representative of the knowledge shared by social media users when conveying specific content and ideas.

This interconnected representation of knowledge can be exploited by simultaneously considering both hashtag-level and word-level information, since words specify the meaning attributed to hashtags.

In this section we use MERCURIAL to analyse the data collected. We do so by characterising the hashtag networks, both in terms of meaning and emotional profiles. Precedence is given to hashtags as they not only convey meaning as individual linguistic units but also represent more general-level topics characterising the online discourse. Then, we inter-relate hashtag networks with word networks. Finally, we perform the emotional profiling of hashtags in specific contexts. The combination of word- and hashtag-networks specifies the perceptions embedded by online users around the same entities, for example, coronavirus, in social discourses coming from different contexts.

3.1. Conceptual Relevance in Hashtag Networks

The largest connected components of the three hashtag networks included: 1000 hashtags and 8923 links for *#italylockdown*; 720 hashtags and 5915 links for *#sciaccalli*; 6665 hashtags and 53395 links for *#italylockdown*. All three networks are found to be highly clustered (mean local clustering coefficient [39] of 0.82) and with an average distance between any two hashtags of 2.1. Only 126 hashtags were present in all the three networks.

Table 1 reports the most central hashtags included in each corpus of tweets thematically revolving around *#iorestocasa*, *#sciaccalli* and *#italylockdown*. Two rankings are considered: (i) hashtag frequency in tweets revolving around a key hashtag, and (ii) closeness centrality, which in here quantifies the tendency for hashtags to co-occur with other hashtags expressing analogous concepts and, therefore, are at short network distance from each other (see Section 2). Hashtags with a higher closeness centrality represent the prominent concepts in the social discourse. This result is similar to those showing that closeness centrality captures concepts which are relevant for early word acquisition [40] and production [47] in language. Additional evidence that closeness can capture semantically central concepts is represented by the closeness ranking, which assigns top-ranked positions to *#coronavirus* and *#COVID-19* in all three Twitter corpora. This is a consequence of the corpora being about the COVID-19 outbreak (and of the network metric being able to capture semantic relevance). Frequency and closeness rankings are only partially correlated, indicating that the two metrics highlight different hashtags of relevance in social discourse (Kendall τ s: 0.213, $p < 10^{-6}$, for *#iorestocasa*, 0.234, $p < 10^{-6}$, for *#sciaccalli* and 0.225, $p < 10^{-6}$, for *#italylockdown*).

In the hashtag network built around *#italylockdown*, the most central hashtags are relative to the coronavirus, including a mix of negative hashtags such as *#pandemia* (English: “pandemic”) and positive ones such as *#italystaystrong*. Frequency misses this combination and highlights only positive hashtags. Similarly, the hashtag network built around *#sciaccalli* highlighted both positive (*#facciamorete* (English: “let’s network”) and negative (*#irresponsabili*—English: “irresponsible”) hashtags. However, the social discourse around *#sciaccalli* also featured prominent hashtags from politics, including references to specific Italian politicians, to the Italian Government, and hashtags expressing protest and shame towards the acts of a prominent Italian politician. Last but not least, closeness highlighted as prominent *#mascherine* or face masks, whereas frequency missed that hashtag. The social discourse around *#iorestocasa* included many positive hashtags, eliciting hope for a better future and the need to act responsibly (e.g., *#andratuttobene* - English: “everything will be fine”, or *#restiamoacasa* - English: “let’s stay at home”). The most prominent hashtags in each network (cf. Table 1) indicate the prevalence of a positive social discourse around *#iorestocasa* and the percolation of strong political debate in relation to the negative topics conveyed by *#sciaccalli*. However, we want to extend these punctual observations of negative/positive valences of single hashtags to the overall global networks. To achieve this, we use emotional profiling.

3.2. Emotional Profiling of Hashtag Networks

Hashtags can be composed of individual or multiple words. By extracting individual words from the hashtags of a given network, it is possible to reconstruct the emotional profile of the social

discourse around the focal hashtags *#sciacalli*, *#italylockdown* and *#iorestoacasa*. We tackle this by using the emotion-based [18] and the dimension based [31] emotional profiles (see Section 2).

Table 1. Top-ranked hashtags in co-occurrence networks based on closeness centrality (bottom) and word frequency (top). In all three rankings, the most central hashtag was the one defining the topic (e.g., *#italylockdown*) and was omitted from the ranking.

Frequency Rank	<i>#italylockdown</i>	<i>#sciacalli</i>	<i>#iorestoacasa</i>
1	iorestoacasa	Meloni	italylockdown
2	coronavirus	SanitaPubblica	coronavirus
3	COVID19	Sciacalli	italystaystrong
4	Covid19	Salvini	COVID19
5	IoRestoACasa	coronavirus	iostoacasa
6	andratuttobene	emergenzaCoronavirus	ItaliaZonaRossa
7	restiamoacasa	ItaliaViva	restiamoacasa
8	grazieanomeditutti	Sciacalle	CoronaVirusitaly
9	Coronavirus	governo	Coronavirus
10	coronavirusitalia	COVID19	Conte
11	andratuttobene	Coronavirus	restaacasa
12	flashmob	terroristi	coronavirusit
13	covid19italia	irresponsabili	Covid19
14	litaliachiamò	Lega	celafaremo
15	Iorestoacasa	Berlusconi	Italy
16	coronarvirusitalia	facciamorete	COVID19italia
17	iostoacasa	FreeJulianAssange	italiazonaprotetta
18	coronavirusitalia	Conte	pandemia
19	covid19	CoronaVirus	coronavirusitalia
20	IORESTOACASA	coronavi	COVID2019
Closeness Rank	<i>#italylockdown</i>	<i>#sciacalli</i>	<i>#iorestoacasa</i>
1	coronavirus	coronavirus	iostoacasa
2	COVID19	COVID19	coronavirus
3	ItaliaZonaRossa	Salvini	COVID19
4	iorestoacasa	Lega	andratuttobene
5	Covid19	iorestoacasa	Covid19
6	Italia	Conte	restiamoacasa
7	italystaystrong	facciamorete	quarantena
8	Italy	Governo	COVID2019
9	pandemia	COVID2019	COVID19italia
10	COVID2019	Meloni	iorestoincasa
11	coronavirusitalia	Covid19	coronarvirusitalia
12	italiazonaprotetta	6marzo	italia
13	restiamoacasa	coronavirusitalia	andratuttobene
14	lockdown	mascherine	covid19italia
15	iostoacasa	zonarossa	pandemia
16	coronarvirusitalia	Salvinivergognati	coronavirusitalia
17	coronavirusitalia	Coronavirus	CoronaVirusitaly
18	CoronaVirusitaly	coronavirusitalia	covid19
19	Conte	irresponsabili	Italia
20	COVID19italia	restiamoacasa	coronavirusitalia

The emotional profiles of hashtags featured in co-occurrence networks are reported in Figure 2 (top). The top section of the figure represents perceived valence and arousal represented as a circumplex model of affect [31]. This 2D space or disk is called *emotional circumplex* and its coordinates represent emotional states that are well-supported by empirical behavioural data and brain research [31]. As explained also in the figure caption, each word is endowed with an (x, y) coordinate expressing its perceived valence (x) and arousal (y). Different points indicate different emotional combinations. For instance, $(1,0)$ is the point of maximum/positive valence and zero arousal, that is, calmness; $(0,-1)$ is the point of zero valence and minimum arousal, that is, lethargy; $(-0.6,+0.6)$ represents a point of strong negative valence and positive arousal, that is, alarm.

