Online Social Networks and Media 37–38 (2023) 100271

Contents lists available at ScienceDirect

Online Social Networks and Media

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

Projection of Socio-Linguistic markers in a semantic context and its application to online social networks

Tomaso Erseghe [a,*], Leonardo Badia [a,1], Lejla Džanko [b,1], Magdalena Formanowicz [b,1], Jan Nikadon [b,1], Caterina Suitner [c,1]

[a] Department of Information Engineering, University of Padova, via Gradenigo, 6/b, Padova, 35131, Italy
[b] Center for Research on Social Relations, University of Social Sciences and Humanities (SWPS), Chodakowska 19/31, Warszaw, 03-815, Poland
[c] Department of Developmental Psychology and Socialization, University of Padova, via Venezia, 8, Padova, 35131, Italy

A R T I C L E I N F O

Keywords:
Data analysis
Computational linguistics
Projection algorithms
Social networking
Sociology–Psychology
#FridaysForFuture
#MeToo
#Covid 19

A B S T R A C T

Relevant socio-psychological processes can be detected in social networks thanks to an analysis of linguistic markers that sheds light on the characteristics and dynamics of the social discourse. Usually, linguistic markers comprise a list of words representative of a given construct; however, this approach does not account for contextual interdependencies of words, which can amplify or diminish the relevance of a particular word. In this paper, we present and leverage a scalable method called PageRank-like marker projection (PLMP) that addresses this problem. Its rationale, inspired by PageRank, is meant to fully exploit the interdependencies in a semantic network to project markers from a social discourse level (tweets) to its semantic elements (words). We show how PLMP is able to associate markers with specific words from their semantic context, which allows for an even richer interpretation of the online sentiment. We demonstrate the effectiveness of PLMP in practice by considering specific instances of social discourse on Twitter for three exemplary calls to collective action.

1. Introduction

Online social networks connect people and convey ideas faster than any real or virtual meeting platform [1]. As social media platforms are increasingly ubiquitous in our daily lives, they have become rich data sources that can be analysed to understand the collective attitude, belief, and behaviour of society [2]. Yet, the online corpus of micro-blogging platforms is a melting pot of content, often bubbly and noisy. Consequently, a prominent challenge for research is to employ analytical methods to extrapolate the underlying meaning, capture the zeitgeist, and predict evolving trends [3].

The effective utilization of data requires a strong multidisciplinary collaboration. First, quantitative analysis is a key component of data mining and involves the handling of large and complex data sets. To this end, technological know-how is essential to access, process, and store the vast amounts of data generated by online social networks [4]. At the same time, social sciences play a critical role in data mining, as they provide the theoretical frameworks and conceptual tools necessary to understand the complex dynamics that underlie online interactions. Social scientists can help identify relevant variables to analyse, develop appropriate sampling methods, and interpret the results of data mining in the context of social theory [5–7]. Finally, mathematical formalization is important for creating models that can capture the dynamics of social interactions and make predictions about future trends [8]. These models can be used to identify patterns and correlations in the data that may not be apparent through qualitative analysis alone. Overall, the success of data analytics in social online networks relies on a balanced blend of quantitative methods, technological know-how, social sciences, and mathematical formalization.

In this article, we apply such a blend to propose a novel approach allowing for the projection of socio-psychological linguistic markers from the holistic perspective of the social discourse to its semantic elements, i.e., words.

Building on an earlier proposal available in [9], we developed a PageRank-like marker projection (PLMP) [1], a methodology that accounts for the interdependencies in the semantic network. In PLMP, a bipartite network of tweets/words is processed through a PageRank-inspired approach [10,11] where information freely flows through the network interdependencies. Unlike PageRank, a row-normalized update matrix (as opposed to column-normalized) is used to maintain coherence with the end goal.

* Corresponding author.
E-mail addresses: tomaso.erseghe@unipd.it (T. Erseghe), leonardo.badia@unipd.it (L. Badia), ldzanko@st.swps.edu.pl (L. Džanko), mformanowicz@swps.edu.pl (M. Formanowicz), jnikadon@swps.edu.pl (J. Nikadon), caterina.suitner@unipd.it (C. Suitner).

[1] These authors contributed equally to this work.

https://doi.org/10.1016/j.osnem.2023.100271
Received 4 April 2023; Received in revised form 11 August 2023; Accepted 30 September 2023
Available online 22 October 2023

In this paper, we develop and expand the contribution of Erseghe et al. [1] along three different avenues. First, we give rigorous mathematical proof of the exactness of PLMP, along the lines of Haveliwala and Kamvar [12], with proper modifications. Second, we show how PLMP enables the investigation of the rhetoric of calls to action in the online social discourse. PLMP allows for a quantification of the syntagms used in terms of certain semantic attributes, based not only on their intrinsic meanings but also on their structural role within the context of the online discourse. In this way, we can show how PLMP is able to detect structural changes of the semantic network from a holistic perspective. Finally, we expand the PLMP evaluations to a much broader set of data, considering three main events that had a strong global impact, reflected in online social communities. We focus on social discourse on Twitter and specifically consider: (1) #MeToo (2017–2018): an initiative encouraging victims of sexual harassment (typically young women) to break the silence; (2) #FridaysForFuture (2018–2019): school strikes demanding action from global leaders to counteract climate change; and (3) #Covid19 (2020): a world-wide pandemic that revolutionized habits and social interactions around the globe. For all these cases, PLMP is used to show how many pertinent keywords in each of the datasets can have a translational change in meaning and relevance, which through a stronger information flow enables collective action. In addition, we demonstrate how PLMP can discover significant variations between the three considered calls to action due to their contextual differences.

The rest of the paper is organized as follows. Section 2 discusses the most relevant findings of the literature that relate to the present work. Section 3 introduces the PLMP approach, and presents a number of possible alternatives that can be devised by exploiting the same rationale. In Section 4, the theoretical/mathematical results related to PLMP convergence are proved. The application to the analysis of call-to-action scenarios is presented in Section 5, showing quantitative results. Section 6 concludes the paper.

2. Related work

2.1. Socio-psychological linguistic markers in calls-to-action

Engaging in a collective action challenges the social system and requires a personal and social investment. As this endeavour runs contrary to the common need for stability, generally operating in favour of the status quo [13], several scholars have investigated socio-psychological factors that drive individuals to such an outstanding and costly devotion to the cause. Notably, the Social Identity Model of Collective Action (SIMCA) [6,14] recognizes three relevant drives for engaging in collective action, namely: *social identity*, *collective efficacy*, and *anger*.

Social identity First, as indicated by the name, collective action requires individuals to act collectively in the interest of the group. In order to do that, individuals need to see themselves as members of a group with shared values, motives, and emotions [15]. The extent to which an individual embraces the group's identity (for example of fellow citizens in case of the pandemic) is predictive of participation in actions aimed at goals relevant for the collective, such as mask wearing [16]. Therefore, the emergence of this common identity, interdependence, and shared meaning is critical to broadening the social scale from an individual perspective focusing on self-interest ("I") to the collective perspective of a common good ("we") [7].

Collective efficacy Having a collective, however, is only one prerequisite for the formation of the collective movements. What is also necessary is a shared idea that collective mobilization can indeed contribute to a broader change. Defined as collective efficacy or agency, it is the sense that the group has adequate resources and sufficient control to achieve the desired social change. In line with resource mobilization theory [17] and its further developments [18], agency is conceptualized as the perceived capacity and the effort of the group to manage collective resources and mobilize its members to achieve an improvement for the group. As this sense of agency has also been identified as a precursor for any human action [19], not surprisingly it is considered a key feature in online mobilizations [4,5].

Anger The third core feature identified as a driver of collective action pertains to the emotional reactions tied to the moral evaluation of the social context. When members of a group assess that their group has been mistreated or that the current status quo runs against group values, they tend to experience individual and group-level emotions, of which anger is the most prominent for collective mobilization. In other words, the emotional aspect of engagement in political action is amplified by a mix of personal and group-based moral motivations. Moral outrage is activated when the issue is of high relevance to one's standards and accordingly leads to doing what is considered right in the situation: collective action [20]. As a result, moral motivations can have a catalyst function that sets people into motion and prompts societal change. Some authors, however, find anger as more harmful than beneficial, because it could cause violence and social division. As moral outrage is more likely to occur among privileged members of the society at the cost of minority groups [21], some results show that anger is less associated with collective action among people of low status than among people of high status [22].

Despite the recent advancements in the understanding of the role of anger in collective action, all three features affecting collective action (i.e., social identity, collective efficacy, and anger) are investigated in the domain of the online movements, see [4,23]. The primary means of that investigation are based on the quantification of their linguistic presence, e.g., through dictionary methods.

2.2 The use of dictionaries

The use of dictionaries is a popular method to investigate socio-cognitive processes in natural language use in general, and specifically their emergence in social media discourse. Only to mention a few, the NCR (National Research Council of Canada) lexicon developed by Mohammad and Turney [24] provides codes for emotions and sentiments; Brysbaert et al. [25] developed concreteness norms for over 40K words, applied for example in tracking temporal changes in language prior to the important societal events [26]; gender stereotypes can be captured by counting words expressing masculine and feminine words [27].

Particularly prevalent in the field is the use of Linguistic Inquiry and Word Count (LIWC) software [28]. LIWC 2015 is a well-established instrument for detecting linguistic proxies of psychological processes in text samples [29]. It contains a wide set of bags of words (over 60 dictionaries), coded and validated for their accuracy in reflecting psychological content. The software counts the number of words that belong to psychologically meaningful categories, and there is ample empirical evidence proving its validity to capture social phenomena on popular media such as Twitter [30–33]. Particularly relevant for our work is the application of LIWC to the language in online collective actions. Gulliver et al. [34] analysed the language used by 497 Australian environmental organizations to promote collective engagement. In line with the literature on the psychological drivers of collective action [6,14], the language of environmental advocacy was characterized by words that reflect collective efficacy and social identity, operationalized through the use of affiliation dictionary and the use of the first person plural pronoun "we". Accordingly, in this research, we rely on LIWC to quantify social identity (through the affiliation dictionary in Study 1 and "we" usage in Study 2 - see also [4]) and anger (through the anger dictionary).

