



HAL
open science

Inferring Attitudinal Spaces in Social Networks

Pedro Ramaciotti Morales, Jean-Philippe Cointet, Gabriel Muñoz Zolotoochin, Antonio Fernández Peralta, Gerardo Iñiguez, Armin Pournaki

► **To cite this version:**

Pedro Ramaciotti Morales, Jean-Philippe Cointet, Gabriel Muñoz Zolotoochin, Antonio Fernández Peralta, Gerardo Iñiguez, et al.. Inferring Attitudinal Spaces in Social Networks. 2022. hal-03573188v1

HAL Id: hal-03573188

<https://hal.science/hal-03573188v1>

Preprint submitted on 14 Feb 2022 (v1), last revised 24 Nov 2022 (v2)

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.

Inferring Attitudinal Spaces in Social Networks

Pedro Ramaciotti Morales^{1*}, Jean-Philippe Cointet¹, Gabriel Muñoz Zolotoochin², Antonio Fernández Peralta³, Gerardo Iñiguez³ and Armin Pournaki^{1,4,5}

^{1*}médialab, Sciences Po, France.

²Dept. of Mathematics, P. Catholic University of Chile.

³Department of Network and Data Science, CEU, Austria.

⁴Max Planck Institute for Mathematics in Sciences, Germany.

⁵Lattice CNRS, ENS-PSL & Sorbonne Nouvelle, France.

*Corresponding author(s). E-mail(s):
pedro.ramaciottimorales@sciencespo.fr;

Abstract

Ideological scaling methods have shown that behavioral traces in social platforms can be used to mine opinions at a massive scale. Current methods exploit one-dimensional left-right opinion scales, best suited for two-party socio-political systems and binary social divides such as those observed in the US. In this article, we introduce a new method to overcome limitations of existing methods by producing multidimensional network embeddings and align them with referential attitudinal for a few nodes. This allows us to infer a larger set of opinion dimensions from social graphs, embedding users in spaces where dimensions stand for indicators of several social dimensions including (in addition to left-right cleavages) attitudes towards elites, or ecology among many other issues. Our method does not rely on text data and is thus language-independent. We illustrate this approach on a Twitter follower network. Finally, we show how our method allows us to analyze the opinions shared within various communities of social networks. Our analyses show that communities of users that have extreme political opinions are also more homogeneous ideologically.

Keywords: Network scaling, graph embedding, ideology, political attitudinal, party systems, polarization.

Contents

1	Introduction	3
2	Related Work	5
3	Social Network and Reference Data	7
3.1	Selecting a set of Twitter users	7
3.2	Social graph data	7
3.3	External referential attitudinal data	8
4	Ideological Embedding	9
4.1	Homophily embedding	9
4.2	Interpreting the dimensions using external attitudinal data	11
4.3	Testing the positions of users with text analysis	13
4.4	Testing the positions of party sympathizers	14
5	Attitudinal Embedding	15
5.1	Mapping ideological space onto attitudinal reference frames	17
5.2	Testing the positions of users with text analysis	18
6	Social Network Analysis in Attitudinal Spaces	21
7	Discussion and Conclusions	23

1 Introduction

Over the last 10 years, social media, with the wealth of granular behavioral data they produce, have been imagined as a privileged source to mine opinions (for a recent extensive review, see the survey by Messaoudi et al., 2022). Inferred opinions may help answer numerous questions about online social and political dynamics. One can use such data to investigate the effect of political preferences on algorithmic recommendation systems (Bakshy et al., 2015; Ramaciotti Morales and Cointet, 2021), or to uncover political motivations driving certain groups of actors and social movements (Budak and Watts, 2015; Cointet et al., 2021) among others. Current methods for opinion mining leverage a diverse set of principles, including the geometrical approaches. We adopt such spatial approaches and position users of a social network in a space where dimensions are informative of opinions, acting as indicators of positive and negative attitudes towards relevant issues. These methods can be traced back to *ideology scaling* methods in political sciences (Poole and Rosenthal, 1985) in recently in online social networks (Barberá, 2015; Bond and Messing, 2015). More recently, we proposed an extension to these methodologies, the *ideological embedding* method (Ramaciotti Morales et al., 2021), with which we were able embed social graphs in multi-dimensional *ideological spaces* in which each dimension acts as indicators of positive or negative attitudes towards several grouped issues of public debate, such as issues related to globalization or immigration, going beyond traditional left-right opinion scales. Building on *ideological space* methods, we further tackle the problem of the interpretability of dimensions and distances in space, and the bounds and reference points of the dimensions used as opinion scales. Mapping attitudinal data from political surveys to inferred *ideological embeddings*, we show that we can interpret the latent dimensions which emerge from social graphs. Put differently, we can determine which dimensions are related with negative and positive attitudes towards which issues. Additionally, we propose a method to map the *ideological space* onto the scales of the external attitudinal data, meaning that all the social media users following representatives can be mapped onto the scales of any external attitudinal data, allowing for comparison between countries, or between opinions at different moments in time.

Opinions may have different forms and functions, which makes them difficult to conceptualize and operationalize. An important type of opinion is that of *evaluative opinions*: that is, to be *for* or *against* something (Bem, 1970). Evaluative opinions are often operationalized as *attitudes*: an individual disposition towards an attitudinal object (*e.g.*, person, institution, issue, event, bill, policy position). Attitudes can also be held towards complementary attitudinal objects. A classical example are attitudes towards liberal and conservative values: positive attitudes towards one set of values imply negative attitudes towards the other, so that individuals can be placed on attitudinal scales ranging from the most liberal to the most conservative positions. Attitudes are, contingently on other factors, important determinants of behavior (Ajzen,

1989; Ajzen and Fishbein, 2005). The traditional method for estimating people’s attitudes is the administration of surveys or polls (*e.g.*, the ANES poll in the US, or the Eurobarometer in Europe). Besides polling, different behavioral traces can be observed to infer the attitudinal positions of individuals. A classical example is found in the work of Poole and Rosenthal (Poole and Rosenthal, 1985, 1984, 2000), in which they estimate attitudes of parliamentarians in the US towards liberal and conservative values using Bayesian inference on roll call data. Recently, Barberá (2015) used similar methods to infer the liberal or conservative attitudes of millions of US Twitter users with statistical inference on observed friendship—who follows whom—networks. This method postulates a model for the formation of a friendship network based on homophily (people with similar attitudes establish network ties; Lazarsfeld et al., 1954), and that the opinions of users are accounted for with a single latent *ideological* variable. The method then uses an observed network structure to perform an inference to estimate the ideological parameter. Importantly, these methods do not rely on textual data, making them language-independent. Empirical validation using external data (such as self-declared political affiliation) has shown that the recovered attitudinal variable in the US coincides with attitudes towards liberal and conservative values.

These methods for the inference of attitudes using social networks, however, have been much less successful in European and other settings (Barberá and Rivero, 2015). This is due to the multi-issue, multi-party socio-political underlying systems that structure public debate (Hix et al., 2006; Benoit and Laver, 2012), and ultimately online social networks such as Twitter. In this article we use methods inspired in network scaling on large social networks, and political science expert survey data with attitudinal positions of very few referential social network nodes. We show that, using both sources of data, it is possible to extract attitudinal positions for several issues for large portions of social networks. The result is an ideological embedding procedure for social networks in which dimensions stand as indicators for attitudes towards different issues, such as taxation, immigration, left-right cleavages, and trade protectionism, to name a few. The resulting spaces are *ideological* in the sense that the position of a user along each dimension provides information about the attitude of that user towards a set of combined issues. These latent ideological spaces reported by Ramaciotti Morales et al. (2021) have known limits, which we discuss and address in this work. While dimensions can be shown to be related to positive and negative attitudes towards certain issues, they do not have a consistent metrical meaning. For example, a user in a position with value equal to 2 cannot be said to be twice as favorable than a user in a position with value equal 1. In other words, in ideological spaces order is readily meaningful but distances are not. Similarly, there are no evident referential positions in space. For example, a user in the position with value equal to 0 in the left-right scale cannot be said to have a central or apolitical position. To overcome these limits, we further map positions of *ideological space* onto the

attitudinal spaces as defined by political surveys, which have scales with meaningful reference points (for example, a left-right ranging from extreme left-wing at 0 to most right-wing at 10) and explicit reference points (for example, a value equal to 5 meaning a central position in a left-right scale ranging from 0 to 10). Through this mapping, we position users in the scales of the survey as if all users in our sample had been placed using the survey. We illustrate this procedure with French Twitter data, and we propose benchmarks to test the validity of the approach. We also show the usefulness of our method by analyzing a fraction of the Twitter social graph. We separate the social graph into communities, and analyze their mean positions and homogeneity along several opinion dimensions. We show that communities that are in the political far-left and far-right extremes are more homogeneous, *i.e.*, their members are less dispersed along the left-right dimension, providing novel experimental evidence to the study of phenomena related to the so-called “echo chambers” (Quattrociocchi et al., 2016).