Figure 2 reports the emotional profiles of all hashtags featured in co-occurrence networks for *#italylockdown* (left), *#sciacalli* (middle) and *#iorestoacasa* (right). To represent the interquartile range of all words for which valence/arousal rating are available, we use a neutrality range.

Histograms falling outside of the neutrality range indicate specific emotional states expressed by words included within hashtags (e.g., #pandemia contains the word “pandemia” with negative valence and high arousal).

3.2.1. Emotional Profiling of Hashtag Networks through the Circumplex Model

In Figure 2 (left, top), the peak of the emotional distribution for hashtags associated with #italylockdown falls within the neutrality range. This finding indicates that hashtags co-occurring with #italylockdown, a neutral hashtag by itself, were also mostly emotionally neutral conceptual entities. Despite this main trend, the distribution also features deviations from the peak mostly in the areas of calmness and tranquillity (positive valence, lower arousal) and excitement (positive valence, higher arousal). Weaker deviations (closer to the neutrality range) were present also in the area of anxiety.

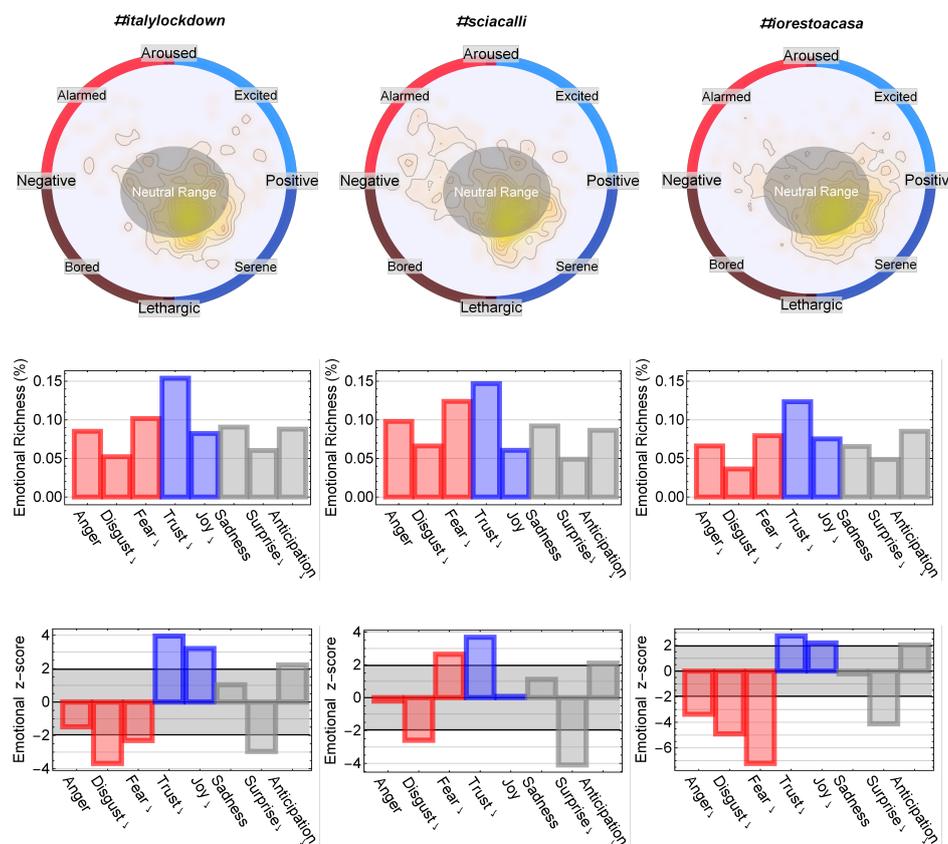


Figure 2. Emotional profiles of all hashtags featured in co-occurrence networks for #italylockdown (left), #sciacalli (middle) and #iorestoacasa (right). **Top:** Circumplex emotional profiling. All hashtags representing one or more words were considered. For each word, valence (x -coordinate) and arousal (y -coordinate) scores were attributed (Section 2) resulting in a 2D density histogram (yellow overlay) relative to the probability of finding an hashtag in a given location in the circumplex, the higher the probability the stronger the colour. Regions with the same probabilities are enclosed in grey lines. A neutrality range indicates where 50% of the words in the underlying valence/arousal dataset would fall and it thus serves as a reference value for detecting abnormal emotional profiles. Distributions falling outside of this range indicate deviations from the median behaviour (i.e., interquartile range, see Section 2). **Middle and Bottom:** NRC-based emotional profiling, detecting how many hashtags inspired a given emotion in a hashtag network. Emotions are colour-coded: red indicates negative emotions while blue and grey indicate positive and neutral emotions, respectively. Results are normalised over the total number of hashtags in a network. Emotions incompatible with random expectation ($|z - score| > 1.96$) were marked with a tick. Z-scores based on random samplings are reported at the bottom, together with the rejection region (grey overlay).

This reconstructed emotional profile indicates that the Italian social discourse featuring *#italylockdown* was mostly calm and quiet, perceiving the lockdown as a positive measure for countering responsibly the COVID-19 outbreak.

Not surprisingly, the social discourse around *#sciacalli* shows a less prominent positive emotional profile, with a higher probability of featuring hashtags eliciting anxiety, negative valence and increased states of arousal, as it can be seen in Figure 2 (center, top). This polarised emotional profile represents quantitative evidence for the coexistence of mildly positive and strongly negative content within the online discourse labelled by *#sciacalli*. This is further evidence that the negative hashtag *#sciacalli* was indeed used by Italian users to denounce or raise alarm over the negative implications of the lockdown, especially in relation to politics and politicians' actions. However, the polarisation of political content and debate over social media platforms has been encountered in many other studies [12,13,22] and cannot be attributed to the COVID-19 outbreak only.

Finally, Figure 2 (top right) shows that positive perception was more prominently reflected in the emotional profile of *#iorestoacasa*, which was the hashtag massively promoted by the Italian Government for supporting the introduction of the nationwide lockdown in Italy. The emotional profile of the 6000 hashtags co-occurring with *#iorestoacasa* indicate a considerably positive and calm perception of domestic confinement, seen as a positive tool to stay safe and healthy. The prominence of hopeful hashtags in association with *#iorestoacasa*, as reported in the previous subsection, indicate that many Italian Twitter users were serene and hopeful about staying at home at the start of lockdown.

3.2.2. Emotional Profiling of Hashtag Networks through Basic Emotions

Hashtag networks were emotionally profiled not only by using the circumplex model (see above) but also by using basic emotional associations taken from the NRC Emotion lexicon (Figure 2, bottom). Across all hashtag networks, we find a statistically significant peak in trust (z -scores > 1.96), analogous of the peaks close to emotions of calmness and serenity an observed in the circumplex models. However, hashtag networks included also negative emotions like fear, which is a natural human response to unknown threats and were observed also with the circumplex representations. All networks featured less disgust eliciting words than random expectation. The intensity of fearful, alarming and angry emotions is stronger in the *#sciacalli* hashtag network, which was used by social users to denounce, complain and express alertness about the consequences of the lockdown.

In addition to the politically-focused jargon highlighted by closeness centrality alone, by combining closeness with graph distance entropy (see Section 2 and Reference [33]) we identify other topics which are uniformly at short distance from others in the social discourse around *#sciacalli*, such as: *#mascherine* (English: "protective masks", which was also ranked high by using closeness only), *#amuchina* (the most popular brand, and synonym of, hand sanitiser), *#supermercati* (English: "supermarkets"). This result suggests an interesting interpretation of the negative emotions around *#sciacalli*. Beside the inflaming political debate and the fear of the health emergency, in fact, a *third* element emerges: Italian twitter users feared and were angry about the raiding and stockpiling of first aid items, symptoms of panic-buying in the wake of the lockdown.