2.3 The limits of dictionaries

There are multiple advantages of using word-count (or dictionary) methodology specifically, their simplicity of implementation and usage. This simplicity, however, comes at a cost, such as the lack of sensitivity to polysemy as well as the lack in the nuance of how well words represent a given construct, since all words in a dictionary are typically assumed to contribute identically.

Deep learning tools are today widely used to overcome these limitations and their superiority in comparison to traditional approaches is postulated (e.g., in [35,36]). State-of-the-art techniques mostly rely on the transformer architecture [37], the core of ChatGPT, and on related models such as Bidirectional Encoder Representations from Transformers (BERT) [38]. The research in this area is very active, especially in relation with sentiment analysis, see [39–42], to cite a few relevant ones. A worth mentioning tool dealing with agency is BERTAgent, a computational language representation model fine-tuned to detect agency in textual data [43]. This model is of particular value since it is substantiated by a set of validation studies and proven to be more powerful than the common dictionary approach (e.g., LIWC). Accordingly, we quantify collective efficacy with the use of this newly developed tool.

At the present time, we are not familiar with any deep learning tools developed for quantifying social identity in language. While, there are tools available for the quantification of emotions (e.g., [44,45]), there is no evidence on their improved performance with respect to dictionary based approaches. Accordingly, as described above, we rely on the dictionary methods to measure social identity and anger.

2.4 Page-rank approaches

It is important to note, that online calls to action are not simply a collection of words, which meanings can be considered in isolation or even in addition to one another. Calls to action comprise complex utterances in which meaning of words is intertwined, and these interdependencies can be well captured by a network formalism. Constraining ourselves to the semantic content, a widely assessed methodology relies on building networks that relate the elements of inspection, i.e., the documents (or tweets if referring to an online social network), through their constituent factors, i.e., the words. This can be done either directly by linking documents via the number of words they share or, equivalently, by linking words via the number of documents they appear in [4,46,47] or by doing it indirectly through their mathematical representation [48]. In both cases, the resulting network is a powerful structure that allows to extract holistic information in an unsupervised manner, i.e., without any previous knowledge on the network properties.

A few methods based on a flow-of-information concept have gained particular interest over the years. In this context, PageRank [10,11], the original core of Google's search algorithm, is an iterative method to spread information from a node to its neighbourhood. It is a way to assess the node importance not only based on the number of its connections, but also as a measure of the quality of such connections. Its mathematical properties [12] ensure fast convergence in few iterations, only 35 for the 80 million nodes network of Kamvar et al. [49], which makes its computational complexity almost linear and largely affordable. It is definitely tractable in the context of online social networks, where a targeted samples of a few million tweets or less are completely sufficient for very detailed analyses [50]. Variations of the PageRank rationale have been used for many purposes: Haveliwala [51] modifies the original idea to evaluate vicinity/similarity to a subset of nodes in the network (a community); Weng et al. [52] modifies the flow-of-information weight in the network to identify relevance with respect to spreading a specific topic in the online discussion; He et al. [53] applies it in a bipartite network setup by modifying the constituent matrix, which controls the flow-of-information, to solve a regularization problem; Cao et al. [54], which is the closest reference to this paper, uses a modified version of PageRank where the constituent matrix is row-normalized (as opposed as to column-normalized) to return a ranked list of related entities for a given query. Hits [55] is another widely cited alternative to PageRank based on similar concepts, but the absence of strong mathematical properties in its constituent matrix fails to make its convergence properties scalable, hence it is of limited appeal in large networks.

3 Method: the PLMP approach

3.1 Rationale

To improve the methods of capturing relevant aspects of the online discourse and its related social dynamics, we introduce the PLMP method and assess its performance. To describe the procedure, assume that we can exploit data extracted from social media, for example Twitter, forming a bipartite network of *tweets* and *words*. Furthermore, assume that a socio-psychological linguistic marker (e.g., agency) is extracted separately for each tweet and each word appearing on tweets, and these are stored in vectors $\tilde{\boldsymbol{m}}_t$ and $\tilde{\boldsymbol{m}}_w$, respectively. Our intent is to define PLMP on this network and then perform a semantic analysis of its content oriented toward the socio-psychological context.

We review the problem at hand and illustrate the PLMP solution with the help of Fig. 1, where, without any loss of generality, we consider a sample semantic network. This network was formed from tweets related to the #MeToo movement in 2018, and as an illustrative linguistic feature we chose *agency*, i.e., action and goal orientation, sense of which is necessary for people to attempt social change [56]. The exact details are explained later in Section 5, but are not necessary for the understanding of this illustration. The word-clouds in Fig. 1 are built in such a way that the nodes size is proportional to their PageRank centrality score (and also roughly proportional to the frequency of each word in the discussion). Thanks to the application of a graph layout – forceatlas2 as implemented in Gephi [57] – and a topic detection algorithm, see [4], also the position of nodes carries a meaning in Fig. 1. Specifically, words belonging to the same topic in the #MeToo discussion, as well as nodes that frequently appear together in tweets, are placed close to one another in the graphical representation. For example, in the lower right corner the words *gender, right, speak, need, take, action* refer to a call-to-action, while the central *help, achieve, equality* identify a discussion topic. Size and position of words is the same in all subplots. The node colour, however, changes in each subplot, according to different projection approaches (Fig. 1.(c) refers to PLMP) as explained in what follows.

Fig. 1.(a) displays, with different colours, the agentic meaning in the absence of a social discourse, i.e., it measures the *in isolation* level of agency of words $\tilde{\boldsymbol{m}}_w$ without resorting to the social discussion expressed by tweets. The distribution of agency values $\tilde{\boldsymbol{m}}_w$ is available in Fig. 2.(a). The limit of such an approach is evident in that, in a specific context (e.g., the #MeToo feminist rhetoric on Twitter) a number of words that are agency-neutral (e.g., *woman*), in this specific social discourse ought to carry a much higher level of agency. This means that an agency level just attributed to words in isolation fails to capture the entire complexity of a social discourse. We argue that the *in isolation* approach fails to capture the social interconnections and relationships expressed in the particular context, whose analysis requires carefully crafted solutions [58,59].

Consequently, a true quantification of agency is better captured through context-based extraction of meaning within the semantic interaction of tweets. We take the vector $\tilde{\boldsymbol{m}}_t$ (tweets markers) to carry this context-specific information. Our goal is to identify an algorithm that reliably assigns agency to words taking into account the available information from $\tilde{\boldsymbol{m}}_w$ and $\tilde{\boldsymbol{m}}_t$. Thus, we leverage the indirect effect of socio-psychological markers, as somehow hinted by the well-known PageRank approach [10,11], i.e., by letting information (iteratively)

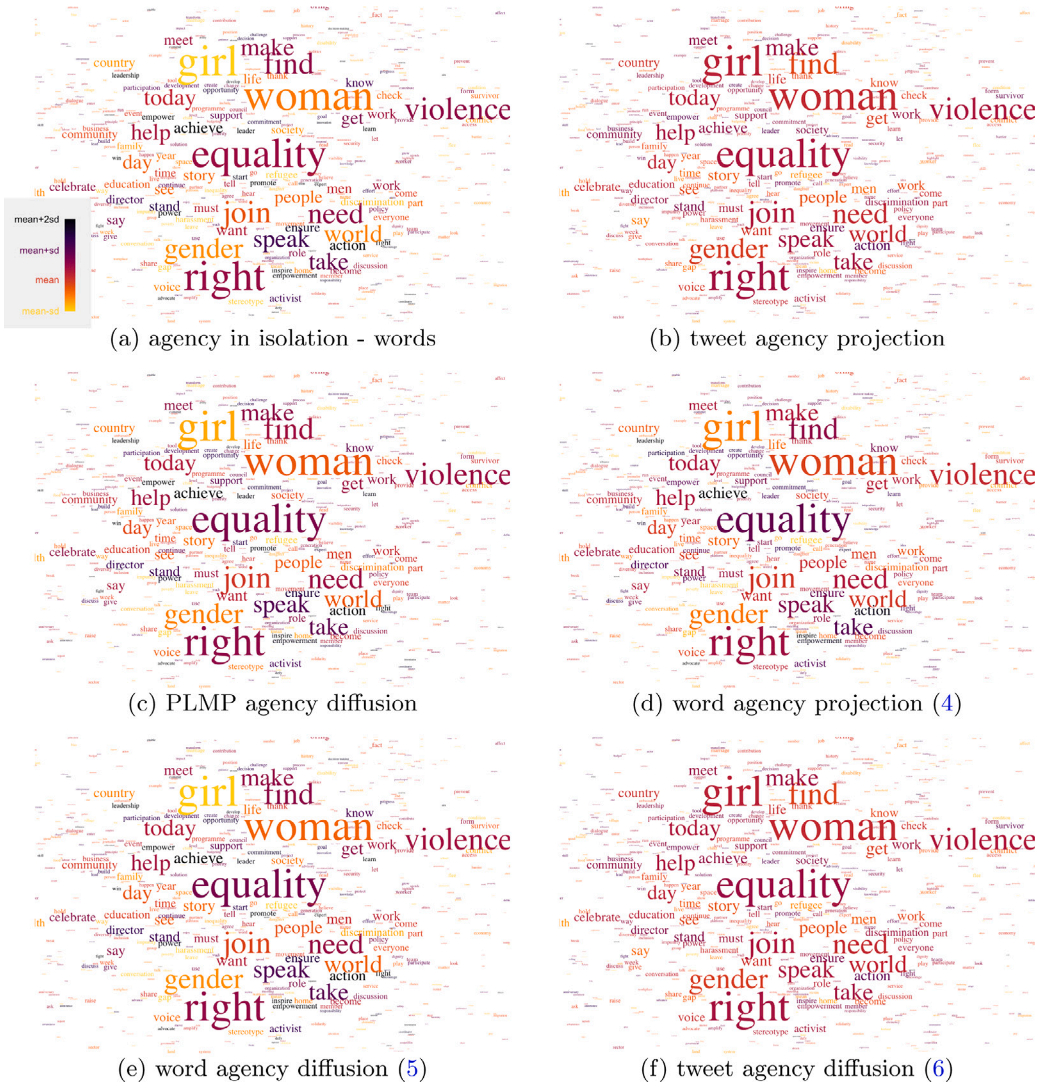

Fig. 1. Comparison of different methods for agency contextualization – Agency word clouds for #MeToo network in 2018; node size is proportional to the PageRank value in the network, colour corresponds to the level of agency. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

traverse the semantic network. This corresponds to arguing that, in the association of syntagms to a given indicator, a word is affected by the mean values of the indicator contained in the tweets it belongs to; conversely, the parent tweet is also affected by the mean values of the words it contains. If the adjacency matrix linking words to tweets is denoted as $\boldsymbol{B}$, then averages can be inferred from the row-normalized ($\boldsymbol{B}_1$) and the column-normalized-and-transposed ($\boldsymbol{B}_2$) counterparts to $\boldsymbol{B}$, respectively, as

$$\begin{aligned} \boldsymbol{B}_1 &= \mathrm{diag}((\boldsymbol{B}\mathbf{1})^{-1}) \cdot \boldsymbol{B} \\ \boldsymbol{B}_2 &= \mathrm{diag}((\boldsymbol{B}^T\mathbf{1})^{-1}) \cdot \boldsymbol{B}^T \; . \end{aligned} \tag{1}$$

