



All Polarized but Still Different: a Multi-factorial Metric to Discriminate between Polarization Behaviors on Social Media

Celina Treuillier, Sylvain Castagnos, Armelle Brun

► To cite this version:

Celina Treuillier, Sylvain Castagnos, Armelle Brun. All Polarized but Still Different: a Multi-factorial Metric to Discriminate between Polarization Behaviors on Social Media. ACM/SIGAPP Symposium On Applied Computing, Apr 2024, Avila, Spain. 10.1145/3605098.3635972 . hal-04327401

HAL Id: hal-04327401

<https://hal.science/hal-04327401>

Submitted on 6 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

ALL POLARIZED BUT STILL DIFFERENT: A MULTI-FACTORIAL METRIC TO DISCRIMINATE BETWEEN POLARIZATION BEHAVIORS ON SOCIAL MEDIA

Céline TREUILLIER

Université de Lorraine, CNRS, LORIA
Nancy, FRANCE
celina.treuillier@loria.fr

Sylvain CASTAGNOS

Université de Lorraine, CNRS, LORIA
Nancy, FRANCE
sylvain.castagnos@loria.fr

Armelle BRUN

Université de Lorraine, CNRS, LORIA
Nancy, FRANCE
armelle.brun@loria.fr

December 6, 2023

ABSTRACT

Online polarization has attracted the attention of researchers for many years. Its effects on society are a cause for concern, and the design of personalized depolarization strategies appears to be a key solution. Such strategies should rely on a fine and accurate measurement, and a clear understanding of polarization behaviors. However, the literature still lacks ways to characterize them finely. We propose GRAIL, the first individual polarization metric, relying on multiple factors. GRAIL assesses these factors through entropy and is based on an adaptable Generalized Additive Model. We evaluate the proposed metric on a Twitter dataset related to the highly controversial debate about the COVID-19 vaccine. Experiments confirm the ability of GRAIL to discriminate between polarization behaviors. To go further, we provide a finer characterization and explanation of the identified behaviors through an innovative evaluation framework.

1 Introduction

Technology has an increasing impact on our everyday life. Its role, associated with the democratization of social media, is more and more questioned. Online media gradually became a persuasive technology through which users can be manipulated, hence undermining social harmony [1]. Digital citizenship [2] promotes education on the use of technologies so that citizens are aware of and able to deal with potential threats. Among them, polarization appears as a key societal issue [3]. In this work, we refer to polarization as political divergence, or more globally, a division into several sharply contrasting groups of opinions. Polarization is widely studied by researchers in various fields, agreeing on the possible harmful effects on society, and the need to limit it in certain cases [4].

Social conflicts sometimes lead to debate and constructive change, provided that everyone is free to express their opinions [5]. However, under certain circumstances, it can lead to a bimodal distribution of opinions and ideological extremes [6].

To limit user polarization, common depolarizing processes bring global diversity into the recommended content [7, 8, 9] to expose users to a broader range of opinions. However, the need for personalized depolarizing strategies has recently been highlighted in literature [10, 5, 11]. The literature does offer many polarization metrics, assessing the global polarization of a group of users, a specific topic, or within a whole network [12, 13, 14]. Besides, only few individual metrics have been proposed [15, 16]. However, these existing metrics, both global and individual, only rely on a single factor and struggle to go beyond a simple binary paradigm. For example, some studies focus on the diversity of

communities with which users interact [12, 15], while others study the diversity of sources [16], but they never combine these multiple factors to their full extent. Paradoxically, the literature is fairly consistent in pointing out that multiple factors, both human and algorithmic, drive polarization [17, 18, 19]. In our view, the implementation of personalized depolarization strategies should thus rely on a fine-grained and multi-factorial assessment of polarization. A recent study carried out in a social media context, strengthened the interest of a multi-factorial approach in modeling and understanding polarization behaviors [20].

Based on these findings, we advocate the need for a new polarization metric, going beyond the existing ones by assessing the polarization of each individual independently and taking into account the multiple factors that influence this polarization. In that respect, we raise two research questions: **(RQ1) How to combine multiple factors to form a discriminative individual polarization metric?** and **(RQ2) Does the proposed metric allow to explain polarization behaviors?**

We propose the GRAIL (GeneRalized AddItive poLarization) metric. GRAIL is evaluated on a social media dataset related to the COVID-19 vaccine debate on Twitter [21], and its ability to be explained by polarization-specific behavioral indicators for different classes of behavior is assessed. The contributions of this work are thus as follows: **(1)** the definition of GRAIL, the first individual and multi-factorial polarization metric, and **(2)** an innovative evaluation framework allowing the explanation of behavioral classes based on individual polarization scores computed with GRAIL.

The remainder of this paper is organized as follows. In Section 2, we provide a literature review of polarization assessment. We then detail the GRAIL metric in Section 3. The evaluated data is presented in Section 4, and GRAIL is evaluated in Section 5. Finally, conclusions and perspectives are drawn in Section 6.

2 Literature Review

2.1 Online Media and Polarization

Operating in many contexts (online/offline, individual/group), polarization has raised the attention of a large variety of fields, *e.g.* mathematicians, physicians, psychologists, computer scientists, . . . Among other things, many mathematical models of opinion dynamics and polarization have been described for years [22]. Physicians such as Baumann *et al.* [23] show that three scenarios occur in a social media context: consensus, unilateral polarization, or bilateral polarization. Besides, psychologists, *e.g.* Geschke *et al.*, propose a three-dimension modeling of polarization: individual, social and technological [17]. The technological dimension is reinforced by the daily use of online media, which has considerably changed the way people access the latest news and get information. Two main types of online content broadcasters can be distinguished: online media and social media.

First, online media and news aggregators offer users the opportunity to access a large amount of information every day. This led to the emergence of News Recommender Systems (NRS), which assist users in finding news of interest. However, NRS have also fostered individual polarization [24, 25] by creating filter bubbles [26, 27]. The role of diversity in NRS regarding polarization decrease is highly discussed and even highlights a potential strengthening of polarisation [28]. Besides, Wu *et al.* claim that the diversification process should be personalized for each user according to her behavior [29]. In our view, this highlights the need for a comprehensive and individual modeling of polarization.

Second, social media, which are the focus of this work, are now commonly used by a large proportion of the population and are undoubtedly involved in the polarization process [4]. The following section is dedicated to social media.

