

Your most telling friends: Propagating latent ideological features on Twitter using neighborhood coherence

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

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

Julio Laborde
julio@recital.ai
reciTAL
Paris, France

Abstract—Multidimensional scaling in networks allows for the discovery of latent information about their structure by embedding nodes in some feature space. Ideological scaling for users in social networks such as Twitter is an example, but similar settings can include diverse applications in other networks and even media platforms or e-commerce. A growing literature of ideology scaling methods in social networks restricts the scaling procedure to nodes that provide interpretability of the feature space: on Twitter, it is common to consider the sub-network of parliamentarians and their followers. This allows to interpret inferred latent features as indices for ideology-related concepts inspecting the position of members of parliament. While effective in inferring meaningful features, this is generally restrained to these sub-networks, limiting interesting applications such as country-wide measurement of polarization and its evolution. We propose two methods to propagate ideological features beyond these sub-networks: one based on homophily (linked users have similar ideology), and the other on structural similarity (nodes with similar neighborhoods have similar ideologies). In our methods, we leverage the concept of neighborhood ideological coherence as a parameter for propagation. Using Twitter data, we produce an ideological scaling for 370K users, and analyze the two families of propagation methods on a population of 6.5M users. We find that, when coherence is considered, the ideology of a user is better estimated from those with similar neighborhoods, than from their immediate neighbors.

Index Terms—Multidimensional scaling, ideological scaling, latent features in social networks, political ideology, propagation in social networks.

I. INTRODUCTION

Methods for embedding networks have become ubiquitous tools for analysis [1], [2]. These methods exploit the geometrical representation of networks in a feature space, which can be used in tasks such as compression [3], clustering [4], visualization [5], link prediction [3], node classification [6], or scaling of a latent property that is determinant in the structure of networks. Examples of this latter application can be found in the use of multidimensional scaling of networks for the retrieval of latent features in music for artist recommendation [7], or in the use of Expectation Maximization algorithms for the estimation of ideological positions of voters, legislators,

or online users [8]. A growing domain of research has sought to exploit these scaling methods to infer ideological feature spaces in which to embed users of social networks. Taking inspiration in ideological scaling of voting data, pioneered by Poole et al. in the 1980s [9], and developed in the 1990s into the widely-used NOMINATE method [10], some adaptations to social network data have been proposed. The first such adaptation was proposed by Bond et al. [11], then popularized by Barbera in the mid- to late 2010s [12], [13]. These new scaling methods have been successful in identifying latent features in social networks related to the ideology of users. They are, however, often limited to small subsets of the whole network of users. The reason for this may be found in computational limitations in computing scaling for large networks, but mostly in the need for interpretability for the found features. On Twitter, traditionally, an ideological scaling is limited to a sub-graph of parliamentarians and their followers, using the embedded features of the first ones to provide an interpretation for the found features [14]. This raises the question: How to compute features for portions of the network outside this *seed* sub-graph? The setting of this problem is similar to that of problems such as latent feature or link prediction, that accounts for a wealth of works to be discussed in Section II. But crucially, it deviates from these settings in that a minority of nodes have *known* features (estimated via scaling), while those of a comparatively large set of nodes must be estimated from this initial seed set. This marks a difference with statistical learning methods, where a flexible model is learned on a majority of nodes, and then applied to a minority of nodes with missing values. The setting of ideological scaling invites the use of less flexible models, including strong assumptions about the link between ideology and the structure of the network.

This article takes inspiration in methods for information propagation in social networks and message coherence in telecommunications networks to propose a method for the propagation of scaled ideological features, from a seed set of nodes, to larger parts of a network. While this problem arises in the case of ideological scaling in social networks in particular, it can also suit a more general class of problems in

which n-dimensional features are only known for a subset of nodes in a network. One may consider that latent ideological features are defined only for users that follow members of parliament (MPs). However, our method supposes that a larger set of nodes could be positioned in the same latent feature space. This underlying hypothesis is leveraged by our method exploiting two strong assumptions: 1) the ideology of a node is similar to that of its neighbors if this neighborhood is ideologically coherent, and 2) the ideology of a node is similar to other structurally similar nodes, that follow the same users for example, if these followed users are ideologically coherent.