Specifically, the flow-of-information rationale discussed above can be formalized as a PageRank-like steady-state equation

$$\underbrace{\begin{bmatrix} \boldsymbol{m}_w \\ \boldsymbol{m}_t \end{bmatrix}}_{\boldsymbol{m}} = \alpha \underbrace{\begin{bmatrix} \mathbf{0} & \boldsymbol{B}_1 \\ \boldsymbol{B}_2 & \mathbf{0} \end{bmatrix}}_{\boldsymbol{M}} \begin{bmatrix} \boldsymbol{m}_w \\ \boldsymbol{m}_t \end{bmatrix} + (1-\alpha) \underbrace{\begin{bmatrix} \tilde{\boldsymbol{m}}_w \\ \tilde{\boldsymbol{m}}_t \end{bmatrix}}_{\boldsymbol{q}} \; , \tag{2}$$

where $\alpha \in (0,1)$ is a mixing parameter that can be seen as the rate of information spreading. It is possible to solve (2) through a standard power iteration, by subsequently applying

$$\boldsymbol{m}_k = \alpha \boldsymbol{M} \boldsymbol{m}_{k-1} + (1-\alpha)\boldsymbol{q} \; , \tag{3}$$

in which averages are iteratively exchanged from tweets to words, and then from words to tweets, starting from an initial state $\boldsymbol{m}_0 = \boldsymbol{q}$, which is a required initialization to achieve the correct solution. Note that, although (2) looks like a standard PageRank equation, here matrix $\boldsymbol{M}$ is row-normalized and not column-normalized, which implies a number of technical complications, discussed in later Section 4, to prove the exactness of (3) and assess the validity of (2). Section 4 also proves that the eigenstructure of $\boldsymbol{M}$ is equivalent to that of a PageRank approach, hence the power iteration (3) exhibits the same properties of PageRank iterations, i.e., it is scalable and easily applicable to very large networks at a low computational cost. As a side note, (2) directly works on the bipartite network, which is a particularly useful approach that avoids a projection onto a network of words, a step that would actually discard essential information [60]. We also observe that the above idea can be strengthened by taking into account any further available information. For example, any metric on the tweets relevance (e.g., number of retweets), or the word relevance, can be used to enhance the effect of more important tweets/words, i.e., to obtain a weighted average effect. This can be done by assigning different weights to the columns of $\boldsymbol{B}_1$ and $\boldsymbol{B}_2$, prior to row-normalization. We finally point out that the reference equation (2), using a row-stochastic matrix $\boldsymbol{M}$, is also used in [54] as a method to refine an initial ranking estimate, which in the present context can be interpreted as a map from $\tilde{\boldsymbol{m}}_t$ to $\boldsymbol{m}_t$. Differently, here we propose to exploit this PageRank-like approach in a completely different scenario, where we are interested in projecting information

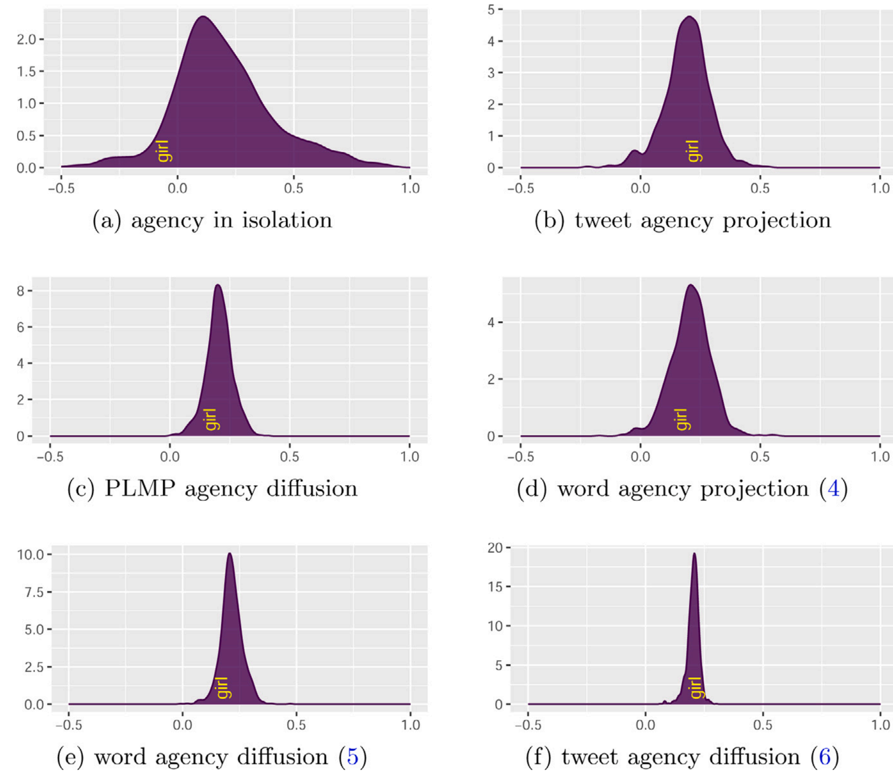

Fig. 2. Comparison of different methods for agency contextualization – Agency densities for the wordclouds of Fig. 1. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

from one class to another, i.e., from $\tilde{\boldsymbol{m}}_t$ to $\boldsymbol{m}_w$, further corroborating our proposal by the key findings of Section 4. The outcome of PLMP is graphically shown in the wordcloud of Fig. 1.(c), whose corresponding density is shown in Fig. 2.(c).

To get a full understanding on the nature of PLMP, we focus on the upper part of Figs. 1 and 2 comparing:

(a) *agency in isolation* $\tilde{\boldsymbol{m}}_w$;
(b) *tweet agency projection*, namely the state-of-the-art solution [9], which corresponds to $\boldsymbol{m}_w = \boldsymbol{B}_1 \tilde{\boldsymbol{m}}_t$, i.e., to assigning to a word the average marker value of its parent tweets;
(c) our proposed solution PLMP, in which averages are iteratively exchanged from tweets to words, and then from words to tweets.

Note that, in the word plots of Fig. 1, the colour scale is normalized, to allow for a comparison of results that are different in absolute terms. Different approaches in fact imply different active ranges, as can be inferred from the density plots of Fig. 2. Namely, as shown in the legend, the orange colour refers to the average value of the agency distribution, whereas darker (purple) or lighter (yellow) colour denote words that deviate from that, up to an increase of two standard deviations or a decrease of one standard deviation, respectively. By observing Fig. 1.(c), it is evident how PLMP is able to combine the agency *in isolation* of Fig. 1.(a) and the *one-hop* projection [9] of Fig. 1.(b), in two different ways. First, words that are agentic in themselves keep their original *in isolation* agency level, e.g., *stand*, *speak*, and *achieve* in dark colour in both Fig. 1.(a) and (c) and in lighter colour in Fig. 1.(b). This effect is visible from the wider colour range of both Fig. 1.(a) and (c), compared to the narrow range of Fig. 1.(b). Second, thanks to the flow through the semantic network, words acquire agency from the social discourse. This is evident from the density plots of Fig. 2.(b) and (c) whose agency values belong to the positive range, while in Fig. 2.(a) we observe the presence of a large fraction of values with negative/zero agency; e.g., the agentic level of *girl* is negative in Fig. 2.(a), and turns to positive in both Fig. 2.(b) and (c).

3.2 Convergence validity and alternative formulations

To fully assess the adequateness of PLMP, we evaluate its divergent and convergent validity. To establish that the projection is not inflated by the specific role that a target word (possibly a relevant one) has on the network, we compare PLMP with PageRank, proving its independence. To verify its convergent validity, we compare its congruence with some alternative projection methods that we formulated, as valid alternatives to PLMP, based on either Badia et al. [9] or the PageRank rationale. These alternatives further validate that our proposal reliably captures the phenomenon under investigation. They are displayed in the plots below in Figs. 1 and 2, as follows:

(d) that we call *word agency projection*, applies the approach of Badia et al. [9] to the sub-network of words whose (projected) adjacency matrix takes the form $\boldsymbol{M}_w = \boldsymbol{B}_1 \boldsymbol{B}_2$ [61], so that agency projection is inferred from

$$\boldsymbol{m}_w = \boldsymbol{M}_w \tilde{\boldsymbol{m}}_w \tag{4}$$

(e) that we call *word agency diffusion*, generalizes (4) by adopting a PageRank diffusion, to identify a steady-state equation of the form

$$\boldsymbol{m}_w = \alpha \boldsymbol{M}_w \boldsymbol{m}_w + (1-\alpha)\tilde{\boldsymbol{m}}_w \tag{5}$$

(f) that we call *tweet agency diffusion*, conversely exploits $\tilde{\boldsymbol{m}}_t$ in a PageRank-like context, to obtain a projection $\boldsymbol{M}_t = \boldsymbol{B}_2 \boldsymbol{B}_1$ on tweets, and have

$$\begin{aligned} \boldsymbol{m}_t &= \alpha \boldsymbol{M}_t \boldsymbol{m}_t + (1-\alpha)\tilde{\boldsymbol{m}}_t \\ \boldsymbol{m}_w &= \boldsymbol{B}_1 \boldsymbol{m}_t\,. \end{aligned} \tag{6}$$

From the plots of Figs. 1 and 2, one can see how word agency projection (d) and word agency diffusion (e) closely relate to PLMP (c) in both wordcloud and density, while tweet agency diffusion (f) gathers a result similar to tweet agency projection (b), although with a more compact active region as evident from the agency density in Fig. 2.(f).