This article is structured as follows. After discussing the related works that are pertinent to this article, we present the datasets on which we will illustrate our method. We then present the *ideological embedding* method as it was introduced by Ramaciotti Morales et al. (2021), detailing the embedding method and the use of external attitudinal data to interpret the substance of the inferred dimensions. Next, we present the *attitudinal embedding* method, which consists in projecting the positions of users in *ideological space* onto the scales of the attitudinal spaces defined in external data. We then show different social network analyses that can be performed using these embedded social networks. In particular, we investigate how detected communities are distributed in our attitudinal spaces. Finally, we present our main conclusions and discuss about several possible applications of our method in diverse fields of research, such as the study of populism, polarization, and the effects of Recommender Systems.

2 Related Work

The broadest field within social network analysis in which this work is inscribed is that of opinion mining (Liu and Zhang, 2012; Messaoudi et al., 2022). This field consists of a diversity of techniques and models. Many methods rely on textual analysis (Gentzkow et al., 2019; Groseclose and Milyo, 2005), and are thus language- and context-dependent. For example, people express themselves differently on different platforms, and a model capturing ideology from parliamentary debates (Rheault and Cochrane, 2020) will necessarily differ from one predicting political slant in tweets (Stefanov et al., 2020). Our work is mainly concerned with a family of methods called ideological scaling. This family consists of methods that produce spatial models that explain choice data (*e.g.*, parliamentarians voting for bills, users liking politicians online), in which dimensions hold some relation with choice, and counts a variety of applications (Imai et al., 2016). Scaling methods rely only on relational traces, producing

ideological spatializations from topological relations: *e.g.*, who is friends with whom, or votes for or clicks on what. The NOMINATE method by Poole and Rosenthal (1985) is a landmark and pioneering example. Using voting data from the US congress, Poole and Rosenthal were able to position Democrat and Republican parliamentarians on ideological left-right, and measure distances between parties as a proxy for polarization (McCarty et al., 2016)

Bond and Messing (2015) were among the first to apply ideological scaling methods to large social network choice data. In their pioneering work, they applied the same principle of NOMINATE on data on how users *liked* politicians on Facebook. The resulting spatial models allowed them to position these politicians present on Facebook on some left-right scale, according to how they were perceived by users. Their method also allowed them to position the nearly 6 million that gave *likes* on the online platform. The work presented in this article is related to a similar network scaling method, proposed by Barberá (2015) for one-dimensional scaling. This scaling computes a single latent ideological parameter ϕ_i for every user i , following an homophily probabilistic law adjusted for activity and popularity:

$$P(A_{ij} = 1 | \alpha_i, \beta_j, \gamma, \phi_i, \phi_j) = \text{logit}^{-1}(\alpha_i + \beta_j - \gamma |\phi_i - \phi_j|^2), \quad (1)$$

where $A_{ij} = 1$ when user i follows user j , α_i and β_j are the “activity” (tendency to follow others) and “popularity” (tendency to be followed) of users i and j , ϕ_i and ϕ_j are their latent ideological parameters, and γ is a normalization constant. When computed in the bipartite network of the US parliamentarians and their followers on Twitter, using Markov Chain Monte Carlo methods, the ideological parameters ϕ_i were shown to act as indicators of attitudes towards liberals and conservatives. The inference of ideological parameters has been shown (Lowe, 2008; Carroll et al., 1997) to be approximated by Correspondence Analysis (CA) (Greenacre, 2017), which has also been verified experimentally on Twitter data (Barberá et al., 2015). CA also allows for the inference of multi-dimensional ideological parameters, embedding users in ideological spaces where positions are informative of their attitudes (D’Esposito et al., 2014). In recent works (Ramaciotti Morales et al., 2020; Cointet et al., 2021) we have suggested that these multi-dimensional parameters might be related to attitudes towards several issues of public debate, beyond the classical one-dimensional attitude scale from left-right or liberal-conservatives cleavages. Several recent works leverage ideological inference, for example for measuring polarization in politics (Flamino et al., 2021) or around particular issues, such as climate change (Falkenberg et al., 2021). In this article, we further explore the extraction of multi-dimensional attitudinal indicators suggested by Ramaciotti Morales et al. (2020), and later by Ramaciotti Morales et al. (2021), detailing the inference of latent multidimensional parameters. In contrast with opinion inference relying on text analysis or news media citations (Baumann et al., 2020; Kwak et al., 2021), ideological embeddings relying only social networks structures have the potential to be language- and context-independent, thus allowing, for example, comparisons between countries, and between users with data collected at different moments in time.

3 Social Network and Reference Data

In this section we describe the data that we use in the rest of the article. These data are organized in three types: Section 3.1) a set of Twitter users for which we will illustrate the proposed method, Section 3.2) the social graph (*i.e.*, the directed *who follows whom* graph) subtended by these users, and Section 3.3) external (to Twitter) data with attitudes of some reference points in several opinion scales as given by political survey.

3.1 Selecting a set of Twitter users

To select a set of users on which to apply the methods proposed in this article we use French Twitter networks of friends¹. We consider the set of the 831 (out of 925) French Members of Parliament (MPs) present on Twitter who are affiliated to 10 parties and their followers². We choose this starting point for collection, as we expect, based on abundant evidence (Barberá, 2015; Barberá and Rivero, 2015; Briatte and Gallic, 2015; Ramaciotti Morales et al., 2020) that the choice of MPs to follow will be revealing of implicit political opinions and preferences. This collection was conducted in May 2019, and resulted in 4.487.430 unique followers. In this bipartite network, some MPs are followed by a handful of users (min. degree = 38) while others are followed by nearly half of the users that follow any MP (max. degree = 2.241.986). On the other hand, many users in this network (2.279.199) follow only one MP, with the most diverse user following 757 MPs (see Fig. 1). To filter out inactive or bot accounts, and accounts without enough ideological referential connections, we follow the criteria proposed by Barberá (2015) and consider only followers that follow at least 3 MPs, and that have at least 25 followers. We also removed followers with a repeated set of followed MPs to obtain 368.831 followers to ensure the full rank of the adjacency matrix. Using the bipartite graph of MPs and their 368.831 followers we will illustrate the ideological and attitudinal embedding methods in the next Sections 4 & 5.

3.2 Social graph data

The previously described bipartite graph between MPs and their followers will be used to mine opinions through our embedding methods. Additionally, we consider the social network formed by these users for which opinions will be inferred. For these 368.831 users, we collected the subtended directed social graph (*who follows whom*). Some users have disabled permissions to have their followers collected, thus resulting in a graph of 230.911 users and 67.217.556 edges (density=0.00126). Having also the social graph subtended by 230.911 of the initial 368.831 nodes, will allow us to produce social network analyses using both opinion spatial data and topological network data in Section 6. We

¹Please refer to the Acknowledgments section to consult the data registry and treatment declaration, related GDPR documents, and access to the respective legal notice.

²Obtained from <http://www2.assemblee-nationale.fr/deputes/liste/reseaux-sociaux> for deputies, and http://www.senat.fr/espace_presse/actualites/201402/les_senateurs_sur_twitter.html for senators.

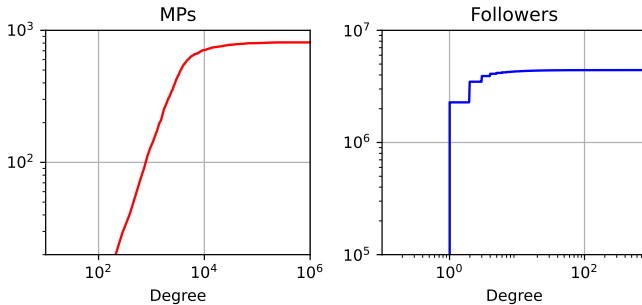


Fig. 1 Cumulative Distribution Function (CDF) of the degree of two types of nodes of the bipartite network of French MPs (in red) and their followers (in blue).

will refer to this friendship network as “social graph”, while we will refer to the links between MPs and their followers (from Section 3.1) as the “bipartite graph”.