2.2 Polarization Metrics on Social Media

Social media have been democratized since the late 2000s. They are now used on a daily basis all around the world. Although social media have strong benefits such as liberty of interactions, free access to information and democratization, they also play a critical role in the polarization process [4, 30]. From the very first steps of Twitter, Conover *et al.* studied polarization on social media [31]. They have shown that public political interaction (mentions and retweets) induces distinct network topologies. The *retweets* network allows to distinguish two well-separated communities, thus showing that there is a high polarization, whereas the *mentions* network results in a unique highly connected community. The authors rely on modularity to quantify polarization [32], thus capturing the strength of the division of a network into modules (communities). While giving some insights about the division of a network, Guerra *et al.* explain that modularity does not seem to be a direct measure of polarization, as non-polarized networks can also be divided into distinct communities [13]. The authors also emphasize the potential role of community boundaries to quantify polarization, by comparing internal and external connections of intermediary nodes. Polarization can also be evaluated

through the *controversy* of a topic [12], *i.e.* its propensity to provoke heated debates. Highly controversial topics automatically lead to well-separated communities in the network, and *controversy* can thus be measured, for example, by examining the information flow between communities. Morales *et al.* propose to model opinions shared in a social media by the probability density function: polarization is then quantified according to the resulting distribution [14]. The authors also propose to consider attractors, *i.e.* a concentration of users around different opinion positions, in polarization metrics [14]. We can conclude that polarization metrics from the social media literature rely on a varied set of information. Though, these measures share one major characteristic: they all quantify polarization or controversiality of a whole network or topic. In the form they are presented, existing polarization metrics thus provide an overview of the state of polarization in a network or about a specific topic, but give no information about the level of polarization of each user, individually.

Recall that polarization is influenced by both individual and group factors [17]. We thus promote individual and group polarization metrics to participate in a refined model of polarization. However, such metrics remain uncommon in the literature. Among the few existing ones, we can find the *polarization score* proposed by Becatti *et al.* [15]. It is based on the differentiation of a set of communities C from the graph built from users' interactions. To evaluate this metric, a user u is represented by a vector $I_{u,C}$ where each element $I_{u,c}$ represents the proportion of interactions of u with community c ($N_{u,c}$), compared to her total number of interactions N_u . The polarization score $\rho(u)$ of user u , is evaluated as the maximum value of vector $I_{u,C}$, as presented in Equation (1).

$$\rho(u) = \max_{c \in C} \left\{ \frac{N_{u,c}}{N_u} \right\} \quad (1)$$

Users with $\rho \approx 1$ are thus users who access a unique community and are highly polarized. Users who access all the communities equally, have $\rho = 1/C$. In a two-community context, the polarization score presented by Schmidt *et al.* is similar but is oriented, in the sense that the value informs which community is accessed the most [33].

Cicchini *et al.* [16] proposed the *Lack of Diversity* (LD) metric. It is highly similar to the *polarization score* of Becatti *et al.* [15], but considers sources of information instead of communities while correcting the bias linked to the popularity of these sources. In their work, sources of information consist in a set of M media outlets, with which users can interact. Thus, each user u is represented by her number of interactions $N_{u,m}$ with news from media m . *LD* is computed as follows:

$$LD(u) = \max_{m \in M} \left\{ N_{u,m} \cdot \log\left(\frac{|U|}{|U_m|}\right) \right\} \quad (2)$$

U is the set of users of size $|U|$, while $|U_m|$ is the number of users interacting with media m . The *log* term corrects a potential bias introduced by an imbalance in the number of interactions with a specific source m . The LD metric is not bounded.

In our view, the aforementioned individual polarization metrics face some limits. First, both *polarization score* [15] and *Lack of Diversity* [16] only exploit the maximum value in each user vector. It is very reductive: two users who behave the same way in their main community/media will get the same metric value, while they may behave quite differently in other communities/media. This is not taken into account in the metrics, which we believe limits the accuracy of the modeling. Indeed, all values can bring useful insight, and they can all contribute to a better understanding of polarization behavior. Second, these metrics are only measured on a single factor (interaction with communities or interaction with sources), while polarization occurs over the influence of multiple factors [24, 34]. As a consequence, if only a single factor is considered, several users may be identified as similarly polarized according to this factor, but may in fact exhibit a wide range of behaviors and distinguish themselves according to other factors. To sum up, despite a strong community dealing with polarization, the literature still lacks an individual and multi-factorial polarization metric.

3 GRAIL: A New Individual and Multi-Factorial Polarization Metric

In this work, we propose GRAIL (*GeneRalized AddItive poLarization*), a new polarization metric designed to finely measure the degree and nature of polarization of each user. The degree corresponds to the intensity of the polarization, while the nature informs about the community a user is closer or belongs to. GRAIL has several features: **(F1)** multi-factorial, *i.e.* it can consider multiple polarization factors, **(F2)** discriminative, *i.e.* it goes beyond a polarized/non-polarized distinction and allows a fine modeling of polarization, and **(F3)** generalizable, *i.e.* it can be used on any type of data and data source.

The following subsections detail the basic components of GRAIL, that rely on 1) an information theory-related way to manage factors of any type and foster fine-grained modeling of user behavior, 2) a polynomial transformation to

heighten the discrimination between users, 3) a Generalized Additive Model to combine multiple factors. To guide the author through this section and illustrate the components that constitute the proposed metric, we punctuate our explanations with a running example.

3.1 An Entropy-Based Metric to Better Match User Behaviors

Let Z be a factor, composed of a set of entities with which users can interact. We propose to represent this factor as a probability distribution to represent the diversity/heterogeneity of a user's interactions with these entities.

Following Information Theory concepts, this diversity can be evaluated through entropy [35], a measure of the uncertainty of a data source. The more homogeneously distributed the probability mass, the higher the entropy and the greater the uncertainty. Entropy is maximum in the case of the equiprobability of entities. Conversely, if an important part of the probability mass is spread over a small set of entities, the entropy is low. The entropy is 0 if an entity has a 1 probability for one of its entities. As the maximal entropy value is a function of the number of entities, we use the normalized entropy (see Equation (3)):

$$H_N(Z) = \frac{-\sum_z P(z)\log(P(z))}{\log(n)} \quad (3)$$

Where Z is a discrete random variable that takes n possible entity values, and $P(z)$ is the probability of entity z . For the sake of simplicity, we will refer to normalized entropy as entropy. Entropy can be applied to any factor that can be seen as a random variable (topics in news, shared data sources, accounts followed, etc.), which meets characteristic *F3*.