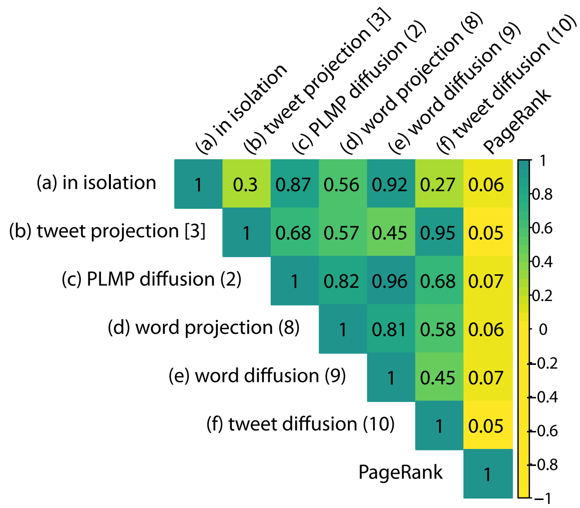

Fig. 3. Pearson's correlation matrix comparing projection methods among themselves and with PageRank centrality; all values are statistically significant.

A full view of the interdependencies between the six approaches is available in Fig. 3, showing Pearson's correlations between different options. With the help of Fig. 3, we can infer how different approaches are able to balance the two main effects of keeping the information *in isolation* as visible from row (a) as well as the averaging effect of *one hop* spreading [9] displayed in row (b). The best compromise is in this case obtained by PLMP, as it outperforms word projection (d) in both correlation values, and it also stays away from the unbalanced correlation values of word diffusion (e) (slightly unbalanced) and tweet diffusion (f) (strongly unbalanced). The consistency of Fig. 3, tested over different semantic networks, confirms the predicted PLMP advantage as the preferred choice among the proposed alternatives implementing the same concept.

Fig. 3 further investigates the relation with the PageRank centrality of the network, highlighting the absence of correlation. This feature, especially for approaches based on a PageRank-like rationale, is a guarantee of adequateness as it implies that the flow of agency across the network is not redundant with the nodes centrality.

4 Mathematical insights

4.1 Solving PLMP through power iterations: a proof

We explore the solutions to the steady-state equation (2) for a row-normalized square matrix $\boldsymbol{M}$, i.e., $\boldsymbol{M}\boldsymbol{1} = \boldsymbol{1}$, and for $\boldsymbol{q} \geq \boldsymbol{0}$, $\boldsymbol{q} \neq \boldsymbol{0}$. If $\boldsymbol{M}$ is *irreducible*, which is the situation of greatest interest in our problem, and the mixing parameter α lies in the open interval $(0,1)$, then one can mimic the PageRank analysis of Haveliwala and Kamvar [12], with some modifications because of the different normalization of $\boldsymbol{M}$. Assume that $\boldsymbol{v}$ is the left eigenvector corresponding to the right eigenvector $\boldsymbol{1}$, i.e., related to eigenvalue 1. From Perron–Frobenius theorem [62], without loss of generality we can assume $\boldsymbol{v} > \boldsymbol{0}$; then, (2) gives

$$\boldsymbol{v}^T\boldsymbol{m} = \alpha\, \underbrace{\boldsymbol{v}^T\boldsymbol{M}}_{\boldsymbol{v}^T}\, \boldsymbol{m} + (1-\alpha)\boldsymbol{v}^T\boldsymbol{q}\,, \tag{7}$$

so that $\boldsymbol{v}^T\boldsymbol{m} = \boldsymbol{v}^T\boldsymbol{q} > 0$. This enables a rewriting of the steady-state equation (2) as

$$\boldsymbol{m} = \underbrace{\left[\alpha\boldsymbol{M} + (1-\alpha)\frac{\boldsymbol{q}\boldsymbol{v}^T}{\boldsymbol{v}^T\boldsymbol{q}}\right]}_{\boldsymbol{M}_1}\boldsymbol{m}\,, \tag{8}$$

where matrix $\boldsymbol{M}_1$ satisfies $\boldsymbol{v}^T\boldsymbol{M}_1 = \boldsymbol{v}^T$ by construction, i.e., $\boldsymbol{v}^T$ is a left eigenvector of $\boldsymbol{M}_1$ associated to eigenvalue 1. Instead, the corresponding right eigenvector, providing the solution $\boldsymbol{m}$, is different from $\boldsymbol{1}$. As the irreducibility of $\boldsymbol{M}$ implies that $\boldsymbol{M}_1$ is irreducible as well, then Perron–Frobenius theorem guarantees that the eigenvalues of $\boldsymbol{M}_1$ satisfy $|\lambda| \leq 1$, and eigenvalue 1 has multiplicity 1.

The *left* eigenvector $\boldsymbol{v}$ can be employed to characterize the *right* eigenvectors of $\boldsymbol{M}_1$. If the Jordan form of $\boldsymbol{M}_1$ is $\boldsymbol{M}_1 = \boldsymbol{R}\boldsymbol{J}\boldsymbol{R}^{-1}$, where $\boldsymbol{R}$ collects the (generalized) right eigenvectors of $\boldsymbol{M}_1$, and, conversely, $\boldsymbol{R}^{-1}$ collects the right eigenvectors, we have

$$\underbrace{\boldsymbol{v}^T\boldsymbol{M}_1}_{\boldsymbol{v}^T}\,\boldsymbol{R} = \boldsymbol{v}^T\boldsymbol{R}\boldsymbol{J}\,. \tag{9}$$

Therefore, $\boldsymbol{v}^T\boldsymbol{R}(\boldsymbol{I}-\boldsymbol{J}) = \boldsymbol{0}$. Since $\boldsymbol{M}_1$ contains only one eigenvalue equal to 1, all of the right eigenvectors, but the one related to $\boldsymbol{v}$, satisfy $\boldsymbol{v}^T\boldsymbol{r}_i = 0$. Thus, we get

$$\boldsymbol{M}_1\boldsymbol{r}_i = \alpha\boldsymbol{M}\boldsymbol{r}_i + (1-\alpha)\frac{\boldsymbol{q}\boldsymbol{v}^T}{\boldsymbol{v}^T\boldsymbol{q}}\boldsymbol{r}_i = \alpha\boldsymbol{M}\boldsymbol{r}_i \tag{10}$$

so that the right (generalized) eigenvector $\boldsymbol{r}_i$ of $\boldsymbol{M}_1$ is also a right (generalized) eigenvector of $\alpha\boldsymbol{M}$; as such, it is related to an eigenvalue $|\lambda_i| \leq \alpha$. Thus, 1 is also included among the eigenvalues of $\boldsymbol{M}_1$, as well as other eigenvalues with absolute value lower than or equal to α. This guarantees the convergence of (3) to the desired (and unique) solution; note the importance of imposing $\boldsymbol{m}_0 = \boldsymbol{q}$ to ensure $\boldsymbol{v}^T\boldsymbol{m}_k = \boldsymbol{v}^T\boldsymbol{q}$ throughout the iterations.

4.2 The reducible matrix case

An equivalent result can be obtained in case that $\boldsymbol{M}$ is *reducible*. If so, it can be organized (by a permutation of its elements) in the upper-triangular block form [63]

$$\boldsymbol{M} \sim \begin{bmatrix} \boldsymbol{B}_{1,1} & \boldsymbol{B}_{1,2} & \cdots & \boldsymbol{B}_{1,K} \\ & \boldsymbol{B}_{2,2} & \cdots & \boldsymbol{B}_{2,K} \\ & & \ddots & \vdots \\ & & & \boldsymbol{B}_{K,K} \end{bmatrix}$$

with *irreducible* diagonal blocks $\boldsymbol{B}_{k,k}$. Incidentally, the diagonal blocks identify sets of nodes that are strongly connected. By construction, all diagonal blocks have eigenvalues $|\lambda| \leq 1$. Moreover, by the Perron–Frobenius theorem, those diagonal blocks that satisfy $\boldsymbol{B}_{k,k}\boldsymbol{1} = \boldsymbol{1}$ have exactly one eigenvalue equal to 1 – these are blocks that correspond to *leaves* in the condensation graph [63]. Conversely, those diagonal blocks that satisfy $\boldsymbol{B}_{k,k}\boldsymbol{1} \neq \boldsymbol{1}$ have all their eigenvalues with $|\lambda| < 1$. In this context, vector $\boldsymbol{v}$ should be active only on those *leaves* that are reachable from blocks where $\boldsymbol{q}$ is active. If this is the case, then matrix $\boldsymbol{M}_1$ assumes the form

$$\boldsymbol{M}_1 \sim \begin{bmatrix} \tilde{\boldsymbol{B}}_{1,1} & \tilde{\boldsymbol{B}}_{1,2} & \cdots & \tilde{\boldsymbol{B}}_{1,K} \\ & \tilde{\boldsymbol{B}}_{2,2} & \cdots & \tilde{\boldsymbol{B}}_{2,K} \\ & & \ddots & \vdots \\ & & & \tilde{\boldsymbol{B}}_{K,K} \end{bmatrix}$$

where all $\tilde{\boldsymbol{B}}_{k,k} = \alpha\boldsymbol{B}_{n_k,n_k}$ for some n_k, i.e., they are related to eigenvalues $|\lambda| \leq \alpha$, while $\tilde{\boldsymbol{B}}_{K,K}$ is an irreducible block containing all the nodes where both $\boldsymbol{q}$ and $\boldsymbol{v}$ are active. Now, to guarantee that $\tilde{\boldsymbol{B}}_{K,K}$ has only one eigenvalue equal to 1, which also corresponds to the spectral radius, we need that $\boldsymbol{v}$ is non-zero in all the nodes belonging to $\tilde{\boldsymbol{B}}_{K,K}$, which is a consequence of the Perron–Frobenius theorem. This is the case if $\boldsymbol{q}$ is active in the leaf nodes of $\boldsymbol{M}$. In our specific setting, where $\boldsymbol{M}$ is not symmetric, this is ensured by the fact that if a link is active from nodes i and j then it is active in both directions, although with a different weight. As a consequence, the same result as above is obtained.