3.3 External referential attitudinal data

To interpret the dimensions of ideological spaces, we use external attitudinal data. These external data contain the positions some referential points in attitudinal scales with predefined bounds (ranging for example from 0, *most opposed*, to 10, *most favorable*), and associated with predefined issues of public debate: *e.g.*, European integration, special rights for minorities, anti-elite sentiments. These external data will also provide a referential attitudinal frame onto which to project social network users. Since our data collection strategy is based on accounts of MPs (belonging to known political parties), we choose to use external referential frames for political parties, which will serve as reference points in space. We use the 2019 Chapel Hill Expert Data (CHES) (Bakker et al., 2020), corresponding to the period on which our Twitter data were collected. The CHES is compiled using the responses of a survey administered to 421 political scientists, in which they place European political parties on scales from 0 to 10 (or from 1 to 7 for the scale of attitudes towards EU integration) for 51 different issues. From the 51 issues, we exclude 3 related to attitudes towards Turkey and the conditions needed for EU accession, which are not available for all parties, thus resulting in 48 dimensions. These data include some scales of interest for comparison with previous studies, such as a left-right scale, but also include new important emerging societal cleavages observed by Ramaciotti Morales et al. (2021). Fig. 2 illustrates the contents of this external referential attitudinal space by showing the position of French political parties in three chosen dimensions: Left-Right, Anti-elite salience, and Anti-immigration stance.

Party positions in attitudinal data

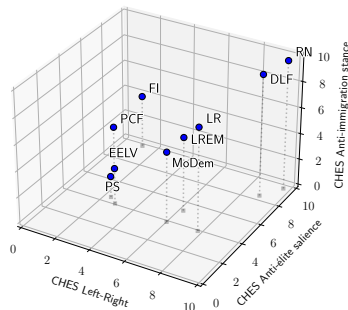


Fig. 2 Position of French political parties in attitudinal reference space provided by the Chapel Hill Expert Survey (CHES) on 3 chose dimensions (out of 51): Left-Right, Anti-elite salience, and Anti-immigration stance.

4 Ideological Embedding

This section describes the *ideological embedding* method, as presented by Ramaciotti Morales et al. (2021). We first show how to exploit generative social network models relying on homophily to infer latent dimensions on which to embed users. Next, we use external attitudinal data to show that some of these dimensions are ideological in the sense that they act as indicators of attitudes towards several political issues. We provide to types of validations to this ideological embedding: 1) comparing how online text utterances are correctly positioned in ideological space (*e.g.*, users declaring themselves as being right-wingers are positioned accordingly in the emerging left-right scale), and 2) by showing that people that declare sympathy for a political party, are positioned in the spatial vicinity of that party.

4.1 Homophily embedding

First we represent the bipartite subgraph of the 831 MP and their 368.831 followers as an adjacency matrix $A \in \{0, 1\}^{|368.831| \times |831|}$ (subfigure A in Fig. 3). Next, we produce a reduced-dimensionality representation of these 831 MPs and the 368.831 followers using the coordinates of accounts in the latent dimensions of the CA of the adjacency matrix (subfigure B in Fig. 3). Lowe (2008) has shown theoretically, and Barberá et al. (2015) experimentally, that CA provides an estimation of parameters ϕ of model from Eq. 1, when these parameters are assumed to be multi-dimensional. This first ideological embedding is said to be an homophily embedding in the sense that the underlying social mechanism leveraged to infer positions is homophily: in Eq. 1, the closer that two users i and j are in the latent space (*i.e.*, the lower the value $\|\phi_i - \phi_j\|$, for ϕ_i, ϕ_j vector parameters), the higher the probability that an edge will be observed between them. We denote the emerging latent dimensions of this bipartite graph by δ_1 to δ_{831} in decreasing order of inertia associated with each dimension. Fig. 4 shows the inertia of the latent dimensions of CA, and the




Fig. 3 Illustration of the ideological and attitudinal embedding methods with the data they take as input and what they produce as output. Ideological embedding takes a bipartite network of the parliamentarians (MPs) and their parties (**A**), to produce a spatialization in a latent ideological space (**B**). Attitudinal embedding takes the ideological positions (**B**) of all users, plus the positions of some reference points in both ideological, to embed all users in this reference space (**C**).




Fig. 4 Inertia of the principal components of the Correspondence Analysis of the bipartite network of MPs on Twitter and their followers (left), and density of users on the first two components, δ_1 and δ_2 (right).

density of the 368.831 followers in the space spanned by the first 2 latent dimensions, *i.e.*, δ_1 and δ_2 , of this reduced-dimensionality space. Keeping the same notation δ , we apply a normal standardization to each dimension to ensure that the dimension-wise mean of users is at $\delta = 0$, and that they have equal variance. When most of the inertia is concentrated in the first dimensions, a spatial representation in the first few dimensions is a suitable representation of the topological data in the sense that a random graph computed using such spatialization and the chosen probabilistic model from Eq. 1 would be similar to the original network (Roberts Jr, 2000). As seen in Fig. 4, the first two dimensions hold relatively more importance in explaining the topological network data.

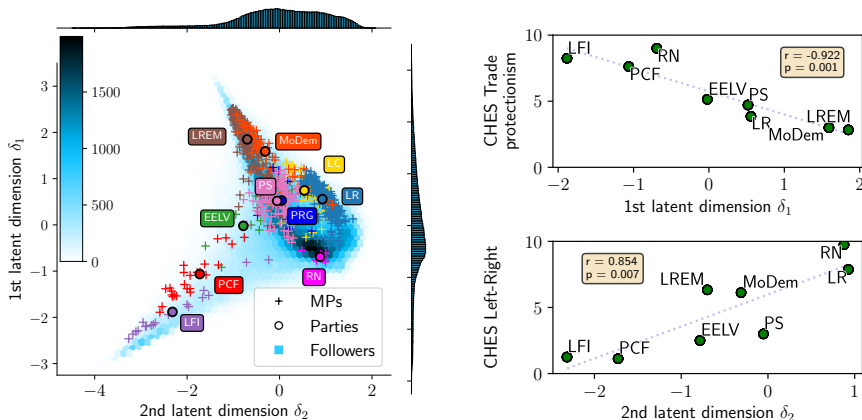


Fig. 5 Positioning of parties in the reduced-dimensionality space spanned by the first two latent dimensions of the ideological space of the bipartite Twitter network, computed as means of the positions of parliamentarians (left), and examples of two external party attitudinal positions (CHES) that are correlated with first two dimensions (right).

4.2 Interpreting the dimensions using external attitudinal data

To link the dimensions of our space with issues of public debate present in the CHES data, we compare the positions of political parties according to our latent dimensions, with their positions in the attitudinal CHES dimensions (scoring parties in a 0 to 10, or 1 to 7 scale, ranging from *very opposed* to *very favorable* to each issue). To compare a latent dimension with a CHES dimension, we compute the Pearson correlation of party positions on both. We compare each one of the first 15 latent dimensions (as suggested by Fig. 4) with all 48 available CHES dimensions. We find that only the first 3 latent dimensions show statistically significant correlations with CHES dimensions. These correlations show that each one of the first 3 latent dimensions is associated (up to a correlation significance of $\alpha = 0.05$, marked in dark blue in Fig. 6) with a set of CHES dimensions.