Let us introduce here our running example. Z is a set of information sources with which users can interact and $n = 4$. According to their interactions with each of these sources, the probability distribution can be computed for each individual user and forms a probability vector of size 4, with $P(z)$ the probability that the user interacts with the source z . Table 1 presents an example of such vectors, computed for four users u_1 to u_4 , with associated probability distributions. The resulting scores computed with existing individual polarization metrics ρ [15] and LD [16], as well as the normalized entropy, defined in Equation (3), can thus be computed for each user and are given in Table 1. To simplify this example, we are assuming that the four sources z are all equally popular. The correcting term of LD (Equation (2)) is therefore equal to 1, and $\rho = LD$.

User	Probability distribution	Baselines ρ and LD	Entropy value $H_N(Z)$	Transformed entropy value $f(H_N(Z))$				
u_1	<table><tr><td>0.5</td><td>0</td><td>0.5</td><td>0</td></tr></table>	0.5	0	0.5	0	0.5	0.50	0.50
0.5	0	0.5	0					
u_2	<table><tr><td>0.5</td><td>0.2</td><td>0.2</td><td>0.1</td></tr></table>	0.5	0.2	0.2	0.1	0.5	0.88	0.73
0.5	0.2	0.2	0.1					
u_3	<table><tr><td>0.3</td><td>0.3</td><td>0.3</td><td>0.1</td></tr></table>	0.3	0.3	0.3	0.1	0.3	0.95	0.81
0.3	0.3	0.3	0.1					
u_4	<table><tr><td>0.25</td><td>0.25</td><td>0.2</td><td>0.3</td></tr></table>	0.25	0.25	0.2	0.3	0.3	0.99	0.90
0.25	0.25	0.2	0.3					

Table 1: Example of four users with different behaviors

Recall that these baseline metrics (ρ and LD) only exploit the maximum probability in each vector. The remaining probabilities, whatever their distribution, do not influence the resulting polarization measure. For example, users u_1 and u_2 have the same maximum value (equals to 0.5) so both users will be associated with the same metric value, while they interact quite differently on other sources. Regarding entropy, which considers the distribution of all entities of the vector, users u_1 and u_2 greatly differ. The use of entropy may thus lead to a more accurate metric.

3.2 A Polynomial Function to Discriminate between Users

Recall that one of the objectives of GRAIL is to better discriminate between users (*F2*). The use of entropy was a first step in that direction. However, part of users may be tightly bunched in the polarization representation space. This limits the differentiation between them. To ensure a wide distribution, we propose to apply a non-linear transformation of the entropy values of the entire set of users to better discriminate between them.

We took inspiration from the logistic function, often used to model non-linear relationships between variables. It represents the cumulative distribution function of the logistic law that is of interest to us. Such a function is bounded, monotonically increasing, has exactly one inflection point at the midpoint, and is differentiable. Its first derivative is a bell shape and its curve has a S-shaped characteristic.

Nevertheless, the standard logistic function has two asymptotes with equations $y = 0$ and $y = 1$. We intend that GRAIL scores range in $[0, 1]$, 0 indicating a lack of polarization, and 1 indicating extreme polarization, in line with measures from the literature [15, 16]. We therefore look for a numerically stable function f with $f(0) = 0$ and $f(1) = 1$, and turn to a function with a similar sigmoid pattern. In accordance with the above-mentioned constraints, we propose to rely on the polynomial function (Equation (4)) to perform the transformation of entropy values.

$$f(x) = \frac{x^a}{x^a + (1-x)^a} \quad (4)$$

where x^a controls the yield curve and replaces the traditional e^x . In other words, the stiffness of Equation (4) is described by the parameter a . Precisely, with $a = 1$, the function is linear. The higher a , the steeper the curve and the more sensitive to small changes in x near the curvatures. With $a < 1$ the curve becomes more sensitive to changes in extreme values of x (See Figure 1a).

Let us return to our running example. Equation (4) is applied to the entropy value of each user. First, the polynomial transformation has no impact in u_1 , whatever the value of parameter a , as the entropy of u_1 is 0.5. Besides, given $a = 0.5$, the transformed entropy value of u_2 , u_3 , and u_4 are given in the last column of Table 1. It is worth noting that u_3 and u_4 , the two users with extreme and similar entropy values, are better distinguished once the polynomial function we propose is applied (See Figure 1b).

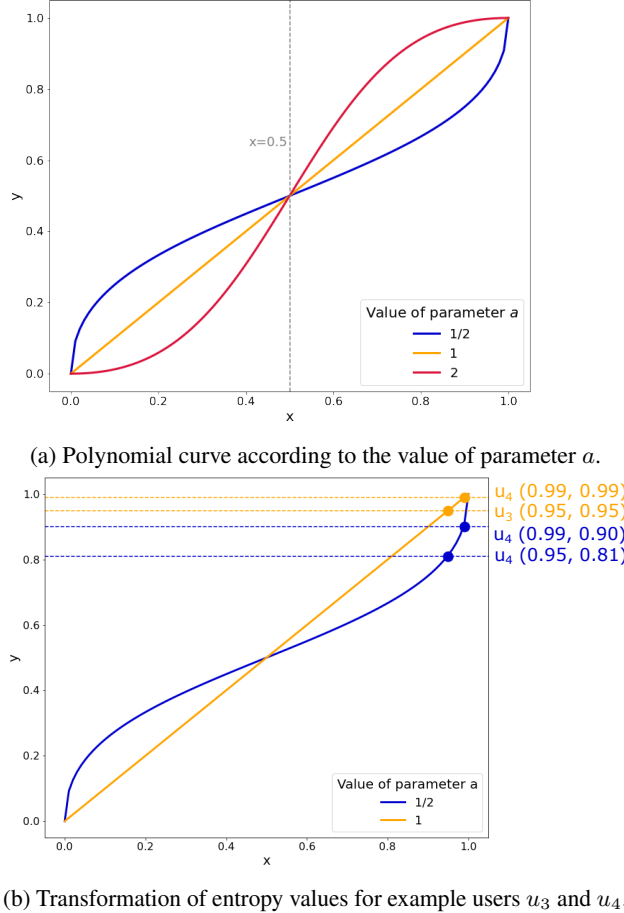


Figure 1: Polynomial function.