3.3. Assessing Conceptual Relevance and Emotional Profiles of Hashtags via Word Networks

The above comparisons indicate consistency between dimension-based (i.e., the circumplex) and emotion-specific emotional profiling. Since the latter offers also a more precise categorisation of words in emotions, we will focus on emotion-specific profiling. Importantly, to fully understand the emotional profiles outlined above, it is necessary to identify the language expressed in tweets using a given combination of hashtags (see also Figure 1, bottom). As the next step of the MERCURIAL analysis, we gather all tweets featuring the focal hashtags *#italylockdown*, *#sciacalli*, or *#iorestoacasa* and any of their co-occurring hashtags and build the corresponding word networks, as explained in the Methods. Closeness centrality over these networks provided the relevance of each single word in

the social discourse around the topic identified by a hashtag. Only words with closeness higher than the median were reported.

Figure 3 shows the cloud of words appearing in all tweets that include #sciacalli, displayed according to their NRC emotional profile. Similar to the emotional profile extracted from hashtags co-occurring with #sciacalli, the words used in tweets with this hashtag also display a polarised emotional profile with high levels of fear and trust. Thanks to the multi-layer analysis, this dichotomy can now be better understood in terms of the individual concepts eliciting it.

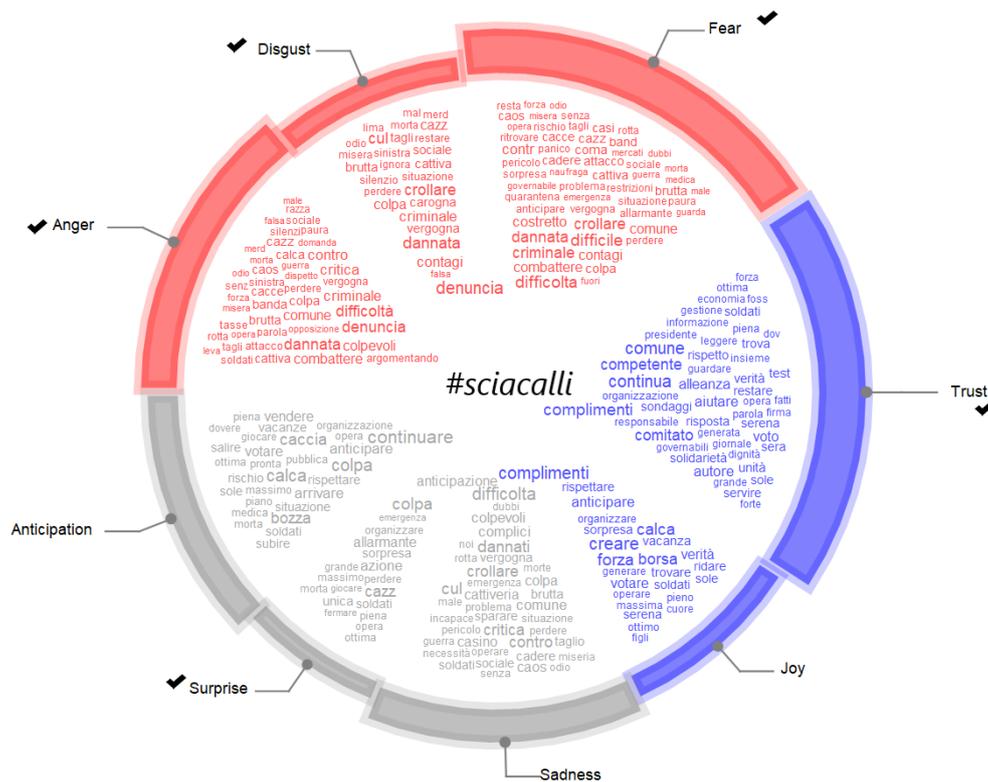


Figure 3. Emotional profile and word cloud of the language used in tweets with #sciacalli. Words are organised according to the emotion they evoke. Font size is larger for words of higher closeness centrality in the word co-occurrence network relative to the hashtag (see Section 2). Every emotion level incompatible with random expectation is highlighted with a check mark. Emotions are colour-coded: red indicates negative emotions, while blue and grey indicate positive and neutral emotions, respectively.

By using closeness on word networks, we identified concepts such as “competente” (English: “competent”), “continua” (English: “continue”, “keep going”), and “comitato” (English: “committee”) to be relevant for the trust-sphere. These words convey trust in the expert committees appointed by the Italian Government to face the pandemic and protect the citizens. We find that other prominent words contributing to make the discourse around #sciacalli trustful are “aiutare” (English: “to help”), “serena” (English: “serene”), “rispetto” (English: “respect”) and “verità” (English: “truth”), which further validate a trustful, open-minded and fair perception of the political and emergency debate outlined above. This perception was mixed with negative elements, mainly eliciting fear but also sadness and anger. The jargon of a political debate emerges in the word cloud of fear: “difficoltà” (English: “difficulty”), “criminale” (English: “criminal”), “dannati” (English: “scoundrel”), “crollare” (English: “to break down”), “banda” (English: “gang”), “panico” (English: “panic”) and “caos” (English: “chaos”). These words indicate that Twitter users felt fear directed to specific targets. A speculative explanation for exorcising fear can be finding

a scapegoat and then target it with anger. The word cloud of such emotion supports the occurrence of such phenomenon by featuring words like “denuncia” (English: “denouncement”), “colpevoli” (English: “guilty”), “vergogna” (English: “shame”), “combattere” (English: “to fight”) and “colpa” (English: “blame”). The above words are reflected also in other emotions like sadness, which features also words like “cadere” (English: “to fall”) and “miseria” (English: “misery”, “out of grace”).

These prominent words in the polarised emotional profile of #sciacalli, suggest that Twitter users feared criminal behaviour, possibly related to unwise political debates or improper stockpiling of supplies (as showed by the hashtag analysis). Our findings also suggest that the reaction to such fearful state, which also projects sadness about negative economic repercussions, was split into a strong, angry denounce of criminal behaviour and messages of trust for the order promoted by competent organisations and committees. It is interesting to note that, according to Ekman’s theory of basic emotions [24], a combination of sadness and fear can be symptomatic of desperation, which is a critical emotional state for people in the midst of a pandemic-induced lockdown.

The same analysis is reported in Figure 4 for the social discourse of #italylockdown (top) and #iorestoacasa (bottom). In agreement with the circumplex profiling, for both #italylockdown and #iorestoacasa the intensity of fear is considerably lower than trust.

However, when investigated in conjunction with words, the overall emotional profile of #italylockdown appears to be more positive, displaying higher trust and joy and lower sadness, than the emotional profile of #iorestoacasa. Although the difference is small, this suggests that hashtags alone are not enough to fully characterise the perception of a conceptual unit, and should always be analysed together with the natural language associated to them.

The trust around #italylockdown comes from concepts like “consigli” (English: “tips”, “advice”), “compagna” (English: “companion”, “partner”), “chiara” (English: “clear”), “abbracci” (English: “hugs”) and “canta” (English: “sing”). These words and the positive emotions they elicit suggest that Italian users reacted to the early stages of the lockdown with a pervasive sense of commonality and companionship, reacting to the pandemic with externalisations of positive outlooks for the future, for example, by playing music on the balconies. This phenomenon was also mimicked in other countries later, and extensively reported by traditional media, see <https://tinyurl.com/balconicovid>, Last Access: 20 April 2020).

Interestingly, this positive perception co-existed with a more complex and nuanced one. Despite the overall positive reaction, in fact, the discourse on #italylockdown also shows fear for the difficult times facing the contagion (“contagi”) and the lockdown restrictions (“restrizioni”), and also anger, identifying the current situation as a fierce battle (“battaglia”) against the virus.