The first latent dimension, δ_1 , is positively correlated with positive attitudes towards the EU, suggesting that the higher the value of the δ_1 position of a user, the more positive their views are towards the EU. δ_1 is also correlated with opposition to redistribution, economic interventionism, and market regulation. It is also correlated with high levels of left-right economic policy dissent (inside parties), and with the relevance granted to Russian interference in politics. Negative δ_1 positions are correlated with positive views on economic protectionism, and with a high level of importance is placed on people-elites cleavages and redistribution. These issues are related to attitudes towards globalization in European settings, and we call δ_1 the “Local-Global” (LG) axis. The second latent dimension, δ_2 , is associated with positions on issues widely attributable to left-right cleavages in France, including: left-right

4.3 Testing the positions of users with text analysis

These associations presented so far were computed and justified using the positions of parliamentarians and their parties. We now seek to validate the definition of our LG, LR, and IM axes using the Twitter profile descriptions of their 368.831 followers, in which users briefly describe themselves. Twitter profiles are short texts in which users describe themselves and are shown on the profile page and profile preview of each user. For each one of the three axes we select two topics, and for each topic we define a minimalist dictionary to classify the profile of followers. Our goal is to show that, even when applying a minimal text analysis, our three ideological axes distribute users (and not only parliamentarians) according to the meaning proposed for these dimensions. For the LG axis we classify users with labels “Europe” (if they include the words “eu” or “europe”) and “People and Elites” (if they include the strings “peuple”, meaning *people* in French, “élite”, or “politicien”). For the LR axis we label users as mentioning the “Left-wing” (if they include the string “gauche”, meaning *left* in French) and as mentioning the “Right-wing” (if they include the string “droite”, meaning *right* in French). For the IM axis we label users as mentioning “Islam” (if they include the strings “musulman”, meaning *muslim* in French, or “islam”) and as talking about “Nationalism” (if they include the string “patriot” or “patriotique” or “patriotisme”). In the context of the French public debate, questions of multiculturalism and immigration policy rapidly revolve around a debate on Islam (Hargreaves, 2007; Freedman, 2017). In the strings that define a topic, we included all possible variants and misspellings. For the labels “Europe”, “Left-wing”, and “Right-wing” we only included profiles with positive sentiment, to differentiate those that express support for the referred topic (users might use the word “right”, *e.g.*, to express criticism of right-winger). We computed sentiments using a BERT-based multilingual model for sentiment analysis³. For the other labels we did not filter by sentiments, as we are trying to detect the importance or salience of the label along the dimensions (which does not involve positive or negative opinion). Fig. 6 shows that axis IM is, for example, correlated with the salience of the issue of immigration restrictions. Fig. 7 shows the proportion of users that use these topics in their profiles and the Clopper-Pearson confidence interval of this proportion at a $\alpha=0.05$ CI. Selection of profiles by positive sentiment is signaled in Fig. 7 with a + symbol. Our results show that, on the sets of users considered for text-based validation (selected by keywords and sentiments when pertinent), ideological dimensions produce an order among users. For example for the LR axis, a user to the left of another has higher probability of being more to the left politically than more to the right, and that this probability increases with distance along that dimension. Users at the extreme left and right edges of the sample have very low probability of being incorrectly positioned: a user describing itself using the word “right” with a positive sentiment, *e.g.*, has very low probability of being in the leftmost edge of the sample.

³<https://huggingface.co/nlptown/bert-base-multilingual-uncased-sentiment>.

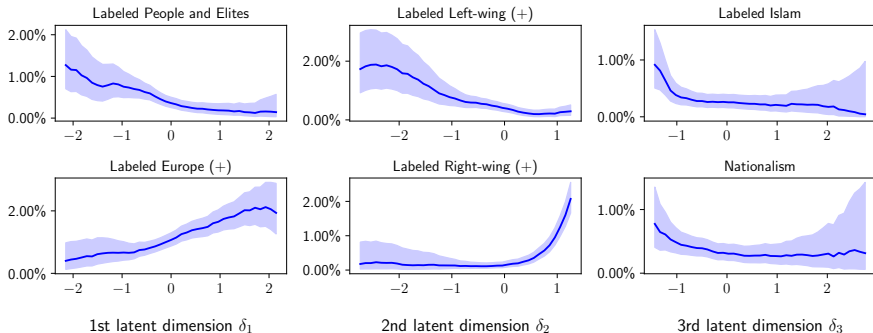


Fig. 7 Proportion of Twitter users labeled as referring to different topics on their text profiles according to their positions along the three axes of ideological space, with $\alpha = 0.05$ confidence intervals. Symbol + indicates whether only those referring to the issue with positive have been considered.

As many studies suggest, left-right cleavages (corresponding to our LR axis) have a structuring role in behavior (*e.g.*, voting behavior; Aldrich et al., 2010) and in particular on French Twitter (Briatte and Gallic, 2015). Indeed, LR axis is among the first two latent dimensions of our ideological space. In our dataset, however, as suggested by other studies (Miklin, 2014; Schön-Quinlivan, 2017; Grossman and Sauger, 2019), social cleavages related to globalization (*e.g.*, European integration, increasing trade openness, and other related to our first latent dimension, LG) take an even more important role. This is a relevant finding, indicating the declining role of left-right cleavages, now observed second to globalization in its importance to determine social choice in digital arenas.

4.4 Testing the positions of party sympathizers

A second way of looking at the position of the followers is by examining the relation between those that declare sympathies towards a party and their position in space with respect to those of the party itself. While it is difficult to account for partisanship, we propose a simple approach based on positive mentions of parties in the Twitter profile description. To do so, we propose again a minimalist approach to classifying users as sympathizers of a given party, based on a few strings for each one of the ten political parties in our dataset:

- EELV (*Europe Écologie - Les Verts*): “eelv”, “les verts”;
- LFI (*La France Insoumise*): “insoumis”, “lfi”;
- LREM (*La République en Marche*): “lrem”, “en marche”;
- MoDem (*Mouvement Démocrate*): “modem”, “mouvement démocrate”;
- LC (*Les Centristes*): “centristes”;
- PCF (*Parti Communiste Français*): “pcf”, “communiste”;
- PRG (*Parti Radical de Gauche*): “radical de gauche”, “parti radical”, “prg”;
- PS (*Parti Socialiste*): “ps”, “socialiste”;
- RN (*Rassemblement National*): “rn”, “rassemblement national”;




Fig. 8 Position of the users identified as sympathizers (mentioning them with positive sentiment) of each one of the ten parties in our dataset (identified with blue dots ●), and the position of the parties computed as the mean position of its MPs (shown in yellow dots ●).

- LR (*Les Républicains*): “lesrepublikains”, “lr”.

As before, we also consider singular and plural, masculine and feminine declensions when possible, as well as upper- and lower-case versions of our keywords. We also consider hashtag for keywords made of a single word (*e.g.*, “pcf” and “#pcf”). Finally, as in the previous section, we filter out possible negative references using the same method for estimating sentiments. Fig. 8 shows the position of party sympathizers in comparison with that of the ensemble of followers and the position of the political parties computed as the mean position of their MPs.

5 Attitudinal Embedding

The ideological space described in the previous section has important properties, but also limitations. Correlation with attitude dimensions from external data such as CHES, and later validation using text analysis, shows that dimensions act as indicators in the sense that they produce a relative order of users along them. This order is identified with a set of group issues that provide the substance of the ideological dimension. In ideological spaces, however, the notion of distance is less clear, and we lack important reference points. To illustrate these points, let us examine the results along the LR ideological axis. Concerning distance, for example, a user at a position with value equal to 2 on the LR axis (thus on the rightmost edge) cannot be meaningfully said to be half as radical or extreme than another user at a position with a value equal to -4 on the LR axis (thus on the leftmost edge of our sample). This is because ideological space has a provable ordering property but not a well-defined metric. Similarly, regarding reference points, a user at a position with value equal to 0 on the LR axis cannot be meaningfully said to be a centrist user, and it can only be said to be in the mean of our sample in that dimension. These distinctions are important, in part because whether a user is a political

radical (and to what degree) or moderate, is at the heart of research questions regarding online behavior (Osmundsen et al., 2021). Ideological spaces are also not comparable: if we take two ideological embeddings from different datasets (at two moments in time, or between networks extracted in two different countries), their positions cannot be readily compared to extract general conclusions regarding comparative order. For example if we collect a bipartite graph of the followers of MPs a month later, and if this bipartite graph has changed, we cannot compare directly the positions of a user that is present on both moments. Finally, because ideological dimensions act as attitudes towards a grouped set of issues, it is difficult to extract conclusions regarding particular individual issues. For example, in Figure 5, LFI (*La France Insoumise* party) is to the left of PCF (*Parti Communiste Français*) in the latent so-called LR dimension, but it is less straightforward to compare these two parties in two attitudinal dimensions that are included in the definition of the LR axis, such as their attitudes towards tax cuts, authoritarianism, or nationalism.