In Section 2, we have highlighted that individual polarization metrics of the literature do not allow for clear discrimination between users. By applying the proposed polynomial function (See Equation (4)), we aim to better discriminate users with close entropy values ($F2$). By mastering the sensitivity of the applied polynomial function thanks to parameter a , we can thus adapt this fine discrimination process to a specific set of users.

3.3 A Generalized Additive Model to Handle Multiple Factors

Polarization can occur in multiple contexts and can be influenced by a wide range of factors. We advocate for polarization metrics that consider this multiplicity. These factors are not necessarily related in a linear way, as they can observe unstable variations, and their relation can be particularly difficult to model. For example, the diversity of sources of information a specific user accesses may not be in a linear relationship with the diversity of the topics she is interested in.

In this respect, Generalized Additive Models (GAMs) appear as a suitable tool to combine multiple dimensions in a unique metric [36] (*F1*). GAMs allow the modeling of complex relationships between a response variable Y , and several predictor variables X_i . The relationship between each variable X_i and Y is modeled using a complex function named a *smooth function* f_i . These smooth functions are often represented as splines which are piecewise polynomial functions. These functions can be combined in a unique model, or be added one at a time for each predictor variable. In our work, we model the relationship between factors X_i , related to users' behavior, and polarization Y . That is why the flexibility offered by GAMs is of particular interest and meets *F1* and *F3*. Taking the structure of GAMs, GRAIL is defined as follows:

$$GRAIL(u) = \sum_i \alpha_i f_i(X_{u,i}) + \beta \text{ with } \sum_i \alpha_i = 1 \quad (5)$$

where u is a user, f_i is a smooth function, X_i is a predictor variable, α_i is the weight of predictor variable X_i and β is a bias term. We use the polynomial function defined in Equation (4) as a smoothing spline. As we expect the value for a non-polarized user to be 0, we choose to set the bias term to $\beta = 0$. Besides, if $\alpha_i = \alpha_j \forall (i, j)$, it comes down to a traditional GAM.

To go further, highly polarized topics most often confront two opposing communities: *pro* and *anti*. For example: pro-life/anti-life, pro-guns/anti-guns, pro-vaccine/anti-vaccine, etc. We thus propose to adapt GRAIL so that it not only informs about a user's degree of polarization but also her community of belonging. Thus, GRAIL is evaluated as:

$$GRAIL(u) = sgn(p) * \sum_i \alpha_i f_i(X_{u,i}) \text{ with } \sum_i \alpha_i = 1 \quad (6)$$

$$sgn(p) = \begin{cases} -1 & \text{if } u \text{ belongs to the anti community,} \\ +1 & \text{if } u \text{ belongs to the pro community.} \end{cases} \quad (7)$$

In the polynomial function f (Equation (4)), we propose to replace the variable x by $(1 - H_N)$, noted $H'(X)$ below, so that higher scores are associated to polarized users. The resulting function is:

$$f_i(X_{u,i}) = \frac{H'(X_{u,i})^{a_i}}{H'(X_{u,i})^{a_i} + (1 - H'(X_{u,i}))^{a_i}} \quad (8)$$

Altogether, the basic constituents of GRAIL make it a good candidate to better discriminate between polarization behaviors. This holds the promise of a better understanding of polarization.

Returning to the running example, additional factors could be computed for users u_1 to u_4 , such as the diversity of topics, communities, sentiment... As a final step, these multiple factors are combined using the GAM to compete GRAIL. The resulting value is thus multi-factorial (*F1*), and allows a clear distinction between users (*F2*).

We would like to point out that both predictor variables X_i and smooth functions f can be instantiated by any other equation ranging $[0, 1]$ to fit any other research context and associated data (*F3*). In a social media analysis context, the factors could for example be graph-related measures (e.g. centrality, PageRank, etc.).

As a final step, GRAIL parameters can be tuned to better fit a given dataset. First, α can be fixed to correspond to research needs or to reflect specific behaviors observed in a dataset. Second, a (in Equation (6)) can also be tuned according to the distribution of data values. Especially, if many users have intermediate values, $a > 1$ will help to better discriminate them, while reassembling users having extreme values. On the contrary, in a context with few intermediate values, $a < 1$ should be preferred so as to keep a differentiating power on extreme values. That way, the tuning phase on the a parameter can also serve to evaluate the polarizing power of a topic within the studied dataset. The optimization phase for these parameters may be based on observations or on expert knowledge in a controlled environment. Optimally, several objective functions can be defined, and parameter values set according to results.

4 Data: A Social Media Context

Social media represent a rich source of data and are commonly studied in the literature to observe polarization behavior. Twitter is one of these social media, where users can share their opinions in concise texts, and where public personalities (politicians, journalists, doctors. . .) are very active. In our work, we use the dataset presented by Treuillier *et al.* [21]. The latter is about the COVID-19 vaccine debate, and data was collected based on the elite users concept¹ [37].

The collected dataset contains 6,697 tweets published during a 7 months period, from January 1st to July 31st, 2022. Among them, 1,869 tweets are from the pro-vaccine elite users, while 4,828 are from anti-vaccine ones. Among all standard users having retweeted these tweets, 1,000 have been studied. The latter have made 16,791 retweets on pro-vaccine tweets, and 283,088 on anti-vaccine tweets.

As an exploratory step, we evaluate to what extent this dataset actually represents a highly polarized environment. We represent associated data as a graph made of 1,020 nodes (1,000 standard users + 20 elite users), connected by 7,242 edges, corresponding to standard users' retweets on elite users' tweets. There are 2 highly connected sets of nodes (*modularity*=0.55 [38]), and few edges between them. Besides, the controversy of the vaccine debate computed using the Random Walk process described by Garimella *et al.* [12] is high (≈ 0.9), reflecting a minimal number of interactions between the two communities. These results thus confirm that data are consistent with a highly polarized environment, where anti-vaccine users are much more active than pro-vaccine users, and with few intermediate users, *i.e.* users at the interface between the two communities.

5 Experiments

This section is dedicated to the evaluation of GRAIL on the previously described dataset.

We choose to use three relevant factors: **(1) An opinion factor** where opinion represents a community (pro or anti-vaccine). This factor is assessed from the standard users' retweets on each community; **(2) A pro-vaccine source factor**, where elite users who publish information in favor of the COVID-19 vaccine act as pro-vaccine sources. This factor is assessed from the standard users' retweets on each pro-source; **(3) An anti-vaccine source factor**, where elite users who publish against the COVID-19 vaccine act as anti-vaccine sources. This factor is assessed from the standard users' retweets on each anti-source.