The analysis of anticipation, the emotional state projecting desires and beliefs into the future, shows the emergence of concepts such as “speranza” (English: “hope”), “possibile” (English: “possible”) and “domani” (English: “tomorrow”), suggesting a hopeful attitude towards a better future.

The social discourse around #iorestoacasa brought to light a similar emotional profile, with a slightly higher fear towards being quarantined at home (quarantena (English: “quarantine”), comando (English: “command”, “order”), emergenza (English: “emergency”). Both surprise and sadness were elicited by the the word “confinamento” (English: “confinement”), which was prominently featured in the network structure arising from the tweets we analysed.

In summary, the above emotional profiles of hashtags and words from the 101,767 tweets suggest that Italians reacted to the lockdown measure with:

1. a fearful denunciation of criminal acts with political nuances and sadness/desperation about negative economic repercussions (from #sciacalli);
2. positive and trustful externalisations of fraternity and affect, combined with hopeful attitudes towards a better future (from #italylockdown and #iorestoacasa);
3. a mournful concern about the psychological weight of being confined at home, inspiring sadness and disgust towards the health emergency (from #iorestoacasa).

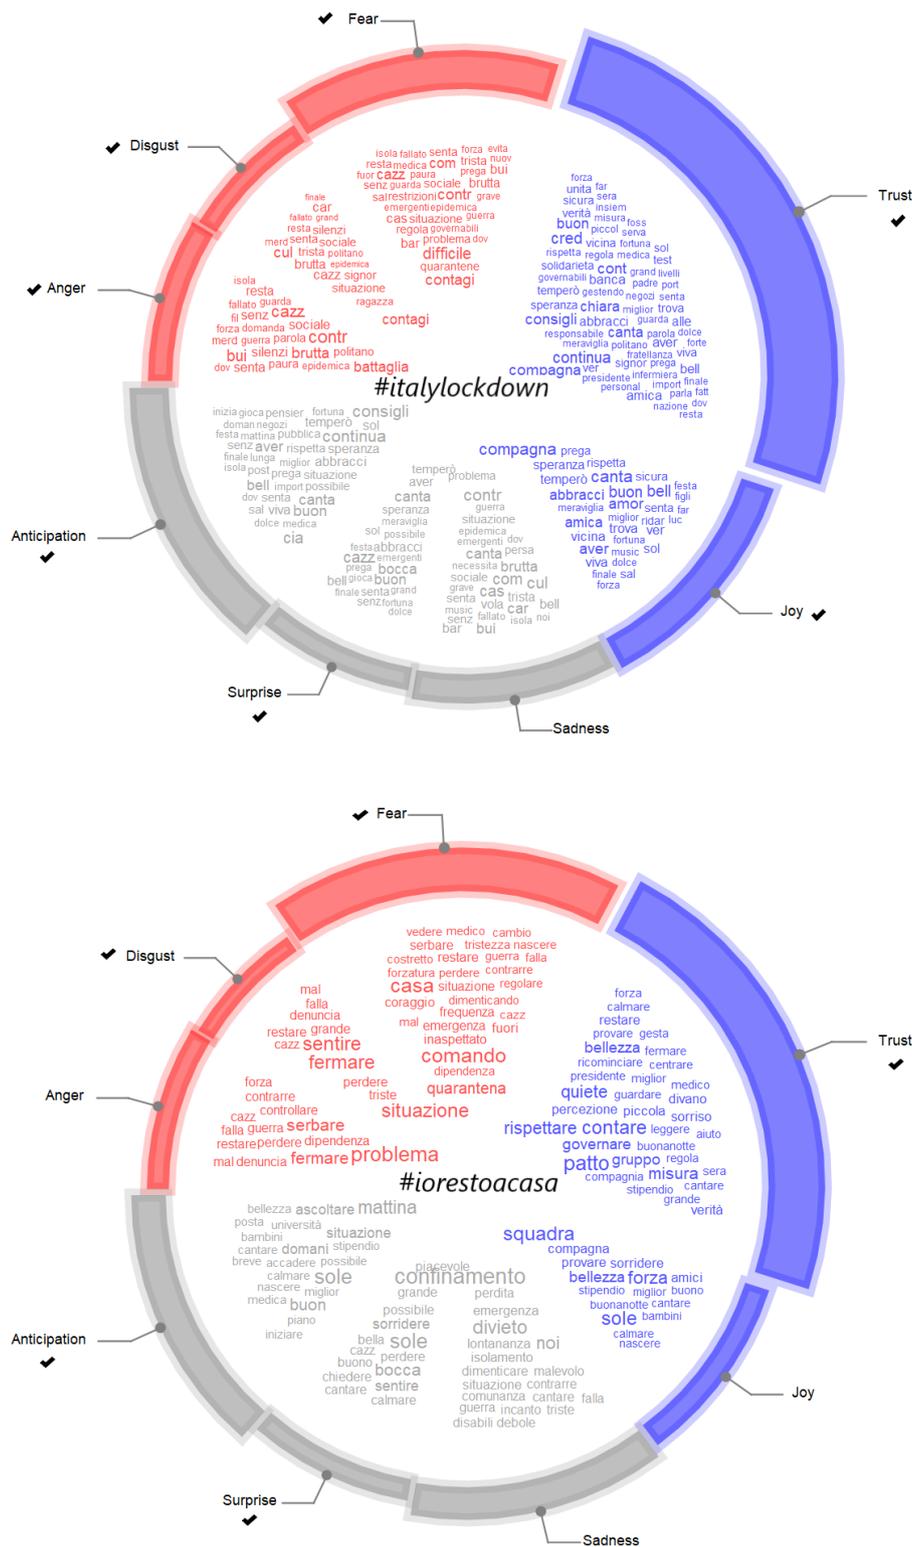


Figure 4. Emotional profile and word cloud of the language used in tweets with **#italylockdown** (top) and **#iorestocasa** (bottom). Words are organised according to the emotion they evoke. Font size is larger for words of higher closeness centrality in the word co-occurrence network relative to the hashtag (Section 2). Every emotional richness incompatible with random expectation is highlighted with a check mark. Emotions are colour-coded: red indicates negative emotions, while blue and grey indicate positive and neutral emotions, respectively.

3.4. Hashtag Co-Occurrence Contextually Influences Hashtag Emotional Profiles

In the previous section we showed our findings on how Italians perceived the early days of lockdown on social media. But what about their perception of the ultimate cause of such lockdown, COVID-19? To better reconstruct the perception of *#coronavirus*, it is necessary to consider the different contexts where this hashtag occurs. Figure 5 displays the reconstruction of the emotional profile of words used in tweets with *#coronavirus* and either *#italylockdown*, *#sciacalli*, or *#iorestoacasa*.