Exploiting the identified relations between the dimensions of our ideological space and the attitude dimensions of instruments such as CHES, we propose to tackle these aforementioned shortcomings of the ideological embedding method by mapping ideological space positions (subfigure B in Fig. 3) onto the attitudinal dimensions predefined by an instrument such as CHES (subfigure C in Fig. 3), using parties as reference points to compute the mapping. We call this procedure *attitudinal embedding*. The attitudinal dimensions of an external instrument such as the CHES have several advantageous properties. To illustrate this, let us take again left-right cleavages as an example. The CHES Left-Right attitudinal dimension specifies 3 spatial reference points: 0 being the extreme left for parties, 10 being the extreme right for parties, and 5 being a centrist position for parties. We emphasize the fact that these references are for parties, as it is conceivable that some users are more extreme than parties. By mapping all users in our dataset onto this CHES Left-Right dimension, we seek to obtain an approximation to the position that each user would have gotten, had it been evaluated by the CHES instrument (*i.e.*, had it been known, examined, and positioned by the same experts, which is impossible for many reasons, including economic ones).

Relying on external attitudinal dimensions also solves the problem of comparison: two ideological embeddings taken at two moments in time can be compared if projected onto the same CHES dimension, and two ideological embeddings taken on two countries can be compared, if the instrument has been applied in those two countries. The CHES dataset would allow this in theory, as it is administered in 32 countries: users from two different countries embedded using the way in which they follow local MPs, could be projected onto a common CHES dimension for comparison.

5.1 Mapping ideological space onto attitudinal reference frames

The results from the previous section, in particular the correlations shown in Fig. 5 and more systematically in Fig. 6, suggest that at least some ideological dimensions could be approximated by affine subspaces of the attitudinal space defined by the 48 attitudinal dimensions available in the CHES dataset. Thus, to map ideological spaces onto this attitudinal reference frame or attitudinal space of 48 dimensions, we seek to establish an affine transformation between ideological and attitudinal spaces (of correspondingly 831 and 48 dimensions). Determination of this affine transformation may be seen as an optimization problem using the positions of the 8 political parties that are present in both spaces (see the right subfigure in Fig.5). Let $P = 8$ be the number of political parties serving as guidance for the fit of the affine transformation, $M = 48$ the number of dimensions of the CHES attitudinal space, and N the number of latent ideological dimensions δ to be considered for the optimization of the affine transformation $T_{\text{aff}} : \mathbb{R}^N \rightarrow \mathbb{R}^M$. Let $Y \in \mathbb{R}^{M \times P}$ be the position of parties in the CHES attitudinal space, and $X \in \mathbb{R}^{N \times P}$ the position of parties in the ideological space. The optimization problem is then to determine an optimal affine transformation T_{aff}^* that minimizes some error between Y and $\hat{Y} = T_{\text{aff}}^* X$. We posit the transformation equation as an augmented matrix one (also called *homogeneous coordinates*), for $\tilde{T}_{\text{aff}} : \mathbb{R}^{N+1} \rightarrow \mathbb{R}^{M+1}$:

$$\underbrace{\begin{pmatrix} Y \\ 1 \end{pmatrix}}_{\tilde{Y} \in \mathbb{R}^{M+1}} = \underbrace{\begin{bmatrix} A & | & B \\ \hline 0 \dots 0 & | & 1 \end{bmatrix}}_{\tilde{T}_{\text{aff}} \in \mathbb{R}^{(M+1) \times (N+1)}} \cdot \underbrace{\begin{pmatrix} X \\ 1 \end{pmatrix}}_{\tilde{X} \in \mathbb{R}^{N+1}}, \quad (2)$$

where $A \in \mathbb{R}^{M \times N}$ and $B \in \mathbb{R}^M$. We choose as error metric the Frobenius norm of $Y - \hat{Y}$:

$$\|Y - \hat{Y}\|_F = \sqrt{\sum_{d=1}^M \sum_{p=1}^P |Y_{dp} - \hat{Y}_{dp}|^2}. \quad (3)$$

Next, we tackle the question of the number of dimensions N that should be used for optimizing \tilde{T}_{aff} , as for any chosen value the error is minimized by the pseudo-inverse $\tilde{T}_{\text{aff}}^* = \tilde{Y} \tilde{X}^T (\tilde{X} \tilde{X}^T)^{-1}$ (Penrose, 1956; see Dokmanić and Gribonval, 2017, for further details). We know that at least the first three dimensions of the ideological space δ_1 , δ_2 and δ_3 (that we named LG, LR and IM) are correlated with some CHES dimension. To explore the gain in error reduction we compute it for different and growing values of N . The result of this exploration (see Fig. 9) is that adding ideological dimensions, beyond the first three identified in the previous section, contributes to reducing the error (up until $N = 7$). The error is of the order of 10^{-12} when $N = 7$, which is also the number of ideological dimension for which the system of (2) is determined: it is easy to check from (2) that there are $M(N+1)$ unknowns and MP equations, and that the system is under-determined for $N < P - 1$, over-determined for

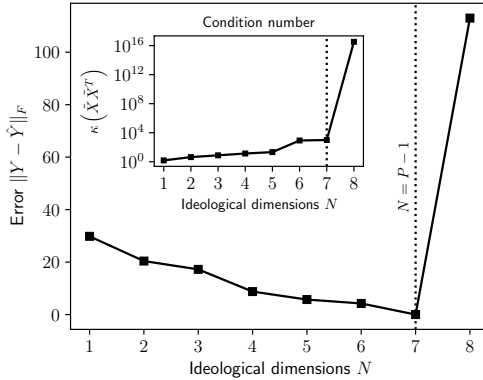


Fig. 9 Error in estimating position of $P = 8$ political parties in attitudinal reference CHES space (\hat{Y}) when using an affine transformation fitted with a varying number N of ideological dimensions, compared with the actual positions (Y), and the condition number of matrix inverted during fitting.

$N > P - 1$, and determined for $N = P - 1$. As seen in Fig. 9, over-determined systems produce much greater errors, accompanied by the deterioration of the condition number of the matrix that is inverted in the fit process.

Using the affine transformation fitted for $N = 7$, we map the position of all of our 368.831 users onto the CHES attitudinal reference space. In Fig. 10 we show the spatial distribution of these users, and of the MPs and their parties using four attitudinal dimensions of the CHES data: CHES Left-Right, CHES Anti-elite salience, CHES EU integration, and CHES Importance of ecology. All parties are inside the bounds of the CHES survey (which is answered by experts positioning parties within these predefined bounds). It is noteworthy that not all MPs are within these bounds, which is natural considering that some have more extreme positions than those of their parties, and that Twitter users might perceive it that way. This supposes that the position of parties are mean of the position of their MPs. Likewise, not all of their followers are within the bounds of the boundaries, as it is possible that many users are more extreme in their positions than the positions held by parties. Before extracting conclusions from the distribution of users in attitudinal space, we proceed to test their positions using text analysis.

5.2 Testing the positions of users with text analysis

As before, we use the self-descriptions written by users to validate that the dimensions are correctly positioning users along attitudinal dimensions. We begin by comparing the already selected labels from Section 4.3 (see Fig. 7) with the corresponding dimensions in the CHES attitudinal reference space.




Fig. 10 Spatial distribution of the MPs, their followers, and their parties using four attitudinal dimensions of the CHES data: Left-Right, Anti-élite salience, EU integration, and Importance of ecology.

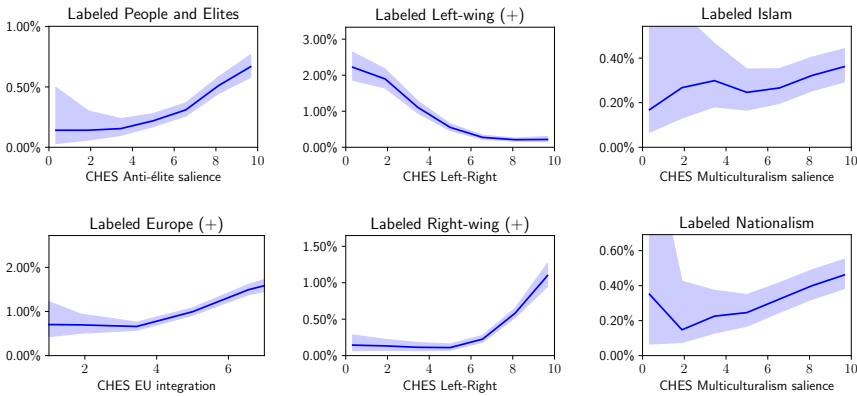


Fig. 11 Proportion of Twitter users labeled as referring to different chosen topics on their bios and according to their positions along corresponding dimensions of attitudinal space. Symbol + indicates that we consider only users indexed with the label and whose descriptions express a positive sentiment.