5 Results and discussion: Application to socio-psychological linguistic markers

5.1 Datasets

We assess the ability of PLMP to capture relevant aspects of online discourse and its related social dynamics by applying it to three different scenarios. To this end, we extracted data from social media related to well-known events, with the intent to perform contextual analyses of their contents through PLMP. In all cases, we chose Twitter as a well-suited reality mirror for our analyses [2], because of its widespread usage and the ease of accessing data through the APIs. From this raw data, we extracted different semantic networks, which we analysed using PLMP. Unlike other common approaches [64], we did not limit our analysis to a bipartite graph of tweets and hashtags, but rather we considered all the words present in the tweets, to better capture the contextual interdependencies in the social discourse. Thus, we always created a network of *tweets* and *words* (or words plus *hashtags*, depending on the context), where the former ones are connected to the latter that appear inside them.

We limited our scope to tweets in the English language, and we sampled two groups of tweets before and after a main event, in three scenarios:

1. #MeToo – Tweets from the @UN_Women pages in the periods of April 1–June 30, 2017 and April 1–June 30, 2018 – 1500 tweets pre and 1500 post the event;
2. #FridaysForFuture – Tweets in the periods of March 1–April 19, 2018 and March 1–April 19, 2019 by using the neutral hashtag #climatechange in the search [4] – 5000 tweets pre and 5000 post the event.
3. #Covid19 – Tweets in the periods of February 11–March 10, 2020 and March 11–April 11, 2020 by using the hashtag #Covid19 in the search; 1000 tweets per week were selected among the most influential ones, i.e., those that received the most reactions, i.e., replies, likes, and quotes.

After getting a sizeable corpus to read as a semantic network, identification and marking were achieved by means of POS non-deterministic tagger [3,65]. We applied a post-processing in which we: discarded all words not tagged as nouns, adjectives, adverbs, or verbs, unless otherwise stated; expanded contracted forms; split non-meaningful composed words; removed stopwords. The remaining words were lemmatized preserving their POS tag.

5.2 Socio-psychological linguistic markers of interest

To establish the validity of the PLMP method in capturing relevant aspects of the online discourse and social dynamics, the resulting semantic networks were analysed with a socio-psychological lens. In particular, we inspected the core features of collective action discussed in Section 2, namely, social identity, collective efficacy, and anger [66], represented in linguistic markers of *affiliation*, *agency*, and *anger* [4], respectively. These were quantitatively inferred from tweets by means of state-of-the-art software tools. Specifically, a validated deep learning approach was chosen for agency, while more standard dictionary approaches were used for affiliation and anger as no validated deep learning tool is today available in this case. The tools are described in what follows.

For affiliation and anger, we leverage Linguistic Inquiry and Word Count (LIWC) 2015 [67], a tool that performs a dictionary-based quantitative content analysis, where every message gets scores relative to several categories; these scores are derived from the number of words belonging to the specific category, adjusted for the overall number of words within the message. LIWC entries of "affiliation" and "anger" were used.

(a) agency

	#MeToo***	#FridaysForFuture	#Covid19***
pre	$m = .1686$ $\sigma = .1121$	$m = .1376$ $\sigma = .1022$	$m = .0801$ $\sigma = .1034$
post	$m = .2061$ $\sigma = .1107$	$m = .1379$ $\sigma = .1133$	$m = .1128$ $\sigma = .1120$
increase	t-ratio= 9.183 Cohen $d = .876$ CI= [.687; 1.06]	t-ratio= .1090 Cohen $d = .0078$ CI= [−.132; .148]	t-ratio= 14.27 Cohen $d = .959$ CI= [.826; 1.09]

(b) affiliation

	#MeToo**	#FridaysForFuture**	#Covid19***
pre	$m = .0338$ $\sigma = .0435$	$m = .0204$ $\sigma = .0308$	$m = .0186$ $\sigma = .0284$
post	$m = .0371$ $\sigma = .0421$	$m = .0224$ $\sigma = .0306$	$m = .0310$ $\sigma = .0382$
increase	t-ratio= 2.115 Cohen $d = .078$ CI= [.006; .149]	t-ratio= 3.106 Cohen $d = .063$ CI= [.023; .103]	t-ratio= 17.19 Cohen $d = .364$ CI= [.322; .406]

(c) anger

	#MeToo	#FridaysForFuture	#Covid19***
pre	$m = .0035$ $\sigma = .0135$	$m = .0039$ $\sigma = .0126$	$m = .0046$ $\sigma = .0130$
post	$m = .0038$ $\sigma = .0126$	$m = .0043$ $\sigma = .0127$	$m = .0057$ $\sigma = .0138$
increase	t-ratio= .599 Cohen $d = .007$ CI= [−.015; .03]	t-ratio= 1.535 Cohen $d = .013$ CI= [−.004; .03]	t-ratio= 3.922 Cohen $d = .033$ CI= [.016; .049]

Fig. 4. Tweets statistics (mean m, standard deviation σ, and t-test t-ratio, Cohen's d and confidence interval CI) for agency (a), affiliation (b), and anger (c) – increase from pre to post the main event; we highlight with *** a statistically highly significant increase, and with ** a statistically significant increase (both also marked in orange).

Linguistic agency was instead measured using BERTAgent, a computational language representation model (LRM) fine-tuned to detect agency in textual data [43]. BERTAgent relies on the transformer architecture [37] and is based on BERT [38], an architecture that has achieved groundbreaking results in a broad range of natural language processing (NLP) tasks [68]. Compared to dictionary- and POS-based methods of agency quantification, BERTAgent offers improved sensitivity to polysemy and negation, because it captures not only agentic words but also the context of their usage. Thus, it provides improved validity and reliability for linguistic agency quantification. The BERTAgent model is available as a Python package [69].

As shown in Fig. 4, for each marker we performed a t-test on tweets to compare the mean values before and after the main event. In the #MeToo discourse, the average per-tweet levels of agency and affiliation markers increased over time, whereas the change of anger was negligible. A different pattern can be observed for #FridaysForFuture as the only marker that was reliably changing over time was affiliation. All markers had instead a significant increase in the #Covid19 outbreak.

5.3 Study #1: PLMP marker increase

A first application of PLMP is displayed in Fig. 5 with respect to the markers of affiliation, agency, and anger for the datasets of #MeToo, #FridaysForFuture, and #Covid19. The semantic networks displayed in Fig. 5 refer to tweets published after the main event, including nouns and verbs for the sake of readability, and colours highlighting the increase of the marker with respect to the semantic network built on tweets published prior to the main event. Moreover, in Fig. 5 the size of nouns and verbs is proportional to their PageRank value, i.e., their centrality in the online discourse, while their position relates horizontally to the strength in the increase of the marker, and vertically

Fig. 5. Wordclouds of the PLMP increase in agency, affiliation, and anger in #MeToo, #FridaysForFuture, and #Covid19 – darker colours identify larger increases; font size proportional to PageRank centrality; PageRank centrality of "woman" was bounded to allow for improved readability. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

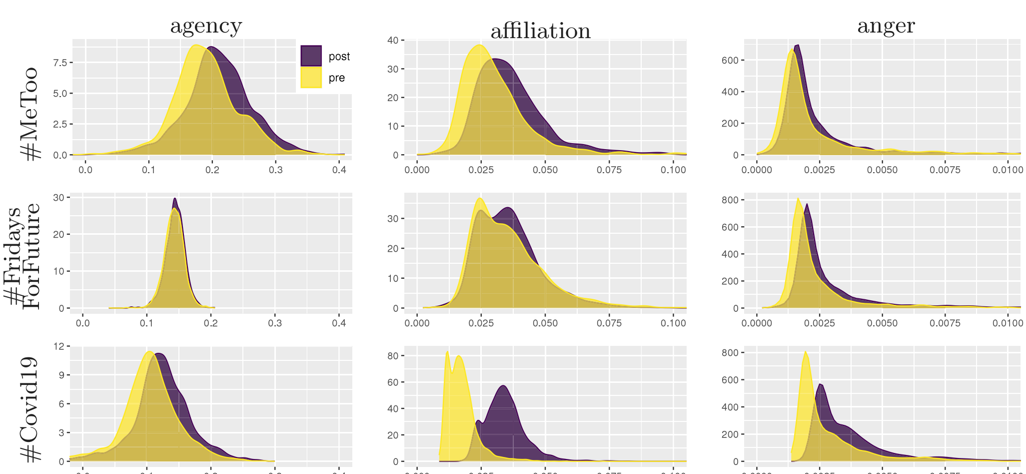

Fig. 6. Density plots of PLMP projected agency, affiliation, and anger on words – pre and post the main event.

to their PageRank centrality strength. Nouns and verbs on the right (left) refer to a stronger increase (decrease) in the marker, while nouns and verbs above (below) refer to more (less) central nodes. This is also reflected by node size (PageRank) and colour (marker increase). Fig. 6 further investigates the density of PLMP projected markers values pre and post the main event, showing agreement with the statistics of Fig. 4, thus certifying the adequacy of the PLMP method to preserve the content of tweets $\tilde{\boldsymbol{m}}_t$.

Fig. 5 shows how these increases relate to words. For #MeToo, the increase pertaining to markers of collective action is strongly associated with the central words of the discourse (e.g., *girl*, *woman*, *violence*, *equality*, *join*, *speak*, *right*, all represented with a darker colour), evident especially for agency and affiliation. For #FridaysForFuture most central words are not generally heavily associated with an increase in the markers (see their brighter colours). #Covid19 shows a general increase in all of the three markers, uniformly spread irrespective of the word centrality (see the dark colours in the entire word cloud).