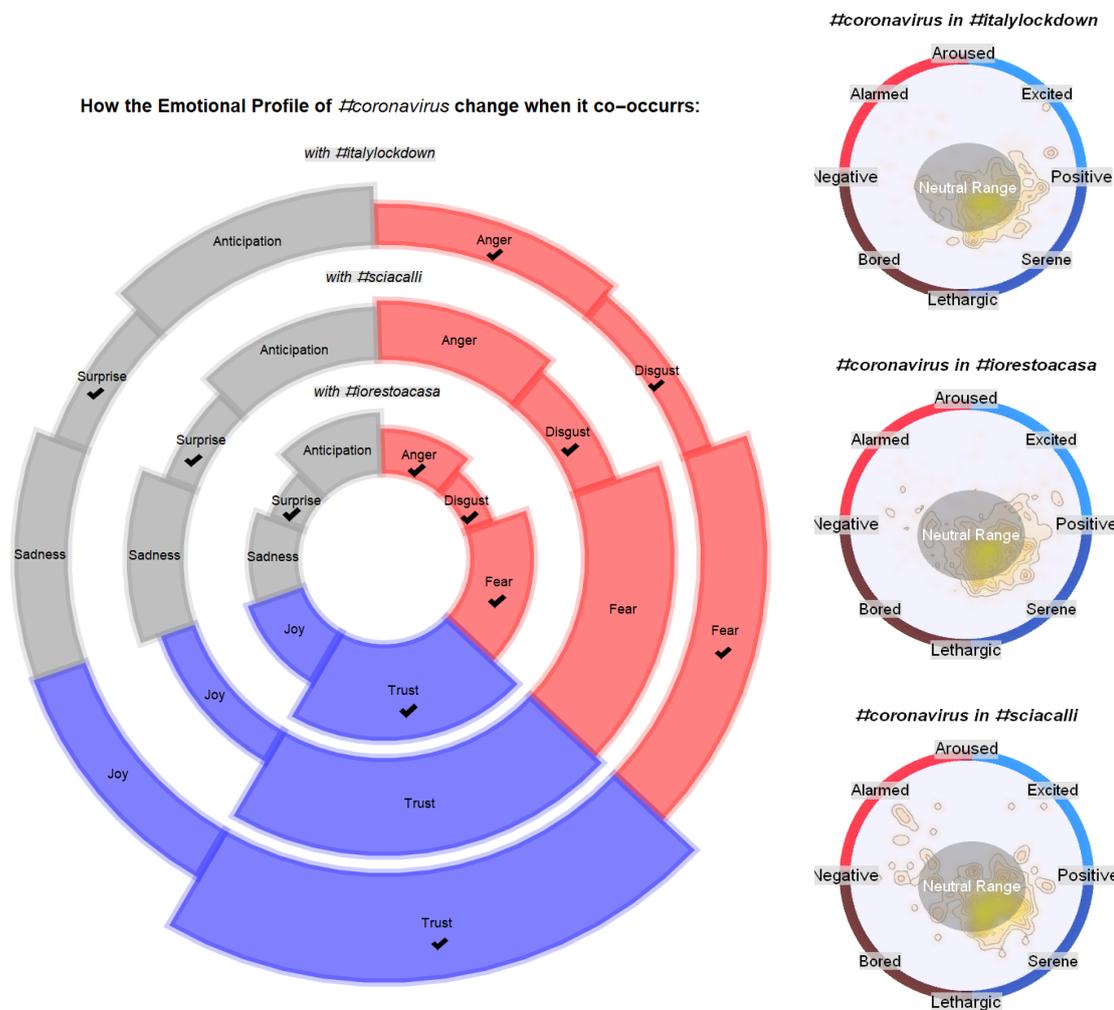


Figure 5. Left: Emotional profiles with the NRC lexicon of the words occurring in tweets with the hashtag *#coronavirus* when it co-occurs with *#italylockdown* (outer circle), *#sciacalli* (middle circle) and *#iorestoacasa* (inner circle). Right: The same emotional profiles detected through the valence/arousal emotional circumplex. Every emotional richness incompatible with random expectation is highlighted with a check mark.

Our results suggest that the emotional profiles of language used in these three categories of tweets are different. For example, when considering tweets including *#sciacalli*, which the previous analysis revealed being influenced by political and social denunces of criminal acts, *#coronavirus* is perceived with a more polarised fear/trust dichotomy.

Although *#coronavirus* was perceived as trustful as random expectations when co-occurring with *#sciacalli* (z-score: 1.69 < 1.96), it was perceived with significantly higher trust when appearing in tweets with *#iorestoacasa* (z-score: 3.05 > 1.96) and *#italylockdown* (z-score: 3.51 > 1.96). To reinforce this picture, the intensity of fear towards *#coronavirus* was statistically significantly lower than

random expectation in the discourse of *#iorestocasa* (z-score: $-2.35 < -1.96$) and *#italylockdown* (z-score: $-3.01 < -1.96$).

This difference is prominently reflected in both the circumplex model (Figure 5, right) and the NRC emotional profile (Figure 5, left), although in the latter both emotional intensities are compatible with random expectation. These quantitative comparisons provide data-driven evidence that Twitter users perceived the same conceptual entity, that is, COVID-19, with a higher trust when associating it to concrete means for hampering pathogen diffusion like lockdown and house confinement, and with a higher fear when denouncing the politics and economics behind the pandemic.

However, social distancing, lockdown and house confinement clearly do not have only positive sides. Rather, as suggested by our analysis, they bear complex emotional profiles, where sadness, anger and fear towards the current situation and future developments have been prominently expressed by Italians on social media.

4. Discussion

This study delved into the massive information flow of Italian social media users in reaction to the declaration of the pandemic status of COVID-19 by WHO, and the announcement of the nationwide lockdown by the Italian Government in the first half of March 2020. We explored the emotional profiles of Italians during this period by analysing the social discourse around the official lockdown hashtag promoted by the Italian Government (*#iorestocasa*), together with a most trending hashtag of social protest (*#sciacalli*), and a general hashtag about the lockdown (*#italylockdown*).

The fundamental premise of this work is that social media opens a window on the minds of millions of people [17]. Monitoring social discourse on online platforms provides unprecedented opportunities for understanding how different categories of people react to real world events [9,12,15].

4.1. Impact of Cognitive Network Science for Social Media Analysis

Here we introduced a new framework, Multi-layer Co-occurrence Networks for Emotional Profiling (MERCURIAL), which is based on cognitive network science and that allowed us to:

- (i) quantitatively structure social discourse as a multi-layer network of hashtag-hashtag and word-word co-occurrences in tweets;
- (ii) identify prominent discourse topics through network metrics backed up by cognitive interpretation [14];
- (iii) reconstruct and cross-validate the emotional profile attributed to each hashtag or topic of conversation through the emotion lexicon and the circumplex model of affect from social psychology and cognitive neuroscience [31].

Our interdisciplinary framework provides a first step in combining network and cognitive science principles to quantify sentiment for specific topics. Our analysis also included extensive robustness checks (e.g., selecting words based on different centrality measures, statistical testing for emotions), further highlighting the potential of the framework.

The analysis of concept network centrality identified hashtags of political denounce and protest against irrational panic buying (e.g., face masks and hand sanitiser) around *#sciacalli* but not in the hashtag networks for *#italylockdown* and *#iorestocasa*. Our results also suggest that the social discourse around *#sciacalli* was further characterised by fear, anger, and trust, whose emotional intensity was significantly stronger than random expectation. We also found that the most prominent concepts eliciting these emotions revolve around social denounce (anger), concern for the collective well-being (fear), and the measures implemented by expert committees and authorities (hope). This interpretation is supported also by Plutchik's wheel of emotions [23], according to which combinations of anger, disgust and anticipation can be symptoms of aggressiveness and contempt. However, within Plutchik's wheel, trust and fear are not in direct opposition.

4.2. Evidence for Emotional Polarisation around COVID-19

The polarisation of positive/negative emotions observed around #sciacalli might be a direct consequence of a polarisation of different social users with heterogeneous beliefs, which is a phenomenon present in many social systems [22] but is also strongly present in social media through the creation of echo chambers enforcing specific narratives and discouraging the discussion of opposing views [2,10,11,13,48].

