



Unfolding the dimensionality structure of social networks in ideological embeddings

Pedro Ramaciotti Morales, Jean-Philippe Cointet, Gabriel Muñoz Zolotoochin

► To cite this version:

Pedro Ramaciotti Morales, Jean-Philippe Cointet, Gabriel Muñoz Zolotoochin. Unfolding the dimensionality structure of social networks in ideological embeddings. 2021 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM), Nov 2021, Amsterdam, Netherlands. hal-03315759v3

HAL Id: hal-03315759

<https://hal.science/hal-03315759v3>

Submitted on 6 Oct 2021

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.

Unfolding the dimensionality structure of social networks in ideological embeddings

Pedro Ramaciotti Morales
pedro.ramaciottimorales@sciencespo.fr
médialab, Sciences Po
Paris, France

Jean-Philippe Cointet
jeanphilippe.cointet@sciencespo.fr
médialab, Sciences Po
Paris, France

Gabriel Muñoz Zolotoochin
gabriel@mat.uc.cl
Facultad de Matemáticas
P. Universidad Católica de Chile

Abstract—Traditionally, public opinion on different issues of public debate has been studied through polls and surveys. Recent advancements in network ideological scaling methods, however, have shown that digital behavioral traces in social media platforms can be used to mine opinions at a massive scale. This has yet to be shown to work beyond one-dimensional opinion scales, which are best suited for two-party systems and binary social divides such as those observed in the US. In this article, we use multidimensional ideological scaling for coupled with referential attitudinal data for some nodes. We show that opinions can be mined in a multitude of issues: from social networks, embedding them in ideological spaces where dimensions stand for indicators of positive and negative opinions, towards issues of public debate. This method does not require text analysis and is thus language independent. We illustrate this approach on the Twitter follower network of French users leveraging political survey data.

Index Terms—Network scaling, graph embedding, ideology, attitudinal data, party systems.

I. INTRODUCTION

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 [1]. 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 [2], [3]. The traditional method for estimating people’s opinions 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 [4], in which they estimate attitudes of parliamentarians in the US towards liberal and conservative values using Bayesian inference on roll call data. Recently Barberà [5] 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 follow each other online), and that users’ opinions 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 [6]. This is due to the multi-issue, multi-party socio-political underlying systems that structure public debate [7], and ultimately social networks such as Twitter. In this article we use similar network scaling methods on large social network topological structures, 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. We illustrate this procedure with French Twitter data, and we propose benchmarks to test the validity of the approach. We also discuss its potential uses in several fields, such as: the study of populism, polarization, and the effects of Recommender Systems.

This work has been funded by the French National Agency for Research under grant ANR-19-CE38-0006 “Geometry of Public Issues” (GOPI) and CIVICA’s “European Polarisation Observatory” (EPO).

International Conference on Advances in Social Networks Analysis and Mining (ASONAM ’21), November 8–11, 2021, Virtual Event, Netherlands 978-1-4503-9128-3/21/11/\$15.00 © 2021

II. RELATED WORK

A diversity of methods for ideology estimation exist [8]. Many methods rely on textual analysis [9], [10], and are thus language- and context-dependent (*e.g.*, people express themselves differently on different platforms, and a model capturing ideology from parliamentary debates [11] will necessarily differ from one predicting political slant in tweets [12]). Other methods rely only on relational traces, producing ideological scaling (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 [4] is a classical example. Bond and Messing were among the first to apply ideological scaling methods to large social network graph data [13]. The work presented in this article is related to a similar network scaling method, proposed by Barberà [5] 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:

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

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 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 [14], [15] to be replaceable by Correspondence Analysis (CA) [16] on Twitter data [17]. CA also allows for the inference of multi-dimensional ideological parameters, embedding users in ideological spaces where positions are informative of their attitudes [18]. Recent works [19], [20] 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. In this article, we further explore the extraction of multi-dimensional attitudinal indicators suggested in [19]. In contrast with attitudinal embeddings relying on text analysis or news media citations [21], ideological embeddings relaying only social networks structures have the potential to be language- and context-independent, thus allowing, for example, comparisons between countries.