This variation is further quantified in Fig. 7. #MeToo shows a dependency between the PageRank centrality of words and their marker increase, as can be inferred from the increasing yellow linear regression lines (increasing especially for agency). For the feminist call #MeToo the most central words in the discourse acquire weight in terms of agency, affiliation, and anger, i.e., they are very prominent hubs embedded in the socio-psychological discourse. Changes between the pre-

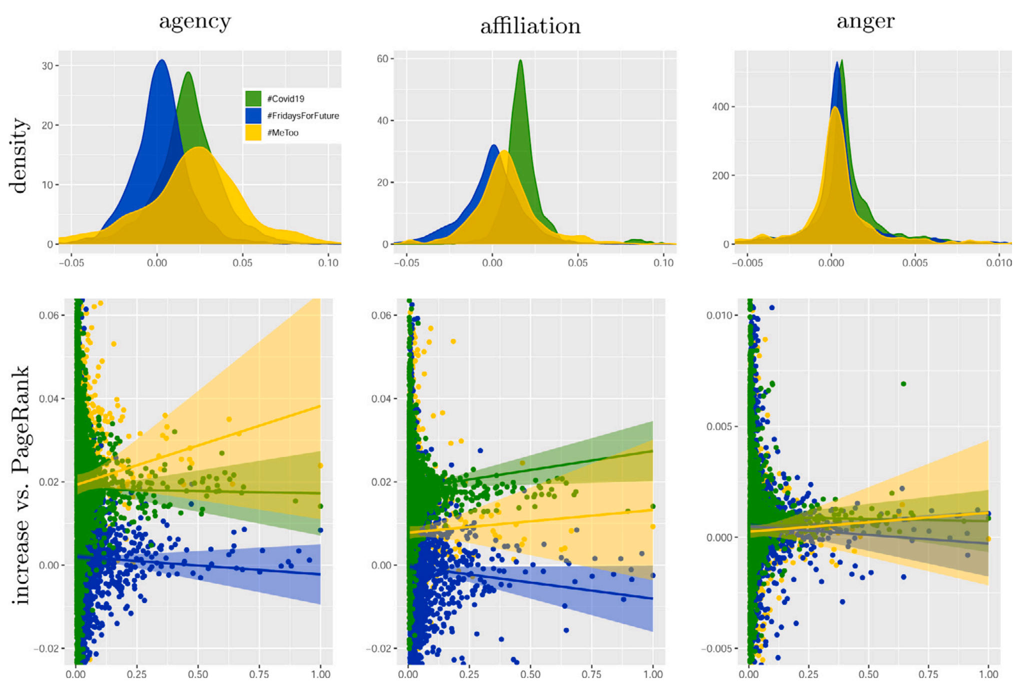

Fig. 7. PLMP markers increase density (above) and increase versus PageRank (below).

and post-event PLMP projected values reflect changes in the discourse, pertaining to socio-psychological concepts under investigation. By contrast, the environmentalist call to action #FridaysForFuture and the #covid19 dataset are not characterized by a strong positive relation between PageRank and collective action markers, i.e., their linear regression lines tend to be horizontal or negatively correlated. This suggests that the evolution of these discourses cannot be traced down to the investigated linguistic markers in relation to specific words. #FridaysForFuture apparently contains a sparser rhetoric that concentrates the increase in agency and affiliation at the periphery of the speech; thus, slightly decreasing linear regression lines indicate stronger relevance for the words with lower PageRank centrality.

These findings enable a finely tuned characterization of the rhetoric of movements, which may offer original insights into social influence phenomena. The feminist discourse appears to be focusing on fewer key target words, with the discussion concentrating on themes that feature relevant emotional and psycho-social qualities. Differently, the climate action call addresses many different topics, and their psychological characterization (in terms of collective action features) is sparser. This finding can be correlated to the Elaboration Likelihood Model [70], according to which people's style of argumentation varies with the relevance of the topic. When people are highly involved in a topic (the self-relevance in the #MeToo is signalled even in the name of the movement), fewer arguments are more persuasive. However, a high quantity of argumentation is more effective when the personal relevance of the topic is lower.

Fig. 5 also offers an overview of the discourse about the pandemic. We can infer that the discussion on #Covid19 is evolving around several important words, which are very diverse in meaning, suggesting that different topics are guiding the social exchange. While this diversity in topics is further investigated in Study #2, it is worthwhile noting that the main keywords evolving over time prove the PLMP usefulness to capture dynamics in a social discourse, likely reflecting the natural evolution of movements and the reality they pertain to. Specifically, we can see that increasing agency is featuring the main actors involved in the event, namely *workers* and *doctors*, but also the *pandemic* itself, and words prompting specific behaviours including *fight*, *test*, *help*, *stay*, *home*. The importance of an increased sense of community and affiliation is carried by semantically congruent keywords, including *help*, *share*, *community*, *family*, and *support*. Finally, the graphical representation of the relevance of anger confirms that this dimension is not very central to the online discussion, and does not characterize a strong dynamic, in line with the fact that the effect size of anger's change is small, as reported in Fig. 4. Interestingly, there is indeed one outstanding word that is not only an important element of the semantic network (as emphasized by its high PageRank) but also carrying an increase in anger, namely the word *fight*, which possibly signals a mobilization of the negative emotion within a specific subgroup of the social exchange. This insight is further investigated in Study #2.

5.4 Study #2: PLMP similarity ranking

In this second study, tweets related to #Covid19 are investigated under a different lens. Van Zomeren et al. [71] suggest that a shared identity can be built through group efficacy, calling for a unitary assessment of the architectural structure of collective action. Along this line, we show that PLMP projection can provide a valid practical tool for operationalizing this holistic view, thanks to its ability to project one feature (e.g., agency) onto the other (e.g., we-ness). In this context, we also test the ability of PMLP to spot topics in the online discussion that specifically call for collective action, so as to offer a more nuanced approach to the psychological analysis of the discourse.

Our choice is to focus on the word "we", because social identification with an ingroup is well signalled by the use of this first-person plural pronoun [72], and this word already proved to be a reliable predictor of online collective actions [73], such as #OccupyWallStreet [74]. Importantly, by projecting anger and agency on the target word "we", one can inspect the interconnection of psycho-social features in online collective action. Such an analysis would be in line with a theoretical perspective advanced in the social psychological literature, which so far has little empirical actualization.

The online discourse regarding #Covid19 is displayed in Fig. 8, which shows the result of the Louvain community detection algorithm [4,75–78] on the bipartite network linking tweets to words and hashtags. Only the words related to the seven most relevant topics (those greater than 5%) are shown in Fig. 8, capturing 80% of the full

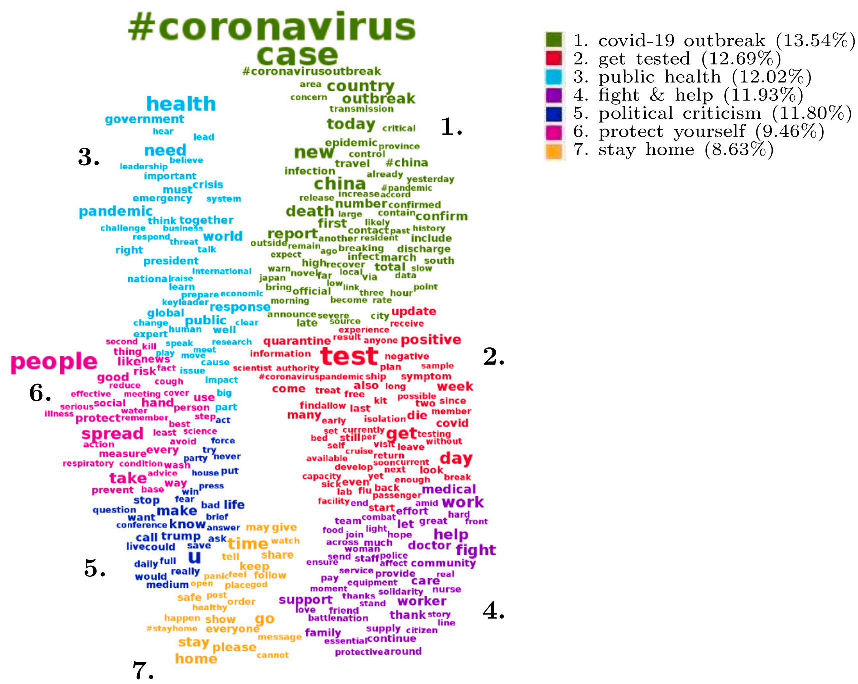

Fig. 8. #Covid19 topics – by Louvain community detection.

semantic network. Even though Fig. 8 displays only words, through this approach we can easily identify, for each topic, both words as well as tweets that pertain to it.

The influence of topics on the target word "we" is measured by resorting to a generalization of the SimRank approach, originally introduced by Haveliwala [51]. In more detail, we exploit the diffusion of PLMP for measuring correlation by activating the vector $\boldsymbol{q}$ in (2) only within a subset of nodes belonging to a specific topic. In case $\boldsymbol{q}$ is only active on a selected tweet i, we measure the influence of the ith tweet on a specific word (e.g., the word "we" at position k), by reading the resulting PLMP vector $\boldsymbol{m}$ at the specific word entry (i.e., k). The influence of a topic, instead, consists of the union (i.e., the sum) of all the contributions from the tweets belonging to the topic.[2] Formally, for each tweet i, the above is equivalent to solving

$$\boldsymbol{m}_i = \alpha \boldsymbol{M} \boldsymbol{m}_i + (1-\alpha) \begin{bmatrix} \boldsymbol{0} \\ \boldsymbol{\delta}_i \end{bmatrix} \tilde{m}_{t,i} \,, \tag{11}$$

where $\boldsymbol{\delta}_i$ is a binary vector active only in position i, and where $\tilde{m}_{t,i}$ corresponds to the ith entry of the agency-in-isolation tweet vector $\tilde{\boldsymbol{m}}_t$. The influence of the jth topic on the word "we"(assumed to be at position k) is then obtained via

$$i_{j,k} = \sum_{i \in \mathcal{T}_j} m_{i,k}$$

where $\mathcal{T}_j$ collects the tweets of topic j, and $m_{i,k}$ is the kth entry of $\boldsymbol{m}_i$. Interestingly, separately knowing the contribution of each tweet, $m_{i,k}$, allows us to apply statistical tests on the increase/decrease of the markers of interest. This is illustrated in Fig. 9, showing, for agency and anger, the statistical variation within each topic from two perspectives, namely: (1) the evaluation of markers on *tweets* for different topics, and (2) an improved investigation that exploits the topic-driven PLMP projection on "we". Statistically relevant changes are highlighted by arrows.