GRAIL is thus evaluated from these three factors. The opinion factor is computed on the probability distribution over the two communities. The pro-vaccine (resp. anti-vaccine) source factor is computed on the probability distribution over pro-vaccine (resp. anti-vaccine) sources of length 10.

To evaluate GRAIL, we proceed in three steps. First, we focus on the distribution of the values of each factor. Second, we study the ability of these factors to collectively discriminate between classes of users. Last, we assess the relevance of GRAIL in each class of users.

To make this evaluation accurate, we compare GRAIL to another three-factor-based metric, that will be referred to as the baseline metric. In this baseline metric, an opinion factor, a pro-vaccine source factor, and an anti-vaccine source factor are also considered, but they are computed by single-factor metric from the literature. Concretely, ρ is used for the opinion factor, LD_{pro} , that corresponds to LD evaluated on pro-vaccine sources, is used for the pro-vaccine source factor and LD_{anti} is used for the anti-vaccine source factor. This comparison will contribute to highlighting the impact of the use of entropy and the transformation with the polynomial function.

5.1 Distribution of Polarization Factors

In this section, we compare the distribution of factors (opinion, pro-source, and anti-source) computed from entropy-based factors (H'_o , $H'_{s,pro}$, $H'_{s,anti}$) and from the baseline metric (ρ , LD_{pro} , LD_{anti}). The distribution of these factors is presented in Figure 2.

Let us start by focusing on the opinion factors. We would first like to remind that both GRAIL and baseline opinion factors are correlated on a length-2 vector. Indeed, even though the *polarization score* ρ only ranges in $[0.5, 1]$, the value of the second element in the vector (used in entropy) can be deduced from the first one, which leads to close distributions. More precisely, the ranking between both factors is the same, even though the associated values are different. Second, we can see that both distributions are tight. This reflects the fact that standard users mainly have an activity in a single community and confirms that the COVID-19 vaccine debate is highly polarized.

¹Popular and legitimate users to address the vaccine debate, being either pro- or anti-vaccine.




Figure 2: Violin plots of the distribution of polarization factors for GRAIL and baseline metrics.

Let us now focus on the source factors. About the GRAIL source factors, we can see that they are more distributed than the GRAIL opinion factor, with the mass of values distributed in the entire range $[0;1]$. In the anti-vaccine source factor, most of the values are around 0.35, whereas they are around 0.55 in the pro-vaccine factor. This means that users tend to retweet a larger number of sources in the pro-vaccine sources than in the anti-vaccine sources (whatever is their respective number of retweets in either community). We can conclude that users tend to be less polarized in the pro-vaccine factor than in the anti-vaccine factor when considering the source factors. Comparing these GRAIL source factors with those of the baseline metric, they are more widely distributed than baseline factors. This is first due to the LD equation, that bounds values in $[0.1;1]$. Beyond this, in the anti-vaccine factor, GRAIL values are significantly more widely distributed than in the baseline anti-vaccine factor.

To summarize, we confirm that GRAIL factors are more widely distributed than baseline factors, which confirms the added value of entropy when we consider to discriminate users. We thus wonder how these factors impact the discrimination of polarization behaviors classes. We propose to compare clusters obtained with the GRAIL factors to those obtained with baseline factors.

5.2 Discrimination of Polarization Behavior Classes

We look for the parameters combination that allows to better discriminate between polarization behaviors. We thus rely on a clustering process, applied to 3D data points: opinion factor, pro-vaccine source factor, and anti-vaccine source factor. We start by tuning GRAIL parameters, as explained in Section 3.3.

5.2.1 Parameters tuning

Two parameters are optimized: parameter a in Equation (4), which influences the shape and the slope of the sigmoid-like curve, and therefore the way entropy-based terms are transformed; and parameter α , which is used to weight the factors combined in GRAIL.

To perform the tuning of α , we make it vary between 0 and 1 with $\Delta_\alpha = 0.1$. The value α is set on the first factor (H'_o), and $(1 - \alpha)$ is evenly distributed on the sources terms ($H'_{s,pro}$ and $H'_{s,anti}$). We make a vary between the following values: $[1/4, 1/3, 1/2, 1, 2, 3, 4]$. In total, 77 combinations of parameters are tested.


Both baseline and GRAIL metrics are studied. Note that, to be consistent with the GRAIL metric which gives information about the community of belonging of the user (See Section 3.3), the opinion factors are oriented and normalized in $[0, 1]$ for the clustering process. The resulting factor equals 0 if the user is polarized in the anti-vaccine community, and equals 1 if she is polarized in the pro-vaccine community.

Once the three factors, either baseline or GRAIL, are computed for each standard user, we run the traditional k -means algorithm [39] and select the parameters combination that allows maximizing clustering performance, evaluated through well-known Davies-Bouldin [40] and Silhouette [41] indexes.


Considering the baseline metric, the optimal value of α is 0.5, allowing the identification of 2 clusters ($k = 2$) with a Silhouette Index = 0.80 and a Davies-Bouldin Index = 0.32. Besides, for GRAIL factors, optimal performance is reached with $\alpha = 0.6$ and $a = 1/2$. It corresponds to a Silhouette Index = 0.85 and Davies-Bouldin Index = 0.35. These values are obtained with $k = 4$ clusters. The optimal value $a = 1/2$ confirms that the transformation of entropy-based factors following the polynomial function in Equation 4 allows to better discriminate between polarization behavior classes. Besides, as the optimal value of $a < 1$, it confirms that users interact on a highly controversial debate, on which most of them are strongly polarized, while users with intermediate scores are few in number (See Section 3.2).

5.2.2 Interpretation of Identified Polarization Behavior Classes

Here, we compare clusters identified through the GRAIL factors to those obtained with the baseline factors. These clusters are presented in Figure 3.



(a) Clusters identified by baseline factors.



(b) Clusters identified by GRAIL factors.

Figure 3: Clusters.