We observe (see Fig. 11) that the users featuring these labels are positioned coherently in the new CHES attitudinal space⁴.

Next, we show that there is a variety of newly available attitudinal dimensions related to new issues that can now be retrieved. To produce validations for additional dimensions, we extracted terms from the Twitter profile self-descriptions of users, and ranked them by term frequency. We then proceeded

⁴In comparison with Fig. 7, these new validations from Fig. 11 have wider confidence intervals for the labels “Islam” and “Nationalism” for values close to 0 in the CHES Multiculturalism salience dimensions. This is due to the fact that there are very few total users in that region of attitudinal space, but that we are including them in the figure for the sake of spanning the [0,10] interval that serves as reference for those CHES dimensions.

to associate these terms with new labels potentially relevant for CHES dimensions whenever possible, using a minimalist approach as in Section 4.3. This led to the identification of 29 labels that could be directly linked to 24 CHES attitudinal dimensions (from the available 48). For example, we identified the label “Ecology”, because 8.180 users included that word in their Twitter profiles, and because as a concept, ecology, is pertinent to two CHES dimensions: “Importance of ecology” (or salience of the issue) and “Economic growth over ecology” (meaning attitudes more favorable towards economic growth when confronted as an alternative to protecting the environment). We present the best 8 cases for which our labeling strategy offers the sharpest contrast (related to 5 new labels, see Fig.12) with their respective CHES dimensions. These 8 cases are related to the following additional labels:

- “Ecology”: defined by the words “écologie” or “écologiste”;
- “Tradition”: defined by the word “tradition”;
- “Agriculture”: defined by the words “agriculture” and “agriculteur”;
- “Social Security”: defined by “sécurité sociale” and “droits sociaux” (meaning *social rights*);
- “Entrepreneur”: defined by “entrepreneur”.

Fig. 12 shows the distribution of the proportion of users indexed with these different labels, according to pertinent related CHES attitudinal dimensions. When the CHES dimension is intended to measure positive or negative attitudes towards an attitudinal object, we further filter by sentiment, indicating the sentiment used to retain users with symbols + and -. For example, the “CHES Opposition to multiculturalism” is intended to distinguish cleavages between favorable and opposite view towards the attitudinal object, namely “multiculturalism”. When the CHES dimension is intended to measure the importance granted to an issue, or its salience, we do not filter by sentiment. This is the case, for example, of “Ecology” and CHES Ecology salience. A second case of use without sentiment is when the keywords already presuppose a stance. This is the case, for example of the label “Entrepreneur”. Many users define themselves as entrepreneurs on Twitter (2.185 in our dataset). We expect that people that are not entrepreneurs will not use the word them in their text profiles. Also we expect that people describing themselves as entrepreneurs might do so without negative or positive sentiment. The CHES GALTAN dimension⁵ refers to a socio-political cleavage dimension between Green-Alternative-Libertarian and Traditional-Authoritarian-Nationalist individuals or parties, introduced in 2002 in the context of research looking for cleavages beyond left-right stances (Hooghe et al., 2002). In the CHES survey, the question regarding the positioning of parties in the GALTAN dimension is framed as 0 being most liberal, 10 being most traditional, and 5 being the center.

⁵The full definition of CHES attitudinal dimensions can be seen in the survey’s *codebook*: <https://www.chesdata.eu/2019-chapel-hill-expert-survey>.

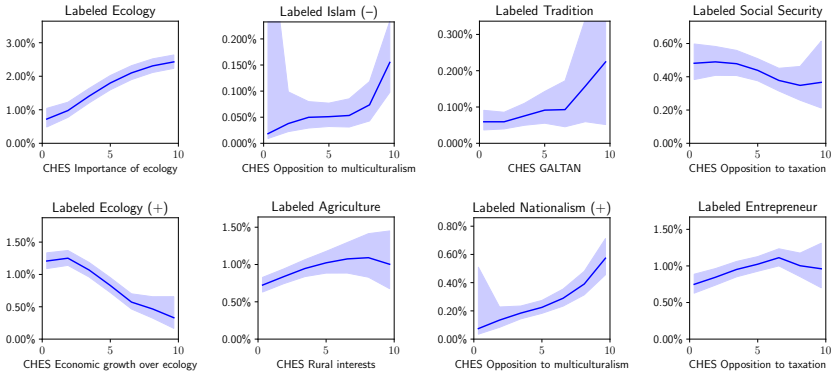


Fig. 12 Proportion of Twitter users labeled as referring to different topics on their profiles according to their positions along selected dimensions of the CHES attitudinal space. Symbols + and - indicate that the users with the label have been further filtered to keep those whose text expresses a positive or negative sentiment.

The recovered CHES dimensions in Fig. 12 confirm the main hypothesis behind our attitudinal embedding method: looking only at friendship networks (to the exclusion of textual data), in this case a bipartite network MPs and their followers, it is possible to infer the attitudes of these followers not only on left-right scales (which is the current state of the art), but also on a larger number of dimensions relating to several issues of public debate that are participating in social choice on internet platforms (in this case Twitter). In what follows, we will focus on a small set of CHES attitudinal dimensions for which the validation is the strongest (“CHES Left-Right”, “CHES Anti-elite salience”, “CHES EU integration” and “CHES Importance of ecology”), to illustrate how these attitudinal spaces can be leveraged in extracting conclusions of social facts from social networks network data.

6 Social Network Analysis in Attitudinal Spaces

Having computed and validated the embedding of the network in an attitudinal space where dimensions have 1) reference positions (such as a extreme-left, -right and center), 2) metrical consistency across space, and 3) dimensions explicitly related to single issues or cleavages, we now turn to the analysis of the 230.911 users for which we also know the subtended social graph, as described in Section 3.2. In particular, we are interested in the relation between topological communities, and the attitudinal stances of the members of these communities along some previously identified CHES attitudinal dimensions. To identify communities in the social graph of the directed friendship network on Twitter, we infer a degree-corrected stochastic block model by minimizing description length using *graph-tool*⁶ (Peixoto, 2014). The result is the partition

⁶https://graph-tool.skewed.de/static/doc/inference.html#graph_tool.inference.minimize_blockmodel_dl




Fig. 13 Scatterplots of the 1.067 communities of the Twitter social graph for four selected CHES attitudinal dimensions, according to the mean and standard deviation of the members of the communities (***) marks a p-value < 0.001 for Pearson correlation r).

of the 230.911 nodes of the network into 1.067 non-overlapping communities. Next, we select four CHES attitudinal dimensions (“CHES Left-Right”, “CHES Anti-elite salience”, “CHES EU integration” and “CHES Importance of ecology”) and compute, for each community and for each dimension, the mean position of the members of the community, and the standard deviation of the positions of the members. For the “CHES Left-Right” dimension, we consider the distance from the central position (defined as 5 on the 0 to 10 scales on the CHES survey), as we are interested on how extreme stances relate to community structures. Fig. 13 shows a scatterplot of the 1.067 community of the social graph, for the four selected CHES dimensions, according to the mean and standard deviation of their members. Red curves show the best polynomial model for the communities, quadratic for CHES Left-Right, linear for the others.

These observations are interesting in the context of the hypothesized but ill-defined *echo chambers* in social networks (Quattrociocchi et al., 2016; Baumann et al., 2020; Cinelli et al., 2021). Broadly speaking, *echo chamber* refers to a situation in which a group of users develops few interactions with contents and other users outside their group. This phenomenon holds importance as these relatively isolated users might develop, through selective exposure

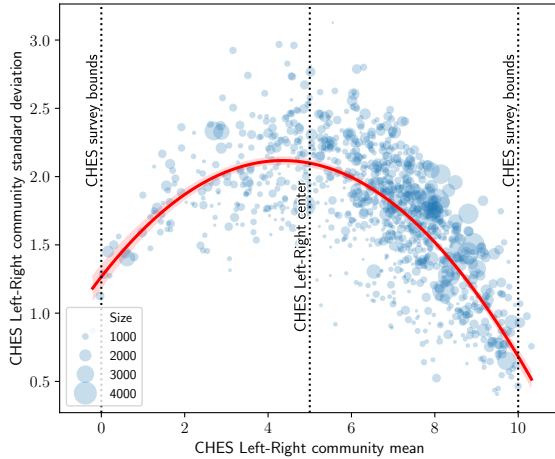


Fig. 14 Scatterplots of the 1.067 communities of the Twitter social graph on the CHES Left-Right dimension, according to the mean and standard deviation of the members of the communities. Extreme political communities are more homogeneous in their stances.