By analysing the psychological characterization in the topics of Fig. 8, and in particular the projection of agency on the target word "we", it is possible to see the social development of the online discourse over time and across specific topics, and capture trends that cannot be spotted without the PLMP projection. Agency is generally increasing from pre to post, but not in every community, and not every word is featured by agency. In particular, while the increase in agency is observable as a general trend in tweets from every community (Fig. 9), the intersection of ingroup salience and agency captured by the projection of agency on the keyword "we" is specifically emerging in four communities (Fig. 9). It is worth noting that these four communities are exactly those characterizing online mobilization, as they either prompt behaviours for the collective good (*get tested*, *fight and help*, *stay home*) or share content that challenges the government (*political criticism*). The PLMP projection of anger on "we" is particularly telling, as the general analyses of anger on the tweets would lead to the gross conclusion that anger is mainly irrelevant, as the only community featuring an increase in this emotion regards getting tested, and possibly is capturing the obvious anger emerging after a positive outcome of the test. PLMP projection allows for a fine-tuned understanding of this psychological feature, which we have already introduced as the most controversial element of collective action. Importantly, we can conclude that an increase of anger on the word "we" appears only in three communities (*fight and help*, *political criticism*, and *stay home*), all being collective action communities. Notably, the community *covid-19 outbreak* is characterized by a decrease of both anger and agency, possibly suggesting a general reduction of the psychological resources in association with the discourse focusing on the data and factual reports signalling the worsening of the pandemic.

6 Conclusions

We introduced PLMP, a method for extrapolating holistic information from a semantic network. PLMP maps linguistic markers onto the network elements, i.e., the words, exploiting a PageRank-like rationale to track the network interdependencies, hence better capturing the subtleties of the social discourse. Its application to the study of calls to collective action revealed that PLMP is able to characterize many call-specific aspects, providing a scented appraisal of the dynamic evolution

[2] In the PageRank original setup, the weight of multiple tweets ought to be averaged, since the mixing matrix is column-normalized and the teleport vector is stochastic. Here, instead, because of the averaging action of matrix $\boldsymbol{M}$ in (2), we must consider their sum.

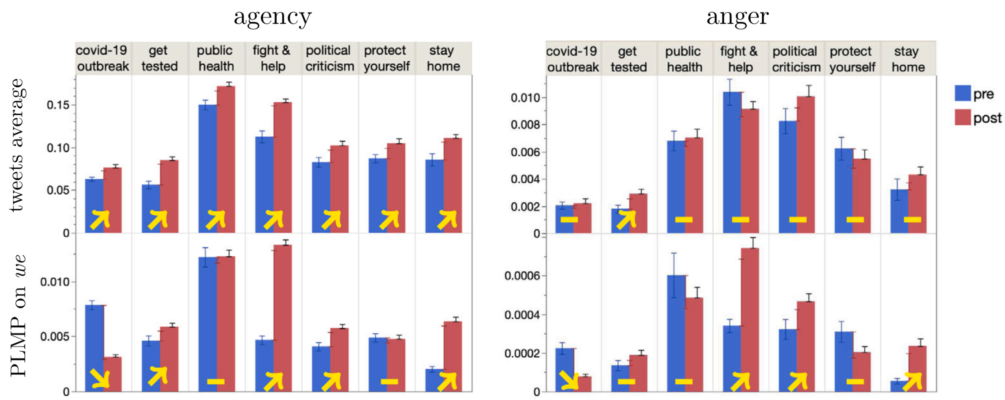

Fig. 9. #Covid19 topics perspective on agency (left) and anger (right) – marker values are displayed as the average of the *tweets* belonging to a community (above), or as the aggregate *PLMP* value separately projected on *we* for each community (below); yellow arrows indicate a statistically relevant increase or decrease (arrow up or down, respectively), while a horizontal dash indicates the absence of a statistically relevant change. Histograms and bars on top represent the average value and the standard error, respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

of the social discourse. A relevant application of the PLMP projection rests on its ability to capture the interplay of the psychological features that characterize social discourse. Applied to the theoretical and practical field of collective action, the projection of agency and anger on ingroup provides relevant insights for the boundary conditions of the evolution of social processes for example confirming agency as a key dimension for collective discourse, and anger being mobilized mainly in specific calls for actions where a key enemy is targeted. We also showed how PLMP is able to keep into account different aspects of individual and in-context meaning of the words, achieving statistical reliability, thus confirming the validity of our approach. Future studies may apply the proposed framework to different datasets, possibly focusing on various social matters and/or worldwide events with a similar variegate discourse that requires carefully targeted interpretation.

CRediT authorship contribution statement

Tomaso Erseghe: Conceptualization, Writing – review & editing, Formal analysis, Methodology, Visualization. **Leonardo Badia:** Conceptualization, Writing – review & editing, Methodology, Validation. **Lejla Džanko:** Conceptualization, Writing – review & editing, Data curation, Visualization. **Magdalena Formanowicz:** Conceptualization, Writing – review & editing, Methodology, Validation. **Jan Nikadon:** Conceptualization, Writing – review & editing, Data curation, Software. **Caterina Suitner:** Conceptualization, Writing – review & editing, Formal analysis, Methodology, Validation.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Funding

This work was supported by the University of Padova, Italy.

Figure generation

All figures were created by RStudio [79], except for Fig. 8 which was created by Gephi [80], and Fig. 9 by JMP [81]. Twitter data was collected using the Twitter API [82].

Code availability

The code used to process the data is available from the corresponding author on reasonable request.

References

[1] T. Erseghe, L. Badia, L. Džanko, C. Suitner, PLMP: A Method to map the linguistic markers of the social discourse onto its semantic network, in: Proc. ASONAM, 2022.

[2] D. Boyd, K. Crawford, Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon, Inf. Commun. Soc. 15 (5) (2012) 662–679.

[3] K. Gimpel, N. Schneider, B. O'Connor, D. Das, D. Mills, J. Eisenstein, M. Heilman, D. Yogatama, J. Flanigan, N. Smith, Part-of-speech tagging for Twitter: Annotation, features, and experiments, in: Ann. Meet. Ass. Computat. Ling., Vol. 2, 2011, pp. 42–47.

[4] C. Suitner, L. Badia, D. Clementel, L. Iacovissi, M. Migliorini, B.G.S. Casara, D. Solimini, M. Formanowicz, T. Erseghe, The rise of# climateaction in the time of the FridaysForFuture movement: A semantic network analysis, Social Networks (2022) URL https://www.sciencedirect.com/science/article/pii/S0378873322000570.

[5] B. Ozkazanc-Pan, On agency and empowerment in a# MeToo world, Gend. Work Organ. 26 (8) (2019) 1212–1220.

[6] M. van Zomeren, M. Kutlaca, F. Turner-Zwinkels, Integrating who "we" are with what "we"(will not) stand for: A further extension of the Social Identity Model of Collective Action, Eur. Rev. Soc. Psychol. 29 (1) (2018) 122–160.

[7] E.F. Thomas, K.I. Mavor, C. McGarty, Social identities facilitate and encapsulate action-relevant constructs: A test of the social identity model of collective action, Group Process. Intergroup Relat. 15 (1) (2012) 75–88.

[8] A. Zancanaro, G. Cisotto, L. Badia, Modeling value of information in remote sensing from correlated sources, in: Proc. IEEE MedComNet, 2022, pp. 47–53.

[9] L. Badia, D. Clementel, T. Erseghe, L. Iacovissi, M. Migliorini, B.G. Salvador Casara, C. Suitner, Structural and semantic impact of online collective action, in: NetSci Conf., 2020, URL https://www.dei.unipd.it/~badia/papers/2020_09_NetSci.pdf.

[10] L. Page, S. Brin, R. Motwani, T. Winograd, The PageRank Citation Ranking: Bringing Order to the Web, Tech. rep., Stanford, 1999.

[11] D.F. Gleich, PageRank beyond the Web, SIAM Rev. 57 (3) (2015) 321–363.

[12] T. Haveliwala, S. Kamvar, The Second Eigenvalue of the Google Matrix, Tech. rep., Stanford, 2003, URL http://www-cs-students.stanford.edu/~taherh/papers/secondeigenvalue.pdf.

[13] D. Osborne, J.T. Jost, J.C. Becker, V. Badaan, C.G. Sibley, Protesting to challenge or defend the system? A system justification perspective on collective action, Eur. J. Soc. Psychol. 49 (2) (2019) 244–269.

[14] M. van Zomeren, T. Postmes, R. Spears, Toward an integrative social identity model of collective action: A quantitative research synthesis of three socio-psychological perspectives, Psychol. Bull. 134 (4) (2008) 504.

[15] S. Stürmer, B. Simon, Collective action: Towards a dual-pathway model, Eur. Rev. Soc. Psychol. 15 (1) (2004) 59–99.

[16] F.G. Neville, A. Templeton, J.R. Smith, W.R. Louis, Social norms, social identities and the COVID-19 pandemic: Theory and recommendations, Soc. Personal. Psychol. Compass 15 (5) (2021) e12596.

[17] J.D. McCarthy, M.N. Zald, Resource mobilization and social movements: A partial theory, Am. J. Sociol. 82 (6) (1977) 1212–1241.