III. SOCIAL NETWORK DATA

To illustrate the extraction of attitudinal indicators using external referential attitudinal data we use French Twitter networks¹. We consider the set of the 831 (out of 925) French Members of Parliament (MPs) present on Twitter who are

¹Data declared the 19 March 2020 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.




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).

affiliated to 10 parties and their followers². 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 most 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 in [5] 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.

IV. HOMOPHILY EMBEDDING

We represent this bipartite subgraph as an adjacency matrix $A \in \{0, 1\}^{[368.831] \times [831]}$. Next, we produce a reduced-dimensionality representation of 831 MPs and the 368,831 followers using the coordinates of accounts in the Principal Components (PCs) of the CA of the adjacency matrix (see Fig. 2). Fig. 3 shows the inertia of the PCs of CA, and the density of the followers in the space spanned by the first 2 PCs of this reduced-dimensionality space. When most of the inertia is concentrated in the first dimensions or components, a spatial representation in those 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 would be similar to the original network [22]. As seen in Fig. 3, only a few dimensions hold relatively more importance in explaining the topological network data.

V. EXTERNAL REFERENTIAL ATTITUDINAL DATA

To inspect the relation between PCs and attitudes towards different issues of the public debate that might be influencing the network structure, we need to identify reference points in the Twitter network (and thus also in the latent embedding space). We then inspect the attitudinal positions of these references as given by external referential attitudinal data

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.

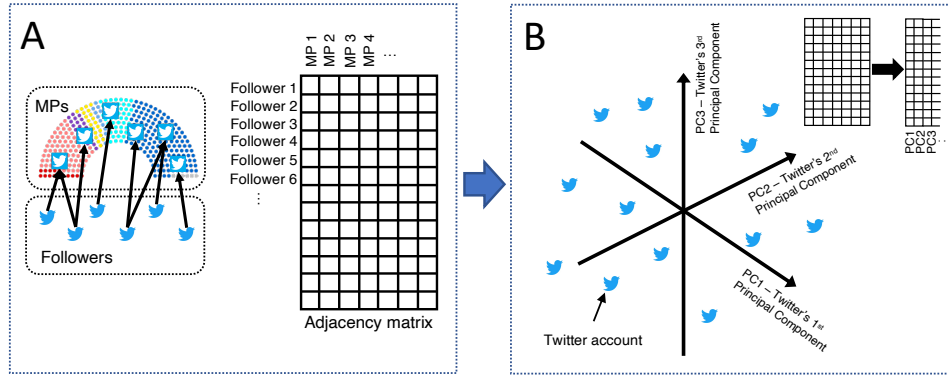


Fig. 2. Illustration of the homophily embedding through which the bipartite network of the French MPs on Twitter and their followers are represented in a multidimensional geometrical space. The bipartite network is represented as an adjacency matrix (A), to which Correspondence Analysis is applied to produce lower-dimensional spatial representations (B).




Fig. 3. Inertia of the principal components of the Correspondence Analysis of the bipartite network of French MPs on Twitter and their followers (left), and density of accounts on the first two components (right).

(e.g., surveys and polls). To do this, we compute the position of French political parties in this space as the centroids of the positions of the Twitter accounts of parliamentarians of these parties (see Fig. 4). We compare the position of parties according to our CA dimensions with the attitudinal positioning of parties on scales for different issues as provided by the 2019 Chapel Hill Expert Data (CHES) [23]. This attitudinal data source 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 for 51 different issues. The positioning in our ideological spaces is related, to different degrees, to the positioning by some of the attitudinal scales available in the CHES data. We measure the relation between positioning along the dimensions of the ideological space and the attitudinal positioning on issues given by CHES by computing Pearson correlations. Fig. 4 illustrates two example CHES attitudinal dimensions that are highly correlated with the positioning of parties in dimensions PC1 and PC2. In these examples, positions along PC1 are –negatively– correlated with attitudes towards trade protectionism: the higher a party is on the PC1 axis, the more opposed it is towards trade protectionism. Similarly, positions along PC2 are –positively– correlated with attitudes towards left- or right-leaning “ideological stances” (as referred to in

the CHES survey): the higher a party is on the PC2 axis, the more favorable it is towards right-wing ideological stances. This is an example of the aforementioned complementary attitudinal objects, as parties with lower positions on the PC2 axis are opposed to right-leaning ideological stances, and thus favorable to left-leaning stances.