First, looking closely at the 2 clusters obtained with baseline factors (See Figure 3a), we can see that they are only differentiated by the baseline opinion factor, with one cluster corresponding to users who access mostly the anti-vaccine community (noted C1 in Figure 3a), and the other corresponding to users who access mostly the pro-vaccine community

(noted C_2 in Figure 3a). This was expected given the tight distribution of this factor (See Figure 2). Due to the heterogeneity of behaviors adopted by users composing C_1 and C_2 , it is difficult to characterize them more precisely. So, although the two baseline sources factors are evenly distributed, they do not contribute to a clear discrimination of polarization behavior classes: users are only split according to their belonging to either community, the diversity of sources they interact with does not help to discriminate between them. The polarization behaviors reflected by these clusters are thus difficult to interpret.

Second, the 4 clusters discriminated based on GRAIL factors (See Figure 3b) are differentiated by the three factors: the opinion factor and the two source factors. The polarization behaviors adopted by users in each of these clusters can be more easily characterized:

1. C_3 (452 users) is made up of highly polarized users in the pro-vaccine community. They interact solely within this community, and the diversity of sources with which they interact varies greatly.
2. C_4 (360 users) is made up of highly polarized users in the anti-vaccine community. They interact solely within this community, and the diversity of sources with which they interact varies greatly.
3. C_5 (141 users) is made up of users who interact mostly in the anti-vaccine community, but have retweeted at least 1 tweet from the pro-vaccine community. In the anti-vaccine community, to which they are closer, users globally interact with a large variety of sources, while they only interact with a limited number of sources in the pro-vaccine community.
4. C_6 (47 users) is made up of users who interact in both communities, with a preference for the pro-vaccine community. The distribution of their interactions in each community is more balanced than for users in C_3 . In the pro-vaccine community, to which they are closer, users globally interact with a variety of sources, while they only interact with a limited number of sources in the anti-vaccine community.

To sum up, we can say that GRAIL factors are not only well-distributed but also that they allow to differentiate and characterize different behaviors (in terms of opinion and source access).

In a nutshell, based on the tuning phase and clustering results, the transformation of entropy-based factors following the polynomial (Equation (4)), as well as their weighting, allow for finely discriminate polarization behaviors and with high clustering performances. Although interesting, we cannot conclude about the reliability of GRAIL to evaluate individual polarization. That is why we propose to go further and complete the evaluation of GRAIL, by explaining the differences of behavior adopted in identified clusters based on behavioral factors.

5.3 Explaining Polarization Scores with Additional Behavioral Indicators

Table 2: Optimal combination of indicators for hierarchical regression model for each cluster with coefficients (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$) and R^2 values.

Cluster	Ordered combination of indicators					R^2
C_3	NRTs	%vaccine	NPro			0.81
	0.01***	0.16***	-0.03***			
ΔR^2	0.05	0.07	0.74			
C_4	NRTs	%vaccine	NAnti			0.65
	-0.01***	-0.10***	0.02***			
ΔR^2	0.01	0.43	0.21			
C_5	NRTs	%vaccine	%weeks	NPro	NAnti	0.80
	-0.01***	-0.24*	-0.03*	0.07***	0.01***	
ΔR^2	0.01	0.16	0.01	0.6	0.02	
C_6	NRTs	%vaccine	%weeks	NPro	NAnti	0.61
	0.01	0.85	-0.15	0.02***	-0.12***	
ΔR^2	0.01	0.03	0.05	0.11	0.41	

To analyze to what extent GRAIL provides a reliable evaluation of polarization, we propose a breakthrough evaluation methodology that relies on hierarchical regression (HR). As users adopt well-discriminated behaviors between clusters, we draw the following hypothesis: if the variance of GRAIL scores in each cluster can be explained by additional behavioral indicators (*i.e.* not used by GRAIL), it will confirm the ability of GRAIL to reflect polarization resulting from different behaviors.

HR is a statistical model that helps analyze the relationship between a dependent variable and one or more independent variable(s). Concretely, HR provides a final R^2 value that represents the proportion of the variability of the dependent variable that can be explained by a combination of independent variables, each allowing to increase this proportion (ΔR^2). It also provides coefficients informing about the relationship between the dependent and independent variables. It gives information about the variation of the dependent variable when the independent varies by one unit. In this work, GRAIL scores correspond to the dependent variable, while additional behavioral indicators (introduced below) are the independent variables.

Appropriately, the studied dataset allows us to infer information, from which indicators, that are different but complementary to GRAIL factors, can be computed. We choose to build indicators, computed from raw data, that represent user engagement, interest, and regularity of interaction in the COVID-19 vaccine debate. For each standard user, we compute the following indicators and apply hierarchical regression:

- *NRT* - Number of retweets about the vaccine debate on elite users' tweets;
- *%vaccine* - Proportion of retweets about the vaccine debate in relation to all retweets;
- *%weeks* - Proportion of active weeks, *i.e.* number of weeks in which the user retweeted;
- *NPro* - Number of pro-vaccine elite users retweeted;
- *NAnti* - Number of anti-vaccine elite users retweeted.

As the size of clusters varies by a factor of 10, we perform a preliminary power test to estimate the required sample size to get robust statistical results and make sure that the results are statistically significant. We set a medium effect size Cohen's $h = 0.5$, and the required minimal power of 0.8, *i.e.* the probability that the hierarchical regression returns an accurate rejection of the null hypothesis. With these parameter values, the minimum required sample size is 34, so results returned by the regression model are robust and well detect statistical differences. Table 2 presents final R^2 values in the rightmost column, as well as the ΔR^2 brought by each indicator and the associated coefficient, whose order and number have been optimized to maximize R^2 .

As a preliminary step, we perform the hierarchical regression on the polarization scores resulting from the combination of baseline factors. The resulting R^2 values are satisfactory (0.35 for C_1 and 0.57 for C_2). However, when we apply it on GRAIL, we notice that R^2 values are very high for every behavioral class, ranging between 0.61 and 0.81. This confirms that clusters can be better explained through GRAIL scores, whatever their size and the polarization behavior they represent.