(Bryant and Miron, 2004; Kwak et al., 2010), stances and ultimately world-views that might be sufficiently different from the ensemble so as to obstruct social deliberation and coordination regarding common issues (*e.g.*, climate change; Williams et al., 2015). Having the ability to inspect the community structure of the social graph in some important attitudinal dimensions provides answers to some relevant questions. Fig. 13 shows that communities of users with extreme-left and extreme-right views, are also more homogeneous in their left-right stances. Communities of users with central left-right positions have a tendency to include a set of broader stances, showing a clear relation between *political polarization* (here understood as distance from the political left-right center) and homogeneity. This is seen more clearly in the distribution of communities along the CHES Left-Right axis (and not just according to distance from the center), as seen in Fig. 14, to which we have fitted a degree 2 polynomial. This is less the case, for example, for attitudes towards European integration or the importance of the ecology, where attitudes are comparatively less related with homogeneity of communities. The situation is yet slightly different regarding attitudes towards anti-elite sentiments. Communities of users that attribute little importance to the societal cleavage between people and elites, are formed by diverse users that have different stances on this issue. Communities of users that attribute great importance to this cleavage, are however much more homogeneous on their views towards this issue.

7 Discussion and Conclusions

In this article we have shown that network topology (without textual metadata of nodes) and external attitudinal data can be used to embed social

networks in ideological spaces where dimensions stand for indicators of attitudes towards sets of issues of public debate. We have illustrated this procedure using the social network formed by the followers of French parliamentarians on Twitter and leveraging the position of political parties in attitudinal dimensions of political surveys. This allowed us to embed the network in an ideological space spanned by three dimensions, positioning users according to their views on globalization and European integration, left-right positions, and views regarding immigration and multiculturalism. This so-called ideological space has the important property of ordering users: a user A to the left of user B in the left-right ideological dimension has high probability of being politically to the left of B. The greater the distance, the higher this probability is. Importantly, users far from moderate positions can be shown to be correctly positioned with high probability. This ideological embedding method already expands on the state of the art by allowing to identify several dimensions, associated with several issues of public debate that might be participating in social choice in internet platforms, where previous methods considered only one-dimensional differences. This method also has some limitations that we discussed in length before presenting a new improved method in Section 5. These are mainly three: 1) ideological spaces do not have reference points in space (*e.g.*, they lack central positions), 2) there is no consistent metric across space, and 3) dimensions act as indicators of attitudes towards a set of grouped issues (thus the name *ideological*). We were able to address these issues by proposing a second method called attitudinal embedding. Leveraging the relation between positions of a few reference points present in both ideological space and attitudinal reference spaces such as political surveys, we were able to map entities from the former onto the latter. Attitudinal embedding solves the aforementioned limitations and thus pave the way towards interesting applications, of which we showed one by examining both, the structure of the social graph and the attitudinal position of the nodes. Computing a community partition of the social graph and measuring the mean position and standard deviation of these communities in several attitudinal dimensions, we were able to show that communities with extreme political positions are more homogeneous than those with less extreme stances.

At this point we stress the main contributions and significance of our work in a summarized fashion. Our work shows that it is possible to infer several indicators of opinions for users related to issues of public debate, by exploiting only topological traces, namely social graphs, to the exclusion of textual data. Independence from textual meta-data means that our method can be applied in any context, and even on users that do not express themselves in written text, audio sound, or with images. It also means that this method can be used to compare users from different settings, countries, and different moments in time. Furthermore, since the method does not require users to provide answers by themselves, this method does not suffer from issues relating to question framing, or by the possible lack of truthfulness on the side of the respondents. This is a significant improvement with respect to the state of

the art, and paves the way for a greater understanding of opinion dynamics and related phenomena. Our attitude estimators are inherently probabilistic: from the probabilistic homophily model at its origin in Eq. (1), to the affine transformation that produces a mapping by interpolation in attitudinal space, to possibly additional undetermined sources or noise (*e.g.*, a user might follow a politician contrary to its own beliefs just to be updated on political issues, or even by error). While this might have obvious limitations in applications where the opinions of a single user are studied, it is certainly a powerful tool in studying the opinions of aggregates of users: the opinion and its dynamics for selected social groups, or even the whole of set of users.

Ideological and attitudinal spaces have many additional potential applications. They can be used to study ideological trajectories in time (Cointet et al., 2021; Ramaciotti Morales et al., 2021), and to study issues related to the meaning of ideological axes (Ramaciotti Morales and Muñoz Zolotoochin, 2022), such as the relative importance of left-right divides and polarization regarding globalization (Grossman and Sauger, 2019). They can also be used to study the dimensional structure (number and relative importance of dimensions) in different digital arenas (Benoit and Laver, 2012), and party systems online (Bakker et al., 2012). A different line of applications involves the study of the effect of Recommender Systems in phenomena such as polarization, now conceptualized in geometrical spaces (Ramaciotti Morales and Cointet, 2021), beyond traditional connectionist approaches. Having estimations of the stances of users to which recommendations are made can be leveraged to understand the effects of algorithmic recommendation in large socio-political digital systems (Bakshy et al., 2015). Finally, many recent studies analyzing the spread of misinformation online leverages political stances of users as an independent variable (Osmundsen et al., 2021). All these suggested applications, and others, may benefit from the inferences of stances offered by the ideological and attitudinal embedding methodologies here presented and tested.

Acknowledgments

This work has been funded by the “European Polarisation Observatory” (EPO) of the CIVICA Consortium, by the Data Intelligence Institute of Paris (diiP), and by the French National Agency for Research (ANR) under grants ANR-19-CE38-0006 “Geometry of Public Issues” (GOPI) and ANR-18-IDEX-0001 “IdEx Université de Paris”.

Data declared the 19 March 2020 and 15 July 2021 at the registry of data processing at the *Fondation Nationale de Sciences Politiques* (Sciences Po) in accordance with General Data Protection Regulation 2016/679 (GDPR) and Twitter policy. For further details and the respective legal notice, please visit <https://medialab.sciencespo.fr/en/activities/epo/>.

References

- Adorno, T., E. Frenkel-Brenswik, D.J. Levinson, and R.N. Sanford. 1950. *The authoritarian personality*. Harper.
- Ajzen, I. 1989. Attitude structure and behavior. *Attitude structure and function* .
- Ajzen, I. and M. Fishbein. 2005. The influence of attitudes on behavior. .
- Aldrich, J., S. Dorobantu, and M.A. Fernandez 2010. The use of the left-right scale in individual's voting decisions. In *Apsa 2010 annual meeting paper*.
- Bakker, R., L. Hooghe, S. Jolly, G. Marks, J. Polk, J. Rovny, M. Steenbergen, and M.A. Vachudova. 2020. 2019 chapel hill expert survey. *Chapel Hill* .
- Bakker, R., S. Jolly, and J. Polk. 2012. Complexity in the european party space: Exploring dimensionality with experts. *European Union Politics* .
- Bakshy, E., S. Messing, and L.A. Adamic. 2015. Exposure to ideologically diverse news and opinion on facebook. *Science* 348(6239): 1130–1132 .
- Barberá, P. 2015. Birds of the same feather tweet together: Bayesian ideal point estimation using twitter data. *Political analysis* .
- Barberá, P., J.T. Jost, J. Nagler, J.A. Tucker, and R. Bonneau. 2015. Tweeting from left to right: Is online political communication more than an echo chamber? *Psychological science* .
- Barberá, P. and G. Rivero. 2015. Understanding the political representativeness of twitter users. *Social Science Computer Review* .
- Baumann, F., P. Lorenz-Spreen, I.M. Sokolov, and M. Starnini. 2020. Modeling echo chambers and polarization dynamics in social networks. *Physical Review Letters* .
- Bem, D.J. 1970. Beliefs, attitudes, and human affairs. .
- Benoit, K. and M. Laver. 2012. The dimensionality of political space: Epistemological and methodological considerations. *European Union Politics* 13(2): 194–218 .
- Bond, R. and S. Messing. 2015. Quantifying social media's political space: Estimating ideology from publicly revealed preferences on facebook. *American Political Science Review* 19 .
- Briatte, F. and E. Gallic 2015. Recovering the french party space from twitter data. In *Science Po Quanti*.