VI. IDEOLOGICAL SPACE

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 PCs, with their position in the attitudinal CHES dimensions (scoring parties in a 0 to 10 scale, ranging from *very opposed* to *very favorable* to each issue). To compare a PC with a CHES dimension, we compute the Pearson correlation of party positions. We compare each one of the first 15 PCs (as suggested by Fig. 3) with all available CHES dimensions. Only the first 3 PCs show correlations with CHES dimensions, and these correlations show that each PC is associated (up to a correlation significance of $\alpha = 0.05$, marked in dark blue in Fig. 5) with a subset of CHES dimensions.

PC1 is positively correlated with positive attitudes towards the EU, suggesting that the higher the value of the PC1 position of a user, the more positive their views are towards the EU. PC1 is also associated with opposition to redistribution, interventionism, and market regulation. It is also associated with high levels of left-right economic policy dissent (inside parties), and with relevance granted to Russian interference in politics. Negative PC1 positions are correlated with positive views on economic protectionism, and 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 it the “Local-Global” (LG) axis. PC2 is associated with positions on issues widely attributable to left-right cleavages in France, including: left-right economic and ideological positions, rural-urban cleavages, religious principles, rights of minorities, authoritarianism, and the balance between crime-fighting and civil liberties among others. We call this second dimension the “Left-Right” (LR)

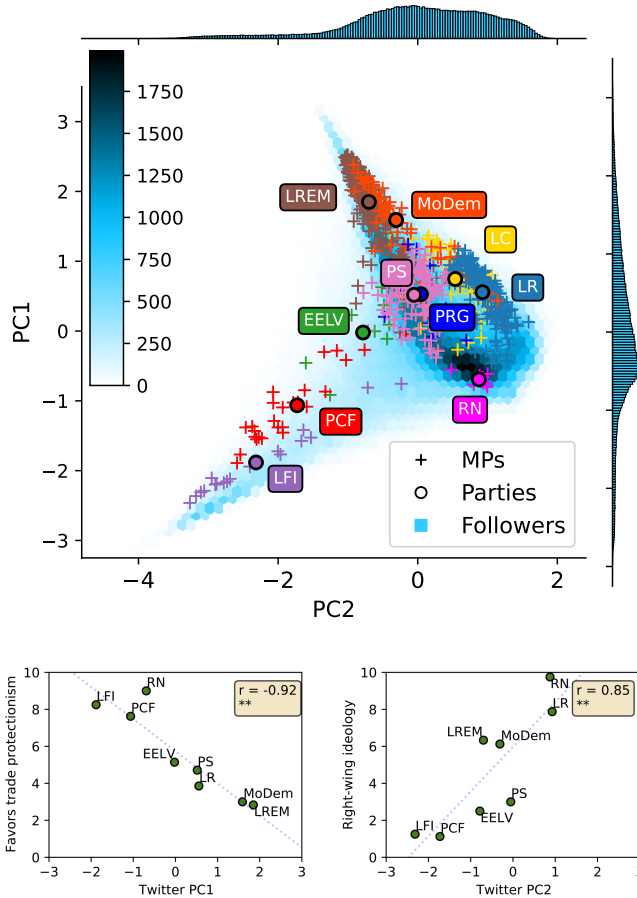


Fig. 4. Positioning of parties in the reduced-dimensionality space spanned by PC1 and PC2 of the CA of the bipartite Twitter network, computed as means of the positions of parliamentarians (top), and examples of two external party attitudinal positions (CHES) that are correlated with first two PC.

axis. Attitudes towards issues related to internationalization (e.g., trade protectionism, EU integration) have long been recognized as growing in importance in structuring individual preferences (e.g., in voting behavior [24]) when compared with the importance of attitudes towards issues distinguished in left-right cleavages (e.g., privatization, taxation, wealth-fare spending). The third PC to exhibit correlations with CHES dimensions is PC3, and is associated with views on the relative importance between leaders and members within parties, with economic policy blurriness, but mostly with importance granted to European integration, multiculturalism, and immigration restrictions. We call this third dimension the “Immigration and multiculturalism” (IM) axis. This third axis acts as an indicator of the importance granted to these issues (salience): lower positions in PC3 relate increased importance. Our axes span an *ideological space* in the sense that they position users according to attitudes towards a set of correlated issues. While the notion of *ideology* refers to “fuzzy” [25] set different concepts (including normative ones relating to power, e.g., the choice of rulers [26] or the justification of power [27]), we chose to use this term in its descriptive dimension:

as an “organization of opinions, attitudes and values” [28], a “structure of attitudes” [29], or as a description of high attitude consistency [30]. When a single variable is informative of a set of attitudes towards some issues because they display high spatial correlation, we call it “ideological dimensions”.