After providing a review of the relevant related work in Section II, we lay out the definitions and the notation needed for the treatment of the problem of coherent feature propagation in networks in Section III. Using this framework, we define the proposed propagation methods in Section IV. To analyze these methods, we develop an application case based on the ideology scaling on Twitter. Section V presents the case of a seed set of Twitter accounts, followers of MPs in France, with known ideological scaling computed using procedures well-established in the literature. Finally, in Section VI, we devise and execute experimental protocols to assess 1) the prediction capabilities using both of our methods following the two hypotheses in estimating ideological features, and 2) the limits of the estimation in terms of the trade-off between precision of the estimation, and its reach as parts of the whole network.

II. RELATED WORK

The methods proposed in this article are based in the propagation of features, from nodes for which they are known, to nodes for which they are not. This is mainly related to three different domains of research here reviewed. We intentionally omit the –vast– scientific literature regarding the identification of important and central nodes in networks, outside the scope of these selected domains.

1) *Latent ideological features in social networks*: The first category of related works include those that compute latent features in social networks. While several such methods exist, this article is concerned by those that leverage the interpretability of a subset of nodes to extract meaningful feature spaces. This stands in contrast to works that produce feature embeddings on which to perform statistical learning such as deep learning, without necessarily providing an interpretation for the embedded features (cf. [15]). A numerous family of methods, connected with link prediction, learn latent features in social networks, but specifically in relation with the probability of existence of a link, to be used in the propagation of local features into other parts of the time-evolving network [3], [16], or the retrieval of missing features for some nodes [6]. The most relevant works for this article, are those that use multidimensional scaling to extract ideological features [12], [13]. Other works seek to predict ideology for large networks, linking prediction to homophily, but use ideological labels and not a –continuous– ideological feature space [17]–[19].

2) *Propagation in social networks*: A second category of works pertinent for this study concerns the propagation in social networks, accounting for numerous results in a long tradition, but organized here in propagation of *information* and of *properties*. Propagation of information counts important cases such as the spread of rumors and misinformation [20], with connection to epidemiological models [21], [22]. This setting is fundamentally different from the one of this article in that it considers the network as the support for the flow of messages that have specific positions in time and space (as opposed to permanent latent features or properties). However, there are possible connections between network structure and homophily on the one hand, and structure in networks and flow of information on the other [23], [24]. This connection is also related, for example, to types of users in rumor propagation cascades in social networks [25].

Finally, propagation in social networks has connections with label propagation on graphs, used for example in heuristic procedures in community detection [26], [27]. However, again, this differs from the setting of this article in that a label can be equated to discrete classification, whereas features can provide (n-dimensional) continuous indices for nodes.

3) *Node coherence*: Node coherence is a concept that can be encountered in domains such as neuroscience [28], artificial intelligence [29], or communications networks [30], and often addresses the degree of dissimilarity of signals coming from different sources. This concept is of particular importance in multi-path networks. Depending on the type of network, coherence can be taken into account for node properties rather than relayed messages. The most relevant domain of works along this line, is found in node coherence in opinion dynamics, where an opinion is often modeled as a continuous feature [31]. In this domain, the concept of opinion coherence is sometimes considered in dynamic models [32]. In opinion dynamics, however, the focus is put in measuring different indices for node coherence to then be used in simulations to study evolution of features in systems.

III. PRELIMINARIES

This section proposes the notation and definitions required to treat the concept of node and neighborhood coherence in networks, before proposing feature propagation methods in Section IV.

A. The universe network

Let us consider a large *universe network* as a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, for some non-empty set \mathcal{V} with directed edges $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$. In social networks, nodes typically represent user accounts, and directed edges represent a relation of *following* between them. Information flows *downstream*, contrary to the direction of the edges: if a user shares information, it will be received by her followers and not by her *followees*¹.