- Bryant, J. and D. Miron. 2004. Theory and research in mass communication. *Journal of communication* .
- Budak, C. and D.J. Watts. 2015. Dissecting the spirit of gezi: Influence vs. selection in the occupy gezi movement. *Sociological Science* 2: 370–397 .
- Campbell, A., P.E. Converse, W.E. Miller, and D.E. Stokes. 1960. *The american voter*. Wiley.
- Carroll, J.D., E. Kumbasar, and A.K. Romney. 1997. An equivalence relation between correspondence analysis and classical metric multidimensional scaling for the recovery of euclidean distances. *British Journal of Mathematical and Statistical Psychology* .
- Cinelli, M., G.D.F. Morales, A. Galeazzi, W. Quattrociocchi, and M. Starnini. 2021. The echo chamber effect on social media. *Proceedings of the National Academy of Sciences* 118(9) .
- Cointet, J.P., P. Ramaciotti Morales, D. Cardon, C. Froio, A. Mogoutov, B. Ooghe-Tabanou, and G. Plique. 2021. What colours are the yellow vests? an ideological scaling of facebook groups. *Statistique et Société* .
- Converse, P.E. 1964. The nature of belief systems in mass publics. *Ideology and Discontent*, .
- Dokmanić, I. and R. Gribonval. 2017. Beyond moore-penrose part i: generalized inverses that minimize matrix norms. *arXiv preprint arXiv:1706.08349* .
- D’Esposito, M.R., D. De Stefano, and G. Ragozini. 2014. On the use of multiple correspondence analysis to visually explore affiliation networks. *Social Networks* .
- Falkenberg, M., A. Galeazzi, M. Torricelli, N. Di Marco, F. Larosa, M. Sas, A. Mekacher, W. Pearce, F. Zollo, W. Quattrociocchi, et al. 2021. Growing climate polarisation on social media. *arXiv preprint arXiv:2112.12137* .
- Flamino, J., A. Galezzi, S. Feldman, M.W. Macy, B. Cross, Z. Zhou, M. Serafino, A. Bovet, H.A. Makse, and B.K. Szymanski. 2021. Shifting polarization and twitter news influencers between two us presidential elections. *arXiv preprint arXiv:2111.02505* .
- Freedman, J. 2017. *Immigration and insecurity in France*. Routledge.
- Gentzkow, M., B. Kelly, and M. Taddy. 2019. Text as data. *Journal of Economic Literature* .
- Greenacre, M. 2017. *Correspondence analysis in practice*.

- Groseclose, T. and J. Milyo. 2005. A measure of media bias. *The Quarterly Journal of Economics* .
- Grossman, E. and N. Sauger. 2019. Economic internationalization and the decline of the left–right dimension. *Party Politics* .
- Hargreaves, A.G. 2007. *Multi-ethnic France: immigration, politics, culture and society*. Routledge.
- Hix, S., A. Noury, and G. Roland. 2006. Dimensions of politics in the european parliament. *American Journal of Political Science* .
- Hooghe, L., G. Marks, and C.J. Wilson. 2002. Does left/right structure party positions on european integration? *Comparative political studies* 35(8): 965–989 .
- Imai, K., J. Lo, J. Olmsted, et al. 2016. Fast estimation of ideal points with massive data. *American Political Science Review* .
- Kwak, H., J. An, E. Jing, and Y.Y. Ahn. 2021. Frameaxis: characterizing microframe bias and intensity with word embedding. *PeerJ Computer Science* 7: e644 .
- Kwak, H., C. Lee, H. Park, and S. Moon 2010. What is twitter, a social network or a news media? In *Proceedings of the 19th international conference on World wide web*, pp. 591–600.
- Lane, R.E. 1962. Political ideology: Why the american common man believes what he does. .
- Lazarsfeld, P.F., R.K. Merton, et al. 1954. Friendship as a social process: A substantive and methodological analysis. *Freedom and control in modern society* 18(1): 18–66 .
- Liu, B. and L. Zhang. 2012. A survey of opinion mining and sentiment analysis, *Mining text data*, 415–463. Springer.
- Lowe, W. 2008. Understanding wordscores. *Political Analysis* .
- McCarty, N., K.T. Poole, and H. Rosenthal. 2016. *Polarized America: The dance of ideology and unequal riches*. mit Press.
- McClosky, H. 1958. Conservatism and personality. *American Political Science Review* .
- Messaoudi, C., Z. Guessoum, and L. Ben Romdhane. 2022. Opinion mining in online social media: a survey. *Social Network Analysis and Mining* 12(1): 1–18 .

- Miklin, E. 2014. From ‘sleeping giant’ to left–right politicization? national party competition on the eu and the e uro crisis. *JCMS: Journal of Common Market Studies* 52(6): 1199–1206 .
- Osmundsen, M., A. Bor, P.B. Vahlstrup, A. Bechmann, and M.B. Petersen. 2021. Partisan polarization is the primary psychological motivation behind political fake news sharing on twitter. *American Political Science Review* 115(3): 999–1015 .
- Peixoto, T.P. 2014. Efficient monte carlo and greedy heuristic for the inference of stochastic block models. *Physical Review E* 89(1): 012804 .
- Penrose, R. 1956. On best approximate solutions of linear matrix equations. In *Mathematical Proceedings of the Cambridge Philosophical Society*, Volume 52, pp. 17–19. Cambridge University Press.
- Poole, K.T. and H. Rosenthal. 1984. The polarization of american politics. *The journal of politics* 46(4): 1061–1079 .
- Poole, K.T. and H. Rosenthal. 1985. A spatial model for legislative roll call analysis. *American Journal of Political Science* .
- Poole, K.T. and H. Rosenthal. 2000. *Congress: A political-economic history of roll call voting*. Oxford University Press on Demand.
- Quattrociocchi, W., A. Scala, and C.R. Sunstein. 2016. Echo chambers on facebook. *Available at SSRN 2795110* .
- Ramaciotti Morales, P. and J.P. Cointet 2021. Auditing the effect of social network recommendations on polarization in geometrical ideological spaces. In *15th ACM Conference on Recommender Systems, RecSys’ 21*.
- Ramaciotti Morales, P., J.P. Cointet, B. Benbouzid, D. Cardon, C. Froio, O.F. Metin, B. Ooghe, and G. Plique. 2021. Atlas multi-plateformes d’un mouvement social: Le cas des gilets jaunes. *Statistique et Société* .
- Ramaciotti Morales, P., J.P. Cointet, and J. Laborde 2020. Your most telling friends: Propagating latent ideological features on twitter using neighborhood coherence. In *2020 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*.
- Ramaciotti Morales, P., J.P. Cointet, and G. Muñoz Zolotoochin 2021. Unfolding the dimensionality structure of social networks in ideological embeddings. In *Proceedings of the 2021 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pp. 333–338.

- Ramaciotti Morales, P. and G. Muñoz Zolotoochin 2022. Measuring the accuracy of social network ideological embeddings using language models. In *2022 International Conference on Information Technology & Systems (ICITS22)*.
- Rheault, L. and C. Cochrane. 2020. Word embeddings for the analysis of ideological placement in parliamentary corpora. *Political Analysis* .
- Roberts Jr, J.M. 2000. Correspondence analysis of two-mode network data. *Social Networks* .
- Schön-Quinlivan, E. 2017. ‘the elephant in the room’no more: Europe as a structuring line of political cleavage in the 2017 presidential election. *French Politics* 15(3): 290–302 .
- Stefanov, P., K. Darwish, A. Atanasov, and P. Nakov 2020. Predicting the topical stance and political leaning of media using tweets. In *58th Annual Meeting of the Association for Computational Linguistics*.
- Van Dijk, T.A. 1998. *Ideology: A multidisciplinary approach*. Sage.
- Williams, H.T., J.R. McMurray, T. Kurz, and F.H. Lambert. 2015. Network analysis reveals open forums and echo chambers in social media discussions of climate change. *Global environmental change* 32: 126–138 .