These associations were computed and justified, however, using the positions of parliamentarians, and we now seek to validate the definition of our LG, LR, and IM axes using the Twitter *bios* of their 368.831 followers, in which users briefly describe themselves. Twitter *bios* 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 bios 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 as mentioning the “EU” (if they include the strings “eu” or “europe”) and as mentioning the cleavage between “People and Elites” (if they include the strings “people”, meaning *people* in French, “élite”, or “politicien”). For the LR axis we classify 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 classify users as mentioning the “Islam” (if they include the strings “musulman”, meaning *muslim* in French, or “islam”) and as mentioning the “Multiculturalism” (if they include the string “multiculturalisme”). In the context of the French public debate, questions of multiculturalism and immigration policy rapidly revolve around a debate on Islam. In the strings that define a topic, we included all possible variants and misspellings. For the topics “EU”, “Left”, and “Right” we excluded bios with negative sentiment using a minimalistic approach by tagging the presence of the strings “anti”, “contre”, and “déteste”. For the other topics we did not filter by sentiments, as we are trying to detect the importance or salience of the topic along the dimensions (which does not involve positive or negative opinion). Fig. 6 shows the proportion of users that use these topics in their bios and the Clopper-Pearson confidence interval of this proportion at a $\alpha=0.05$ CI. As many studies suggest, left-right cleavages have a structuring role in behavior (e.g., voting behavior [31]) and in particular on French Twitter [32] (LR axis). In our dataset, however, the most structuring ideological dimension (LG) relates to globalization and European integration; it is the first PC that is related with these issues. In the rest of this article we will focus on the first two dimensions: PC1 and PC2.