For a node $v \in \mathcal{V}$, we consider the set of its *downstream* or in-neighbors $n_{\text{down}}(v) = \{u \in \mathcal{V} : (u, v) \in \mathcal{E}\}$ and its

¹Followees are also called *friends* on Twitter




Fig. 1. Convention used for *upstream* and *downstream* directions on directed social networks, where edges indicate following/friendship.

upstream or out-neighbors $n_{\text{up}}(v) = \{u \in \mathcal{V} : (v, u) \in \mathcal{E}\}$. Abusing notation, we define the *downstream* neighborhood of a set of nodes $V \subset \mathcal{V}$ as $n_{\text{down}}(V) = \{u \in \mathcal{V} : \exists v \in V ((u, v) \in \mathcal{E})\}$, and its *upstream* neighborhood as $n_{\text{up}}(V) = \{u \in \mathcal{V} : \exists v \in V ((v, u) \in \mathcal{E})\}$. It is worth noticing that, in general, we cannot assure $V \cap n_d(V) = \emptyset$, for $d \in \{\text{up}, \text{down}\}$. Given a direction d , we will denote its opposite direction with a bar, as \bar{d} .

B. Features of nodes and estimation error

Let us suppose that we can attribute N -dimensional features for some limited set of nodes $V \subset \mathcal{V}$. When available for a node $v \in \mathcal{V}$, we denote their known N -dimensional features by $e(v) \in \mathbb{R}^N$, and its i -th component by $e_i(v)$. When initially unavailable for a node $v \in \mathcal{V}$, N -dimensional features may be attributed through feature propagation from nodes for which they are known. Methods for doing so are at the core of this article, and will be proposed in Section IV. We denote estimated features for a node $v \in \mathcal{V}$ by $\hat{e}(v) \in \mathbb{R}^N$, and by $\hat{e}_i(v)$ its i -th component.

Whenever we can estimate features $\hat{e}(v)$ through propagation methods for a node $v \in \mathcal{V}$ for which we also know its *true* features $e(v)$, we can compute the estimation error $E(\hat{e}(v), e(v))$. We measure this estimation error as the p -norm in the feature space:

$$E(\hat{e}(v), e(v)) = \|\hat{e}(v) - e(v)\|_p. \quad (1)$$

For a set of nodes V , the mean estimation error will be denoted by $E(V) = (1/|V|) \sum_{v \in V} E(\hat{e}(v), e(v))$.

C. Node coherence

Node coherence is a key concept in this article, and seeks to capture the degree of similarity of a set of nodes in the feature space. Given $V \subset \mathcal{V}$ for which features are known or estimated, we denote by $I(V)$ the incoherence of V . Several incoherence metrics are possible, of which we propose to use the standard deviation of the distance to the centroid of V in the feature space:

$$I(V) = \sqrt{\frac{1}{|V|} \sum_{v \in V} \|e(v) - c_V\|_p^2}, \quad (2)$$

where $c_V = (1/|V|) \sum_{v \in V} e(v)$ is the centroid of V .

We prefer the definition of Equation (2) over other intuitive common metrics such as entropy [33]: a set V can have minimal entropy and high values of deviation $I(V)$ if the underlying probability distribution of V is multi-modal in the feature space.

D. ε -coherent neighborhoods of a set of nodes

Whenever we have a set of nodes $V \subset \mathcal{V}$, we can consider its up- or downstream neighborhoods of nodes that are also coherent: i.e., nodes that are linked to a set, and whose incoherence with respect to this set is below a given threshold ε . We formalize this notion, for both directions, as the ε -*coherent upstream* and *downstream neighborhood* of set V :

$$n_d^\varepsilon(V) = \{u \in n_d(V) : I(n_{\bar{d}}(u) \cap V) \leq \varepsilon\}$$

for $d \in \{\text{up}, \text{down}\}$, and \bar{d} the direction opposite to d .




Fig. 2. Illustration of the construction of a ε -coherent downstream neighborhood $n_{\text{down}}^\varepsilon(V)$ of a set V embedded in a 1-dimensional feature space.