[18] B. Klandermans, Mobilization and participation: Social-psychological expansisons of resource mobilization theory, Am. Sociol. Rev. (1984) 583–600.
[19] A. Bandura, Exercise of human agency through collective efficacy, Curr. Dir. Psychol. Sci. 9 (3) (2000) 75–78.
[20] V.L. Spring, C.D. Cameron, M. Cikara, The upside of outrage, Trends Cogn. Sci. 22 (12) (2018) 1067–1069.
[21] W.J. Brady, M.J. Crockett, How effective is online outrage? Trends Cogn. Sci. 23 (2) (2019).
[22] K. Li, Y. Xu, S. Yang, Y. Guo, Social class, group-based anger, and collective action intentions in China, J. Pac. Rim Psychol. 13 (2019) e13.
[23] A. Goldenberg, J.J. Gross, Digital emotion contagion, Trends Cogn. Sci. 24 (4) (2020) 316–328.
[24] S.M. Mohammad, P.D. Turney, Crowdsourcing a word–emotion association lexicon, Comput. Intell. 29 (3) (2013) 436–465.
[25] M. Brysbaert, A.B. Warriner, V. Kuperman, Concreteness ratings for 40 thousand generally known English word lemmas, Behav. Res. Methods 46 (2014) 904–911.
[26] S. Bhatia, L. Walasek, Event construal and temporal distance in natural language, Cognition 152 (2016) 1–8.
[27] D. Gaucher, J. Friesen, A.C. Kay, Evidence that gendered wording in job advertisements exists and sustains gender inequality., J. Personal. Soc. Psychol. 101 (1) (2011) 109.
[28] Y.R. Tausczik, J.W. Pennebaker, The psychological meaning of words: LIWC and computerized text analysis methods, J. Lang. Soc. Psychol. 29 (1) (2010) 24–54.
[29] R.C. Hawkins II, R.L. Boyd, Such stuff as dreams are made on: Dream language, LIWC norms, and personality correlates, Dreaming 27 (2) (2017) 102–121.
[30] J. Hirschberg, C.D. Manning, Advances in natural language processing, Science 349 (6245) (2015) 261–266.
[31] Y. Bae, H. Lee, Sentiment analysis of twitter audiences: Measuring the positive or negative influence of popular twitterers, J. Am. Soc. Inf. Sci. Technol. 63 (12) (2012) 2521–2535.
[32] J. Patro, R. Bhaskaran, A. Mukherjee, What propels celebrity follower counts? Language use or social connectivity, 2018, arXiv preprint arXiv:1811.07169.
[33] A. Fong, J. Roozenbeek, D. Goldwert, S. Rathje, S. van der Linden, The language of conspiracy: A psychological analysis of speech used by conspiracy theorists and their followers on Twitter, Group Process. Intergroup Relat. 24 (4) (2021) 606–623.
[34] R. Gulliver, K.S. Fielding, W.R. Louis, Assessing the mobilization potential of environmental advocacy communication, J. Environ. Psychol. 74 (2021) 101563.
[35] S. Alaparthi, M. Mishra, BERT: A sentiment analysis odyssey, J. Mark. Anal. 9 (2) (2021) 118–126.
[36] R. Catelli, S. Pelosi, M. Esposito, Lexicon-based vs. Bert-based sentiment analysis: A comparative study in Italian, Electronics 11 (3) (2022) 374.
[37] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, Adv. Neural Inf. Process. Syst. 30 (2017).
[38] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, 2018, arXiv preprint arXiv:1810.04805.
[39] Z. Gao, A. Feng, X. Song, X. Wu, Target-dependent sentiment classification with BERT, IEEE Access 7 (2019) 154290–154299.
[40] M. Hoang, O.A. Bihorac, J. Rouces, Aspect-based sentiment analysis using bert, in: Proceedings of the 22nd Nordic Conference on Computational Linguistics, 2019, pp. 187–196.
[41] M. Munikar, S. Shakya, A. Shrestha, Fine-grained sentiment classification using BERT, in: 2019 Artificial Intelligence for Transforming Business and Society (AITB), Vol. 1, IEEE, 2019, pp. 1–5.
[42] M. Singh, A.K. Jakhar, S. Pandey, Sentiment analysis on the impact of coronavirus in social life using the BERT model, Soc. Netw. Anal. Min. 11 (1) (2021) 33.
[43] J. Nikadon, C. Suitner, T. Erseghe, L. Džanko, M. Formanowicz, BERTAgent: The development of a novel tool to quantify agency in textual data, J. Exp. Psychol.: Gen. (2023) submitted for publication. URL https://psyarxiv.com/qw6u3/, available online.
[44] A.F. Adoma, N.-M. Henry, W. Chen, Comparative analyses of bert, roberta, distilbert, and xlnet for text-based emotion recognition, in: 2020 17th International Computer Conference on Wavelet Active Media Technology and Information Processing (ICCWAMTIP), IEEE, 2020, pp. 117–121.
[45] A.R. Abas, I. Elhenawy, M. Zidan, M. Othman, BERT-CNN: A deep learning model for detecting emotions from text, Comput. Mater. Contin. 71 (2) (2022).
[46] Y. Xiong, M. Cho, B. Boatwright, Hashtag activism and message frames among social movement organizations: Semantic network analysis and thematic analysis of Twitter during the# MeToo movement, Public Relat. Rev. 45 (1) (2019) 10–23.
[47] W. Shi, H. Fu, P. Wang, C. Chen, J. Xiong, Climatechange vs. globalwarming: Characterizing two competing climate discourses on Twitter with semantic network and temporal analyses, Int. J. Environ. Res. Public Health 17 (3) (2020) 1062.
[48] M. Grootendorst, BERTopic: Neural topic modeling with a class-based TF-IDF procedure, 2022, arXiv preprint arXiv:2203.05794.
[49] S.D. Kamvar, T.H. Haveliwala, C.D. Manning, G.H. Golub, Extrapolation methods for accelerating PageRank computations, in: Proceedings of the 12th International Conference on World Wide Web, 2003, pp. 261–270.
[50] S.C. Murphy, A hands-on guide to conducting psychological research on Twitter, Soc. Psychol. Personal. Sci. 8 (4) (2017) 396–412.
[51] T.H. Haveliwala, Topic-sensitive PageRank: A context-sensitive ranking algorithm for web search, IEEE Trans. Knowl. Data Eng. 15 (4) (2003) 784–796.
[52] J. Weng, E.-P. Lim, J. Jiang, Q. He, Twitterrank: finding topic-sensitive influential twitterers, in: Proceedings of the Third ACM International Conference on Web Search and Data Mining, 2010, pp. 261–270.
[53] X. He, M. Gao, M.-Y. Kan, D. Wang, Birank: Towards ranking on bipartite graphs, IEEE Trans. Knowl. Data Eng. 29 (1) (2016) 57–71.
[54] L. Cao, J. Guo, X. Cheng, Bipartite graph based entity ranking for related entity finding, in: 2011 IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology, Vol. 1, IEEE, 2011, pp. 130–137.
[55] J.M. Kleinberg, Authoritative sources in a hyperlinked environment, J. ACM 46 (5) (1999) 604–632.
[56] A. Abele, B. Wojciszke, Agency and Communion in Social Psychology, Routledge, 2018.
[57] M. Jacomy, T. Venturini, S. Heymann, M. Bastian, ForceAtlas2, a continuous graph layout algorithm for handy network visualization designed for the Gephi software, PLoS One 9 (6) (2014) e98679.
[58] A.V. Guglielmi, L. Badia, Social communication to improve group recognition in mobile networks, in: Proc. IEEE Globecom Wkshps, 2015.
[59] T. Erseghe, A distributed and scalable processing method based upon ADMM, IEEE Signal Process. Lett. 19 (9) (2012) 563–566.
[60] Y. Fan, M. Li, P. Zhang, J. Wu, Z. Di, The effect of weight on community structure of networks, Physica A 378 (2) (2007) 583–590.
[61] T. Zhou, J. Ren, M. Medo, Y.C. Zhang, Bipartite network projection and personal recommendation, Phys. Rev. E 76 (4) (2007) 046115.
[62] F.R. Gantmacher, The Theory of Matrices, Vol. 2, AMS Chelsea Publishing Company, New York, 2000.
[63] F.R. Gantmacher, The Theory of Matrices, Chelsea Publishing Company, New York, 1977.
[64] I. Hellsten, L. Leydesdorff, Automated analysis of actor–topic networks on Twitter: New approaches to the analysis of socio-semantic networks, J. Assoc. Inf. Sci. Technol. 71 (1) (2020) 3–15.
[65] O. Owoputi, B. O'Connor, C. Dyer, K. Gimpel, N. Schneider, N. Smith, Improved part-of-speech tagging for online conversational text with word clusters, in: Proc. NAACL-HLT, 2013, pp. 380–390.
[66] M. van Zomeren, A. Iyer, Introduction to the social and psychological dynamics of collective action, J. Soc. Issues 65 (4) (2009) 645–660.
[67] J.W. Pennebaker, R.L. Boyd, K. Jordan, K. Blackburn, The Development and Psychometric Properties of LIWC2015, Tech. rep., University of Texas at Austin, 2015.
[68] S. González-Carvajal, E.C. Garrido-Merchán, Comparing BERT against traditional machine learning text classification, 2020, arXiv preprint arXiv:2005.13012.
[69] J. Nikadon, Python BERTAgent package, 2023, URL https://github.com/EnchantedStardust/BERTAgent-Quick-Start.
[70] R.E. Petty, J.T. Cacioppo, The effects of involvement on responses to argument quantity and quality: Central and peripheral routes to persuasion, J. Personal. Soc. Psychol. 46 (1) (1984) 69.
[71] M. Van Zomeren, C.W. Leach, R. Spears, Does group efficacy increase group identification? Resolving their paradoxical relationship, J. Exp. Soc. Psychol. 46 (6) (2010) 1055–1060.
[72] J.W. Pennebaker, M.R. Mehl, K.G. Niederhoffer, Psychological aspects of natural language use: Our words, our selves, Annu. Rev. Psychol. 54 (1) (2003) 547–577.
[73] J. Adam-Troian, E. Bonetto, T. Arciszewski, "We shall overcome": First-person plural pronouns from search volume data predict protest mobilization across the United States, Soc. Psychol. Personal. Sci. 12 (8) (2021) 1476–1485.
[74] L.G. Smith, J. Gavin, E. Sharp, Social identity formation during the emergence of the occupy movement, Eur. J. Soc. Psychol. 45 (7) (2015) 818–832.
[75] V.D. Blondel, J.-L. Guillaume, R. Lambiotte, E. Lefebvre, Fast unfolding of communities in large networks, J. Stat. Mech.: Theory Exp. 2008 (10) (2008) P10008.
[76] R. Lambiotte, J.-C. Delvenne, M. Barahona, Laplacian dynamics and multiscale modular structure in networks, 2008, arXiv preprint arXiv:0812.1770.
[77] A. Lancichinetti, S. Fortunato, Community detection algorithms: a comparative analysis, Phys. Rev. E 80 (5) (2009) 056117.
[78] S. Fortunato, Community detection in graphs, Phys. Rep. 486 (3–5) (2010) 75–174.
[79] RStudio Team, RStudio: Integrated Development Environment for R, RStudio, PBC., Boston, MA, 2020, URL http://www.rstudio.com.
[80] M. Bastian, S. Heymann, M. Jacomy, Gephi: An open source software for exploring and manipulating networks, 2009, URL http://gephi.org.
[81] JMP, JMP©Pro 17.0.0, JMP, Cary, NC, 2022, URL https://www.jmp.com.
[82] Twitter Developer Platform, Twitter API, Twitter, San Francisco, CA, 2021, URL http://developer.twitter.com/en/docs/twitter-api.