VII. 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 non-

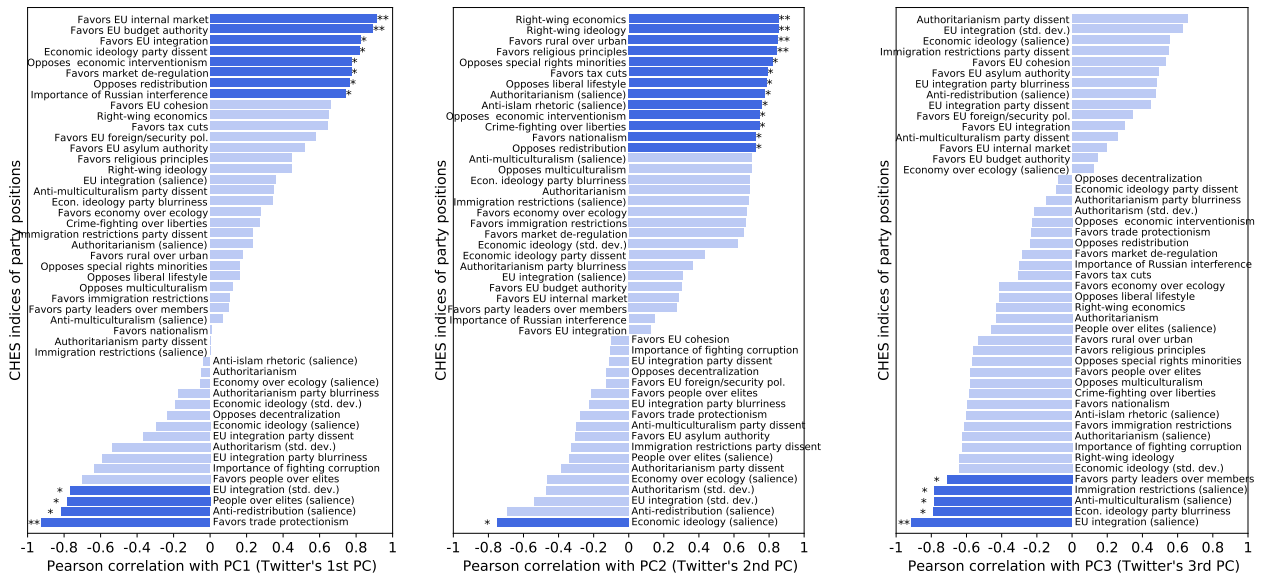


Fig. 5. Pearson correlations of the positioning of French parties on Twitter and according to CHES data for the first three Principal Components, and for all CHES dimensions, each associated with an issue of public debate. * marks 0.05 confidence and ** marks 0.01 confidence levels.

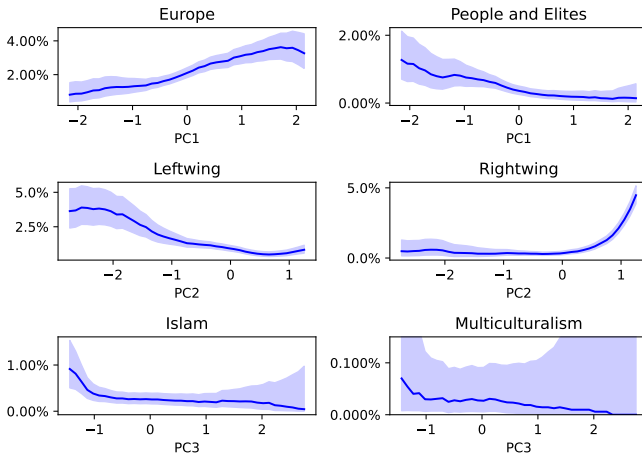


Fig. 6. Proportion of French Twitter users referring to different topics on their bios according to their positions along the three axes of ideological space.

negative mentions of parties in the Twitter *bio* or profile description. To do so, we again propose again a minimalistic approach to classifying users and 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”.

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 set of minimal identifiers of negative sentiment. Fig. 7 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.

VIII. DISCUSSION AND CONCLUSIONS

In this article we have shown that network topological data (to the exclusion of textual data) 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 French Twitter network, using the accounts of parliamentarians, and the CHES data on attitudinal positions for European political parties. This allowed us to embed the Twitter 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.

Ideological spaces have many potential applications. They can be used to study ideological trajectories in time [20], and to study issues related to the meaning of ideological axes, such as the relative importance of left-right divides and polarization regarding globalization [24]. They can also be used to study the dimensional structure (number and relative importance of dimensions) in different digital arenas [33], and party systems online [34]. A third line of applications involves




Fig. 7. Position of the users identified as sympathizers of each one of the ten parties in our dataset (identified with blue crosses +), and the position of the parties computed as the mean position of its MPs (shown in yellow dots •).

the study of the effect of Recommender Systems in phenomena such as polarization, now conceptualized in geometrical spaces [35], beyond traditional connectionist approaches. These applications must consider the representativity of the users of a given digital social network in the population [6]. The study of ideological configurations and dynamics of Twitter in itself presents interesting possibilities due to its importance in setting the political and media agenda.