Emotional polarisation might therefore be a symptom of a severe lack of social consensus across Italian users in the early stages of the lockdown induced by COVID-19. In social psychology, social consensus is a self-built perception that the beliefs, feelings, and actions of others are analogous to one's own [49]. Destabilising this perception can have detrimental effects such as reducing social commitment towards public good or even lead to a distorted perception of society, favouring self-distrust and even conditions such as social anxiety [49]. Instead, acts such as singing from the balconies together can reduce fear and enhance self-trust [43], as well as promote commitment and social bonding [50], which is also an evolutionary response to help coping with a threat, in this case a pandemic, through social consensus. When interpreted under the lens of social psychology, the flash mobs documented by traditional media and identified here as relevant by semantic network analysis for #italylockdown and #iorestoacasa become important means of facing the distress induced by confinement [43,49,50].

4.3. Implications of the Detected Anger and Fear over Self-Awareness and Violence

Anger and fear permeated not only #sciacalli but were found, to a lesser extent, also in association with other hashtags such as #iorestoacasa or #italylockdown. Recent studies (cf. Reference [51]) found that anger and fear can drastically reduce individuals' sense of agency, a subjective experience of being in control of our own actions, linking this behavioural/emotional pattern also to alteration in brain states. In turn, a reduced sense of agency can lead to losing control, potentially committing violent, irrational acts [51]. Consequently, the strong signals of anger and fear detected here represent red flags about a building tension manifested by social users which might contribute to the outbreak of violent acts or end up in serious psychological distress due to lowered self-control.

One of the most direct implications of the detected strong signals of fear, anger and sadness is represented by increased violent behaviour. In cognitive psychology, the General Aggression Model (GAM) [52] is a well-studied model for predicting and understanding violent behaviour as the outcome of a variety of factors, including personality, situational context and the personal internal state of emotion and knowledge. According to GAM, feeling emotions of anger in a situation of confinement can strongly promote violent behaviour. In Italy, the emotions of anger and anxiety we detected through social media are well reflected in the dramatic rise in reported cases of domestic violence. For instance, the anti-violence centers of D.i.Re (*Donne in Rete Contro la Violenza*) reported an anomalous increase of +74.5% in the number of women looking for help for domestic violence in March 2020 in Italy (see the official report in Italian <https://tinyurl.com/direcontrolaviolenza>, Last Access: 20 April 2020). Hence, monitoring social media can be insightful about potential tensions mediated and discussed by large populations, a topic in need for further research and with practical prominent repercussions for fighting COVID-19.

4.4. Contextual Shifts in Emotions around the Novel Coronavirus

As discussed, we found the hashtag #coronavirus to be central across all considered hashtag networks. However, our analysis outlined different emotional nuances of #coronavirus across different networks. In psycholinguistics, contextual valence shifting [53] is a well-known phenomenon whereby the very same conceptual unit can be perceived wildly differently by people according to its context. This phenomenon suggests the importance of considering words in a contextual manner, by comparison to each other, as it was performed in this study, rather than alone. Indeed, contexts can

change the meaning and emotional perception of many words in language. We showed here that the same connotation shifting phenomenon [53] can happen also for hashtags. Online users perceived *#coronavirus* with stronger intensities of trust and lower fear (than random expectation) when using that hashtag in the context of *#iorestoacasa* and *#italylockdown*, but not when associated to *#sciacalli*. This shifting underlines the importance of considering contextual information surrounding a hashtag in order to better interpret its nuanced perception. To this aim, cognitive networks represent a powerful tool, providing quantitative metrics (such as graph distance entropy) that would be otherwise not applicable with mainstream frequency approaches in psycholinguistics.

4.5. Limitations and Future Research

MERCURIAL facilitates a quantitative characterisation of the emotions attributed to hashtags and discourses. Thanks to a better contextualisation of words by using hashtag-word multi-layer networks, MERCURIAL can also be used for all those tasks that are normally undertaken in NLP and revolve around gauging emotions from textual data. Some applications of MERCURIAL could include (but are not limited to) characterising emotional reactions to specific events, anticipate potential unrest or violence, or monitor the political discourse before an election.

Nonetheless, it is important to bear in mind that the analysis we conducted relies on some assumptions and limitations. For instance, following previous work [12], we built unweighted and undirected networks, neglecting information on how many times hashtags co-occurred. Including these weights would be important for detecting communities of hashtags, beyond network centrality. Notice that including weights would come at the cost of not being able to use graph distance entropy, which is defined over unweighted networks and was successfully used here for exposing the denounce of panic buying in *#sciacalli*. Another limitation is relative to the emotional profiling performed with the NRC lexicon, in which the same word can elicit multiple emotions. Since we measured emotional intensity by counting words eliciting a given emotion (plus the negations, see Section 2), a consequence was the repetition of the same words across the sectors of the above word clouds. Building or exploiting additional data about the predominance of a word in a given emotion would enable us to identify words which are peripheral to a given emotion, reduce repetitions and offer even more detailed emotional profiles. Recently, forma mentis networks [30,33] have been introduced as a method to detect the organisation of positive/negative words in the mindsets of different individuals. A similar approach might be followed for emotions in future research. Acting upon specific emotions rather than using the circumplex model would also solve another problem, in that the attribution of arousal to individual words is prone to more noise, even in mega-studies, compared to detecting word valence [54]. Notice also that the current analysis used uniform random sampling in the NRC lexicon for obtaining z-scores. This reference model keeps into account that some emotions might be elicited by fewer/more words than others but it also neglects potential idiosyncratic biases in the Italian twitter discourse. These biases should be detected and cross-validated through massive amounts of Italian twitter messages, beyond the dataset gathered here. The public release of massive multi-language datasets including Italian and reporting the COVID-19 Twitter discourse would enable more refined null models controlling also for idiosyncratic word under- or super-representation.

Another limitation is that emotional profiles might fluctuate over time. The insightful results outlined and discussed here were aggregated over a short time window, thus reducing the impact of aggregation itself. Future analyses on longer time windows should adopt time-series for investigating emotional patterns, addressing key issues like non-stationary tweeting patterns over time and statistical scarcity due to tweet crawling (see also Reference [12]). The current analysis has focused on aggregated tweets, but previous studies have shown both stable individual and intercultural differences in affect [55], especially for dimensions such as arousal. Similarly, some emotions are harder to measure than others, which might affect reliability and thus underestimate their contribution. The current approach estimates emotional profiles on the basis of a large set of words, which will reduce some

language-specific differences. The collection of currently missing large-scale Italian normative datasets for lexical sentiment could further improve the accuracy of the findings.

This study approaches the relation between emotions and mental distress mostly from the perspective that attitudes and emotions of the author are conveyed in the linguistic content. However, the emotion profile might also have implications for readers as well, as recent research suggests that even just reading words of strong valence/arousal can have deep somatic and visceral effects, for example, raising heart beat or promoting involuntary muscle tension [56]. Furthermore, authors and readers participate in an information network, and quantifying which tweets are liked or retweeted depending on the structure of social network can provide further insight on their potential impact [4,10,12,22,57], which calls for future approaches merging social networks, cognitive networks and emotional profiling.

Finally, understanding the impact of nuanced emotional appraisals would also benefit from investigating how these are related to behavioural and societal outcomes including the numbers of the contagion (e.g., hospitalisations, death rate, etc.) and compliance with physical distancing [58].