Looking closer at the hierarchical regression results on GRAIL, we notice that the number and type of indicators in the optimal combination differ between clusters: clusters of highly polarized users (C_3 and C_4) can be explained only with the number of retweets, the proportion of retweets on the COVID-19 vaccine debate, and the number of elite users retweeted in the community of belonging of polarized users. For clusters of users interacting in both communities (C_5 and C_6), the proportion of active weeks and the number of elite users retweeted in both communities matter. This means that behaviors adopted in each of those clusters can be explained differently, thus behaviors actually differ between identified clusters. In class C_3 (highly polarized pro-vaccine users), as we might have expected, GRAIL scores are mainly explained by the number of pro-vaccine elites retweeted, with GRAIL scores dropping by 0.03 when one additional pro-vaccine elite user is retweeted (more diverse sources). In the same way, in class C_4 (highly polarized users in the anti-vaccine community), GRAIL scores increase (getting closer to 0) when the number of anti-vaccine users retweeted increases, as expected. Let us now look at C_5 and C_6 . About C_5 , which is made up of users who mainly interact with the anti-vaccine community, and to a lesser extent with the pro-vaccine community, numerous indicators contribute to the explanation of the variation of GRAIL scores. The interest in the vaccine debate, evaluated through the proportion of interactions on the topic (*%vaccine*), is important: users interacting a lot on the debate tend to get lower GRAIL scores. This means that in C_5 , the more users interact about the vaccine debate, the more they become polarized towards their majority community. About C_6 , made up of users with more balanced interactions with both communities and a slight imbalance in favor of the pro-vaccine community, we can notice that the same type of results are obtained: when their proportion of retweets on the debate increases (*%vaccine*), the polarization scores increases drastically, with an increase of 0.85. Considering now the frequency of interactions (*%weeks*), the impact is not the same for users in C_5 and C_6 : with more regular interactions, users of C_5 see their polarization score decreasing by -0.03 , thus translating a higher polarization. On the contrary, when users of C_6 , whose polarization scores are positive, more frequent interactions imply a decrease of GRAIL scores by 0.15. The effect of the frequency of interactions seems thus to depend on the community to which the user is closer. Besides, results confirm that for C_5 and C_6 , an increase in the number of retweeted pro- or anti-vaccine elite users leads to less extreme GRAIL scores.

6 Conclusion and Perspectives

In this work, we proposed GRAIL (*GeneRalized Additive poLarization*), the first individual and multi-factorial polarization metric. GRAIL has three specificities. Each factor is modeled through entropy. The relationship between factors is not linear, it relies on a generalized additive model that uses a polynomial smooth function. GRAIL is generalizable and can be adapted to the studied dataset thanks to two main tunable parameters: the shape and the slope of the sigmoid-like curve and the weight of each factor.

GRAIL has been evaluated on a COVID-19 vaccine Twitter dataset, where it is implemented as a 3-factorial metric. GRAIL has been compared to a 3-factorial baseline metric, that exploits measures from the literature. Experiments confirm that the use of entropy contributes to a wider distribution of factors. In addition, GRAIL factors, transformed by a polynomial function, allow the identification of well-separated behavioral clusters. This emphasizes that an adequate combination of factors leads to a reliable assessment of polarization. Besides, through an innovative evaluation framework, based on hierarchical regression, we confirm the relevance of GRAIL by finely characterizing each polarization behavioral class. These experiments contribute to answer the two RQs raised in the introduction. First, the use of a generalizable individual polarization metric relying on a GAM allows to combine polarization factors accurately (RQ1). Second, GRAIL scores can be well explained in each behavioral class, and thus enables a fine and accurate characterization of adopted behaviors (RQ2).

In this work, we have focused on polarization in the specific context of social media, by exploiting a Twitter dataset. Nevertheless, GRAIL's components make it a good candidate for assessing polarization in a broader context. Especially in a wider perspective of access to information, notably via recommendation systems. The multiple polarization factors processed by GRAIL can be adapted and determined according to the context of the study and associated data. Going further, as the polarization is highly dynamic, we also wonder how GRAIL could be exploited to assess the temporal process of polarization. Experiments on this temporal dimension are currently being carried out.

In the meantime, the version of GRAIL proposed in this paper allows a finer understanding of polarization behaviors on social media and brings new perspectives to take charge of online polarization. Among other things, the need for personalized depolarization strategies has been put forward in the literature: a joint approach is no longer sufficient [5, 10, 11]. In this respect, the unprecedented characterization of polarization behaviors based on GRAIL is of great interest. It will contribute to our future works related to the development of depolarization algorithms, used by recommender systems. Such systems can propose recommendations tailored to each user's behavior, with the goal of controlling the effect of their depolarizing capacity.

Acknowledgments This research was supported by BOOM (Modeling and Opening Opinion Bubbles) (ANR-20-CE23-0024).

References

- [1] Yochai Benkler, Robert Faris, and Hal Roberts. *Network propaganda: Manipulation, disinformation, and radicalization in American politics*. Oxford University Press, 2018.
- [2] Moonsun Choi. A concept analysis of digital citizenship for democratic citizenship education in the internet age. *Theory & research in social education*, 44(4):565–607, 2016.
- [3] Joshua A Tucker, Andrew Guess, Pablo Barberá, Cristian Vaccari, Alexandra Siegel, Sergey Sanovich, Denis Stukal, and Brendan Nyhan. Social media, political polarization, and political disinformation: A review of the scientific literature. *Political polarization, and political disinformation: a review of the scientific literature (March 19, 2018)*, 2018.
- [4] Emily Kubin and Christian von Sikorski. The role of (social) media in political polarization: a systematic review. *Annals of the International Communication Association*, 45(3):188–206, 2021.
- [5] Jonathan Stray. Designing recommender systems to depolarize. *First Monday*, 27, 2021.
- [6] Markus Prior. Media and political polarization. *Annual Review of Political Science*, 16:101–127, 2013.
- [7] Natali Helberger, Kari Karppinen, and Lucia D'Acunto. Exposure diversity as a design principle for recommender systems. *Information, Communication & Society*, 21(2):191–207, 2018.
- [8] Gabriel Machado Lunardi, Guilherme Medeiros Machado, Vinicius Maran, and José Palazzo M. de Oliveira. A metric for filter bubble measurement in recommender algorithms considering the news domain. *Applied Soft Computing*, 97:106771, 2020.