REFERENCES

- [1] D. J. Bem, “Beliefs, attitudes, and human affairs.” 1970.
- [2] I. Ajzen, “Attitude structure and behavior,” *Attitude structure and function*, 1989.
- [3] I. Ajzen and M. Fishbein, “The influence of attitudes on behavior.” 2005.
- [4] K. T. Poole and H. Rosenthal, “A spatial model for legislative roll call analysis,” *American Journal of Political Science*, 1985.
- [5] P. Barberá, “Birds of the same feather tweet together: Bayesian ideal point estimation using twitter data,” *Political analysis*, 2015.
- [6] P. Barberá and G. Rivero, “Understanding the political representativeness of twitter users,” *Social Science Computer Review*, 2015.
- [7] S. Hix, A. Noury, and G. Roland, “Dimensions of politics in the european parliament,” *American Journal of Political Science*, 2006.
- [8] K. Imai, J. Lo, J. Olmsted *et al.*, “Fast estimation of ideal points with massive data,” *American Political Science Review*, 2016.
- [9] M. Gentzkow, B. Kelly, and M. Taddy, “Text as data,” *Journal of Economic Literature*, 2019.
- [10] T. Groseclose and J. Milyo, “A measure of media bias,” *The Quarterly Journal of Economics*, 2005.
- [11] L. Rheault and C. Cochrane, “Word embeddings for the analysis of ideological placement in parliamentary corpora,” *Political Analysis*, 2020.
- [12] P. Stefanov, K. Darwish, A. Atanasov, and P. Nakov, “Predicting the topical stance and political leaning of media using tweets,” in *58th Annual Meeting of the Association for Computational Linguistics*, 2020.
- [13] R. Bond and S. Messing, “Quantifying social media’s political space: Estimating ideology from publicly revealed preferences on facebook,” *American Political Science Review*, vol. 19, 2015.
- [14] W. Lowe, “Understanding wordscores,” *Political Analysis*, 2008.
- [15] J. D. Carroll, E. Kumbasar, and A. K. Romney, “An equivalence relation between correspondence analysis and classical metric multidimensional scaling for the recovery of euclidean distances,” *British Journal of Mathematical and Statistical Psychology*, 1997.
- [16] M. Greenacre, *Correspondence analysis in practice*, 2017.
- [17] P. Barberá, J. T. Jost, J. Nagler, J. A. Tucker, and R. Bonneau, “Tweeting from left to right: Is online political communication more than an echo chamber?” *Psychological science*, 2015.
- [18] M. R. D’Esposito, D. De Stefano, and G. Ragozini, “On the use of multiple correspondence analysis to visually explore affiliation networks,” *Social Networks*, 2014.
- [19] P. Ramaciotti Morales, J.-P. Cointet, and J. Laborde, “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 (ASONAM)*, 2020.
- [20] J.-P. Cointet, P. Ramaciotti Morales, D. Cardon, C. Froio, A. Mogoutov, B. Ooghe-Tabanou, and G. Plique, “What colours are the yellow vests? an ideological scaling of facebook groups,” *Statistique et Société*, 2020.
- [21] F. Baumann, P. Lorenz-Spreen, I. M. Sokolov, and M. Starnini, “Modeling echo chambers and polarization dynamics in social networks,” *Physical Review Letters*, 2020.
- [22] J. M. Roberts Jr, “Correspondence analysis of two-mode network data,” *Social Networks*, 2000.
- [23] R. Bakker, L. Hooghe, S. Jolly, G. Marks, J. Polk, J. Rovny, M. Steenbergen, and M. A. Vachudova, “2019 chapel hill expert survey,” *Chapel Hill*, 2020, www.chesdata.eu.
- [24] E. Grossman and N. Sauger, “Economic internationalization and the decline of the left-right dimension,” *Party Politics*, 2019.
- [25] T. A. Van Dijk, *Ideology: A multidisciplinary approach*. Sage, 1998.
- [26] R. E. Lane, “Political ideology: Why the american common man believes what he does.” 1962.
- [27] H. McClosky, “Conservatism and personality,” *American Political Science Review*, 1958.
- [28] T. Adorno, E. Frenkel-Brenswik, D. J. Levinson, and R. N. Sanford, *The authoritarian personality*. Harper, 1950.
- [29] A. Campbell, P. E. Converse, W. E. Miller, and D. E. Stokes, *The american voter*. Wiley, 1960.
- [30] P. E. Converse, “The nature of belief systems in mass publics,” *Ideology and Discontent*, 1964.
- [31] J. Aldrich, S. Dorobantu, and M. A. Fernandez, “The use of the left-right scale in individual’s voting decisions,” in *Apsa 2010 annual meeting paper*, 2010.
- [32] F. Briatte and E. Gallic, “Recovering the french party space from twitter data,” in *Science Po Quanti*, 2015.
- [33] K. Benoit and M. Laver, “The dimensionality of political space: Epistemological and methodological considerations,” *European Union Politics*, 2012.
- [34] R. Bakker, S. Jolly, and J. Polk, “Complexity in the european party space: Exploring dimensionality with experts,” *European Union Politics*, 2012.
- [35] P. Ramaciotti Morales and J.-P. Cointet, “Auditing the effect of social network recommendations on polarization in geometrical ideological spaces,” in *15th ACM Conference on Recommender Systems, RecSys’21*, 2021.