5. Conclusions

Given the massive attention devoted to the COVID-19 pandemic by social media, monitoring online discourse can offer an insightful thermometer of how individuals discussed and perceived the pandemic and the subsequent lockdown. Our MERCURIAL framework offered quantitative readings of the emotional profiles among Italian twitter users during early COVID-19 diffusion. The detected emotional signals of political and social denounce, the trust in local authorities, the fear and anger towards the health and economic repercussions, and the positive initiatives of fraternity, all outline a rich picture of emotional reactions from Italians. Importantly, the psychological interpretation of MERCURIAL's results identified early signals of mental health distress and antisocial behaviour, both linked to violence and relevant for explaining increments in domestic abuse. Future research will further explore and consolidate the behavioural implications of online cognitive and emotional profiles, relying on the promising significance of our current results. Our cognitive network science approach offers decision-makers the prospect of being able to successfully detect global issues and design timely, data-informed policies. Especially under a crisis, when time constraints and pressure prevent even the richest and most organised governments from fully understanding the implications of their choices, an ethical and accurate monitoring of online discourses and emotional profiles constitutes an incredibly powerful support for facing global threats.

6. Data Availability

The IDs of the tweets analysed in this study are available on the Open Science Foundation repository: <https://osf.io/jy5kz/>.

Author Contributions: Conceptualization, M.S. and V.R.; methodology, M.S. and S.D.D.; validation, M.S., V.R. and S.D.D.; formal analysis, M.S.; investigation, M.S., V.R. and S.D.D.; resources, M.S.; data curation, M.S.; writing—original draft preparation, M.S., V.R. and S.D.D.; writing—review and editing, M.S., V.R. and S.D.D.; visualization, M.S.; supervision, M.S.; project administration, M.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Acknowledgments: M.S. acknowledges Daniele Quercia, Nicola Perra and Andrea Baronchelli for stimulating discussion.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Zarocostas, J. How to fight an infodemic. *Lancet* **2020**, *395*, 676. [[CrossRef](#)]
2. Cinelli, M.; Quattrocioni, W.; Galeazzi, A.; Valensise, C.M.; Brugnoli, E.; Schmidt, A.L.; Zola, P.; Zollo, F.; Scala, A. The covid-19 social media infodemic. *arXiv* **2020**, arXiv:2003.05004.
3. Gallotti, R.; Valle, F.; Castaldo, N.; Sacco, P.; De Domenico, M. Assessing the risks of “infodemics” in response to COVID-19 epidemics. *arXiv* **2020**, arXiv:2004.03997.
4. Pulido, C.M.; Villarejo-Carballido, B.; Redondo-Sama, G.; Gomez, A. COVID-19 infodemic: More retweets for science-based information on coronavirus than for false information. *Int. Soc.* **2020**. [[CrossRef](#)]
5. Wang, C.; Pan, R.; Wan, X.; Tan, Y.; Xu, L.; Ho, C.; Ho, R. Immediate Psychological Responses and Associated Factors during the Initial Stage of the 2019 Coronavirus Disease (COVID-19) Epidemic among the General Population in China. *Int. J. Environ. Res. Public Health* **2020**, *17*, 1729. [[CrossRef](#)]
6. World Health Organization. Mental Health During COVID-19 Outbreak. *Lancet Psychiatry* **2020**, *7*, e15–e16. [[CrossRef](#)]
7. Zhu, S. Wu, Y.; Zhu, C.; Hong, W.; Yu, Z.; Chen, Z.; Wang, Y. The immediate mental health impacts of the COVID-19 pandemic among people with or without quarantine managements. *Brain Behav. Immun.* **2020**, in press. [[CrossRef](#)]
8. Wang, C.; Pan, R.; Wan, X.; Tan, Y.; Xu, L.; McIntyre, R.; Choo, F.; Tran, B.; Ho, R.; Sharma, V.; et al. A longitudinal study on the mental health of general population during the COVID-19 epidemic in China. *Brain Behav. Immun.* **2020**, in press. [[CrossRef](#)]
9. Ferrara, E.; Yang, Z. Quantifying the effect of sentiment on information diffusion in social media. *PeerJ Comput. Sci.* **2015**, *1*, e26. [[CrossRef](#)]
10. Davis, J.T.; Perra, N.; Zhang, Q.; Moreno, Y.; Vespignani, A. Phase transitions in information spreading on structured populations. *Nat. Phys.* **2020**, *16*, 590–596. [[CrossRef](#)]
11. Ciulla, F.; Mocanu, D.; Baronchelli, A.; Gonçalves, B.; Perra, N.; Vespignani, A. Beating the news using social media: The case study of American Idol. *EPJ Data Sci.* **2012**, *1*, 8. [[CrossRef](#)]
12. Stella, M.; Ferrara, E.; De Domenico, M. Bots increase exposure to negative and inflammatory content in online social systems. *Proc. Natl. Acad. Sci. USA* **2018**, *115*, 12435–12440. [[CrossRef](#)]
13. Bail, C.A.; Argyle, L.P.; Brown, T.W.; Bumpus, J.P.; Chen, H.; Hunzaker, M.F.; Lee, J.; Mann, M.; Merhout, F.; Volfovsky, A. Exposure to opposing views on social media can increase political polarization. *Proc. Natl. Acad. Sci. USA* **2018**, *115*, 9216–9221. [[CrossRef](#)] [[PubMed](#)]
14. Siew, C.S.; Wulff, D.U.; Beckage, N.M.; Kenett, Y.N. Cognitive Network Science: A review of research on cognition through the lens of network representations, processes, and dynamics. *Complexity* **2019**, *2019*, 2108423. [[CrossRef](#)]
15. Stella, M. Text-mining forma mentis networks reconstruct public perception of the STEM gender gap in social media. *arXiv* **2020**, arXiv:2003.08835.
16. Amancio, D.R. Probing the topological properties of complex networks modeling short written texts. *PLoS ONE* **2015**, *10*, e0118394. [[CrossRef](#)]
17. Quercia, D.; Kosinski, M.; Stillwell, D.; Crowcroft, J. Our twitter profiles, our selves: Predicting personality with twitter. In *Proceedings of the 2011 IEEE 3rd International Conference on Privacy, Security, Risk and Trust and 2011 IEEE Third International Conference on Social Computing*; IEEE: Piscataway, NJ, USA, 2011; pp. 180–185.
18. Mohammad, S.M.; Turney, P.D. Emotions evoked by common words and phrases: Using mechanical turk to create an emotion lexicon. In *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*; Association for Computational Linguistics: Stroudsburg, PA, USA, 2010; pp. 26–34.
19. Dodds, P.S.; Harris, K.D.; Kloumann, I.M.; Bliss, C.A.; Danforth, C.M. Temporal patterns of happiness and information in a global social network: Hedonometrics and Twitter. *PLoS ONE* **2011**, *6*, e26752. [[CrossRef](#)]
20. Mohammad, S.; Bravo-Marquez, F.; Salameh, M.; Kiritchenko, S. Semeval-2018 task 1: Affect in tweets. In *Proceedings of the 12th International Workshop on Semantic Evaluation*, New Orleans, LA, USA, 5–6 June 2018; pp. 1–17.
21. Kleinberg, B.; van der Vegt, I.; Mozes, M. Measuring Emotions in the COVID-19 Real World Worry Dataset. *arXiv* **2020**, arXiv:2004.04225.