- [9] Lucien Heitz, Juliane A Lischka, Alena Birrer, Bibek Paudel, Suzanne Tolmeijer, Laura Laugwitz, and Abraham Bernstein. Benefits of diverse news recommendations for democracy: A user study. *Digital Journalism*, pages 1–21, 2022.
- [10] Abraham Bernstein, Claes De Vreese, Natali Helberger, Wolfgang Schulz, Katharina Zweig, Christian Baden, Michael A Beam, Marc P Hauer, Lucien Heitz, Pascal Jürgens, et al. Diversity in news recommendations. *arXiv preprint arXiv:2005.09495*, 2020.
- [11] Celina Treuillier, Sylvain Castagnos, Evan Dufraisie, and Armelle Brun. Being diverse is not enough: Rethinking diversity evaluation to meet challenges of news recommender systems. In *Fairness in User Modeling, Adaptation and Personalization (FairUMAP 2022)*, 2022.
- [12] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. Quantifying controversy on social media. *ACM Transactions on Social Computing*, 1(1):1–27, 2018.
- [13] Pedro Guerra, Wagner Meira Jr, Claire Cardie, and Robert Kleinberg. A measure of polarization on social media networks based on community boundaries. In *Proceedings of the international AAAI conference on web and social media*, volume 7, pages 215–224, 2013.
- [14] Alfredo Jose Morales, Javier Borondo, Juan Carlos Losada, and Rosa M Benito. Measuring political polarization: Twitter shows the two sides of venezuela. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 25(3):033114, 2015.
- [15] Carolina Becatti, Guido Caldarelli, Renaud Lambiotte, and Fabio Saracco. Extracting significant signal of news consumption from social networks: the case of twitter in italian political elections. *Palgrave Communications*, 5(1):1–16, 2019.
- [16] Tomas Cicchini, Sofia Morena Del Pozo, Enzo Tagliazucchi, and Pablo Balenzuela. News sharing on twitter reveals emergent fragmentation of media agenda and persistent polarization. *EPJ Data Science*, 11(1):48, 2022.
- [17] Daniel Geschke, Jan Lorenz, and Peter Holtz. The triple-filter bubble: Using agent-based modelling to test a meta-theoretical framework for the emergence of filter bubbles and echo chambers. *British Journal of Social Psychology*, 58(1):129–149, 2019.
- [18] John T Jost, Delia S Baldassarri, and James N Druckman. Cognitive–motivational mechanisms of political polarization in social-communicative contexts. *Nature Reviews Psychology*, 1(10):560–576, 2022.
- [19] Carlo M Valensise, Matteo Cinelli, and Walter Quattrociocchi. The drivers of online polarization: fitting models to data. *Information Sciences*, page 119152, 2023.
- [20] Samantha C Phillips and Kathleen M Carley. An organizational form framework to measure and interpret online polarization. *Information, Communication & Society*, pages 1–33, 2023.
- [21] Celina Treuillier, Sylvain Castagnos, and Armelle Brun. A Multi-Factorial Analysis of Polarization on Social Media. In *UMAP’23*, Limassol, Cyprus, June 2023.
- [22] Alina Sirbu, Vittorio Loreto, Vito DP Servedio, and Francesca Tria. Opinion dynamics: models, extensions and external effects. In *Participatory sensing, opinions and collective awareness*, pages 363–401. Springer, 2017.
- [23] Fabian Baumann, Philipp Lorenz-Spreen, Igor M Sokolov, and Michele Starnini. Modeling echo chambers and polarization dynamics in social networks. *Physical Review Letters*, 124(4):048301, 2020.
- [24] Tinggui Chen, Qianqian Li, Jianjun Yang, Guodong Cong, and Gongfa Li. Modeling of the public opinion polarization process with the considerations of individual heterogeneity and dynamic conformity. *Mathematics*, 7(10):917, 2019.
- [25] Kiran Garimella, Tim Smith, Rebecca Weiss, and Robert West. Political polarization in online news consumption. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 15, pages 152–162, 2021.
- [26] Lien Michiels, Jens Leysen, Annelien Smets, and Bart Goethals. What are filter bubbles really? a review of the conceptual and empirical work. In *Adjunct Proceedings of the 30th ACM Conference on User Modeling, Adaptation and Personalization*, pages 274–279, 2022.
- [27] Eli Pariser. *The filter bubble: How the new personalized web is changing what we read and how we think*. Penguin, 2011.
- [28] Christopher A. Bail, Lisa P. Argyle, Taylor W. Brown, John P. Bumpus, Haohan Chen, M. B. Fallin Hunzaker, Jaemin Lee, Marcus Mann, Friedolin Merhout, and Alexander Volfovsky. Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences*, 115(37):9216–9221, 2018.
- [29] Wen Wu, Li Chen, and Yu Zhao. Personalizing recommendation diversity based on user personality. *User Modeling and User-Adapted Interaction*, 28(3):237–276, 2018.

- [30] Jay J. Van Bavel, Steve Rathje, Elizabeth Harris, Claire Robertson, and Anni Sternisko. How social media shapes polarization. *Trends in Cognitive Sciences*, 25(11):913–916, 2021.
- [31] Michael Conover, Jacob Ratkiewicz, Matthew Francisco, Bruno Gonçalves, Filippo Menczer, and Alessandro Flammini. Political polarization on twitter. In *Proceedings of the international aaai conference on web and social media*, volume 5, pages 89–96, 2011.
- [32] Mark EJ Newman. Modularity and community structure in networks. *Proceedings of the national academy of sciences*, 103(23):8577–8582, 2006.
- [33] Ana Lucía Schmidt, Fabiana Zollo, Antonio Scala, Cornelia Betsch, and Walter Quattrociocchi. Polarization of the vaccination debate on facebook. *Vaccine*, 36(25):3606–3612, 2018.
- [34] Cass R Sunstein. The law of group polarization. *University of Chicago Law School, John M. Olin Law & Economics Working Paper*, (91), 1999.
- [35] Claude Elwood Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948.
- [36] Trevor J Hastie. Generalized additive models. In *Statistical models in S*, pages 249–307. Routledge, 2017.
- [37] Simonetta Primario, Dario Borrelli, Luca Iandoli, Giuseppe Zollo, and Carlo Lipizzi. Measuring polarization in twitter enabled in online political conversation: The case of 2016 us presidential election. In *2017 IEEE international conference on information reuse and integration (IRI)*, pages 607–613. IEEE, 2017.
- [38] Aaron Clauset, Mark EJ Newman, and Cristopher Moore. Finding community structure in very large networks. *Physical review E*, 70(6):066111, 2004.
- [39] Aristidis Likas, Nikos Vlassis, and Jakob J. Verbeek. The global k-means clustering algorithm. *Pattern Recognition*, pages 451–461, 2003.
- [40] David L Davies and Donald W Bouldin. A cluster separation measure. *IEEE transactions on pattern analysis and machine intelligence*, (2):224–227, 1979.
- [41] Peter J. Rousseeuw. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, pages 53–65, 1987.