22. Brito, A.C.M.; Silva, F.N.; Amancio, D.R. A complex network approach to political analysis: Application to the Brazilian Chamber of Deputies. *PLoS ONE* **2020**, *15*, e0229928. [[CrossRef](#)]
23. Plutchik, R. *The Emotions*; University Press of America: Lanham, MA, USA, 1991.
24. Ekman, P.E.; Davidson, R.J. *The Nature of Emotion: Fundamental Questions*; Oxford University Press: Oxford, UK, 1994.
25. Hatfield, E.; Cacioppo, J.T.; Rapson, R.L. Emotional contagion. *Curr. Direct. Psychol. Sci.* **1993**, *2*, 96–100. [[CrossRef](#)]
26. Barsade, S.G. The ripple effect: Emotional contagion and its influence on group behavior. *Adm. Sci. Quart.* **2002**, *47*, 644–675. [[CrossRef](#)]
27. Kramer, A.D.; Guillory, J.E.; Hancock, J.T. Experimental evidence of massive-scale emotional contagion through social networks. *Proc. Natl. Acad. Sci. USA* **2014**, *111*, 8788–8790. [[CrossRef](#)] [[PubMed](#)]
28. Frey, S.; Donnay, K.; Helbing, D.; Sumner, R.W.; Bos, M.W. The rippling dynamics of valenced messages in naturalistic youth chat. *Behavi. Res. Methods* **2019**, *51*, 1737–1753. [[CrossRef](#)] [[PubMed](#)]
29. Jasper, J.M. Emotions and social movements: Twenty years of theory and research. *Ann. Rev. Soc.* **2011**, *37*, 285–303. [[CrossRef](#)]
30. Stella, M.; De Nigris, S.; Aloric, A.; Siew, C.S. Forma mentis networks quantify crucial differences in STEM perception between students and experts. *PLoS ONE* **2019**, *14*, e0222870. [[CrossRef](#)]
31. Posner, J.; Russell, J.A.; Peterson, B.S. The circumplex model of affect: An integrative approach to affective neuroscience, cognitive development, and psychopathology. *Dev. Psychopathol.* **2005**, *17*, 715–734. [[CrossRef](#)] [[PubMed](#)]
32. De Deyne, S.; Kenett, Y.N.; Anaki, D.; Faust, M.; Navarro, D.J. Large-scale network representations of semantics in the mental lexicon. In *Big Data in Cognitive Science: From Methods to Insights*; Routledge/Taylor & Francis Group: Abingdon, UK, 2016; pp. 174–202.
33. Stella, M.; Zaytseva, A. Forma mentis networks map how nursing and engineering students enhance their mindsets about innovation and health during professional growth. *PeerJ Comput. Sci.* **2020**, *6*, e255. [[CrossRef](#)]
34. Mehler, A.; Gleim, R.; Gaitsch, R.; Hemati, W.; Uslu, T. From Topic Networks to Distributed Cognitive Maps: Zipfian Topic Universes in the Area of Volunteered Geographic Information. *Complexity* **2020**, *2020*, 4607025. [[CrossRef](#)]
35. Stella, M.; De Domenico, M. Distance entropy cartography characterises centrality in complex networks. *Entropy* **2018**, *20*, 268. [[CrossRef](#)]
36. Chen, E.; Lerman, K.; Ferrara, E. Tracking Social Media Discourse About the COVID-19 Pandemic: Development of a Public Coronavirus Twitter Data Set. *JMIR Pub. Health Surv.* **2020**, *6*, e19273. [[CrossRef](#)]
37. Marinho, V.Q.; Hirst, G.; Amancio, D.R. Labelled network subgraphs reveal stylistic subtleties in written texts. *J. Complex Netw.* **2018**, *6*, 620–638. [[CrossRef](#)]
38. Vankrunkelsven, H.; Verheyen, S.; Storms, G.; De Deyne, S. Predicting lexical norms: A comparison between a word association model and text-based word co-occurrence models. *J. Cogn.* **2018**, *1*, 45. [[CrossRef](#)] [[PubMed](#)]
39. Newman, M. *Networks*; Oxford University Press: Oxford, UK, 2018.
40. Stella, M. Modelling early word acquisition through multiplex lexical networks and machine learning. *Big Data Cogn. Comput.* **2019**, *3*, 10. [[CrossRef](#)]
41. Warriner, A.B.; Kuperman, V.; Brysbaert, M. Norms of valence, arousal, and dominance for 13,915 English lemmas. *Behav. Res. Methods* **2013**, *45*, 1191–1207. [[CrossRef](#)] [[PubMed](#)]
42. Russell, J.A. A circumplex model of affect. *J. Pers. Soc. Psychol.* **1980**, *39*, 1161. [[CrossRef](#)]
43. Unwin, M.M.; Kenny, D.T.; Davis, P.J. The effects of group singing on mood. *Psychol. Music* **2002**, *30*, 175–185. [[CrossRef](#)]
44. Fonagy, P.; Gergely, G.; Jurist, E.L. *Affect Regulation, Mentalization and the Development of the Self*; Routledge: Abingdon, UK, 2018.
45. Stella, M. Forma mentis networks reconstruct how Italian high schoolers and international STEM experts perceive teachers, students, scientists, and school. *Educ. Sci.* **2020**, *10*, 17. [[CrossRef](#)]
46. Miller, G.A. *WordNet: An Electronic Lexical Database*; MIT Press: Cambridge, MA, USA, 1998.
47. Castro, N.; Stella, M. The multiplex structure of the mental lexicon influences picture naming in people with aphasia. *J. Compl. Netw.* **2019**, *7*, 913–931. [[CrossRef](#)]

48. Brugnoli, E.; Cinelli, M.; Quattrociocchi, W.; Scala, A. Recursive patterns in online echo chambers. *Sci. Rep.* **2019**, *9*, 20118. [[CrossRef](#)]
49. Krueger, J. On the perception of social consensus. In *Advances in Experimental Social Psychology*; Elsevier: Amsterdam, The Netherlands, 1998; Volume 30, pp. 163–240.
50. Pearce, E.; Launay, J.; Dunbar, R.I. The ice-breaker effect: Singing mediates fast social bonding. *R. Soc. Open Sci.* **2015**, *2*, 150221. [[CrossRef](#)]
51. Christensen, J.F.; Di Costa, S.; Beck, B.; Haggard, P. I just lost it! Fear and anger reduce the sense of agency: A study using intentional binding. *Exp. Brain Res.* **2019**, *237*, 1205–1212. [[CrossRef](#)] [[PubMed](#)]
52. DeWall, C.N.; Anderson, C.A.; Bushman, B.J. The general aggression model: Theoretical extensions to violence. *Psychol. Viol.* **2011**, *1*, 245. [[CrossRef](#)]
53. Polanyi, L.; Zaenen, A. Contextual valence shifters. In *Computing Attitude and Affect in Text: Theory and Applications*; Springer: Berlin, Germany, 2006.
54. Mohammad, S.M. Obtaining Reliable Human Ratings of Valence, Arousal, and Dominance for 20,000 English Words. In Proceedings of The Annual Conference of the Association for Computational Linguistics (ACL), Melbourne, VIC, Australia, 15–20 July 2018.
55. Kuppens, P.; Tuerlinckx, F.; Yik, M.; Koval, P.; Coosemans, J.; Zeng, K.J.; Russell, J.A. The relation between valence and arousal in subjective experience varies with personality and culture. *J. Pers.* **2017**, *85*, 530–542. [[CrossRef](#)] [[PubMed](#)]
56. Vergallito, A.; Petilli, M.A.; Cattaneo, L.; Marelli, M. Somatic and visceral effects of word valence, arousal and concreteness in a continuum lexical space. *Sci. Rep.* **2019**, *9*, 20254. [[CrossRef](#)]
57. Thelwall, M.; Thelwall, S. Retweeting for COVID-19: Consensus building, information sharing, dissent, and lockdown life. *arXiv* **2020**, arXiv:2004.02793.
58. Bedford, J.; Enria, D.; Giesecke, J.; Heymann, D.L.; Ihekweazu, C.; Kobinger, G.; Lane, H.C.; Memish, Z.; Oh, M.D.; Schuchat, A.; et al. COVID-19: Towards controlling of a pandemic. *Lancet* **2020**, *395*, 1015–1018. [[CrossRef](#)]



© 2020 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).