IV. PROPAGATION METHODS FOR LATENT FEATURES IN SOCIAL NETWORKS

We consider two different approaches for propagating ideological features in a larger set of the network: Method A) directed sequences of ε -coherent neighborhoods, and Method B) sequences of projections using ε -coherent neighborhoods. Both families of methods are based on the concept of coherent neighborhoods, but motivated by distinct intuitions. Method A assumes that the features attached to nodes are “diffusing” along the social network formed by follower/followee relationships. This method is rooted in the larger homophily hypothesis “that birds of a feather flock together” [24], meaning that two connected users are likely to demonstrate homophilic behavior, sharing a similar ideological position. Method B has different premises. Now, if unknown, ideological features of a node are estimated using those of other nodes that occupy a similar position in the network. This method is based on the structural equivalence hypothesis, which posits that two nodes sharing the same neighborhood are similar [34].

The approach of this second family of methods is, for example, at the core collaborative filtering approaches in algorithmic recommendation. Here we will follow the same principle, estimating the ideological features of a user on Twitter, as the aggregation of features of users following, or being followed by the same users.

A. Directed sequences of ε -coherent neighborhoods

This method generates two sequences of sets of nodes: one of coherent nodes $\{V_i\}_{i \geq 0}$, for which we estimate the ideology, and one of incoherent nodes $\{\bar{V}_i\}_{i \geq 0}$, that we avoid using in ideology estimation.

Let us consider a seed set $V_0 \subset \mathcal{V}$ for which features $e(v)$ for $v \in V_0$ are known, and a direction $d \in \{\text{up}, \text{down}\}$. Starting at V_0 , a directed sequence of ε -coherent sets of nodes V_0, V_1, V_2, \dots is computed as $V_{i+1} = V_i \cup \Delta V_i$ for $i = 0, 1, 2, \dots$, with

$$\Delta V_i = \{v \in n_d^\varepsilon(V_i) : v \notin (V_i \cup \bar{V}_i)\}, \quad (3)$$

where $\bar{V}_{i+1} = \bar{V}_i \cup \Delta \bar{V}_i$ for $i = 0, 1, 2, \dots$, with $\bar{V}_0 = \emptyset$ and

$$\Delta \bar{V}_i = \{v \in n_d(V_i) : v \notin (\bar{V}_i \cup \Delta V_i)\}. \quad (4)$$

By definition, $\{V_i\}_{i \geq 0}$ is an incremental ($V_i \subset V_{i+1}$) by disjoint additions ($V_i \cap \Delta V_i = \emptyset$) sequence of ε -coherent neighbors. Sequence $\{\bar{V}_i\}_{i \geq 0}$ is also increasing ($\bar{V}_i \subset \bar{V}_{i+1}$) by disjoint additions ($\bar{V}_i \cap \Delta \bar{V}_i = \emptyset$), but of incoherent neighbors. Nodes deemed incoherent are stored to avoid coherence collisions: an incoherent node cannot become coherent by virtue of new nodes who had been attributed features by propagation. Whenever new ε -coherent neighbors ΔV_i are discovered at the i -th step, the features of their nodes are estimated as

$$\hat{e}(v) = \frac{1}{|n_{\bar{d}}(v) \cap V_i|} \sum_{u \in n_{\bar{d}}(v) \cap V_i} \hat{e}(u), \quad \text{for } v \in \Delta V_i, \quad (5)$$

setting $\hat{e}(v) = e(v)$ for $v \in V_0$.




Fig. 3. Schematic representation of the computation of ΔV_i and $\Delta \bar{V}_i$ on the i -th step of an upstream directed sequence of ε -coherent neighborhoods using **Method A**. In the upstream variant, users from V_i follow users from ΔV_i .

B. Sequences of projections using ε -coherent neighborhoods

Inspired by the way in which features are computed in multi-dimensional scaling for users depending on who they follow, we propose a second family of methods, Method B, for feature propagation using ε -coherent neighborhoods. Most ideology scaling methods exploit the collective structural relations between nodes that are important for the latent features of a network (typically politicians) and their followers [12], [13]. While intuitive, the first family of methods proposed in Section IV-A overlooks this aspect. This can have disadvantages: for example, some nodes can attract followers without themselves following many nodes coherently. This could arguably be the case of Twitter accounts of media outlets, institutions, or public figures. Accordingly, we propose a second family of methods centered around the notion of

structural similarity: nodes that are similar because they follow (upstream) or are followed (downstream) by the same users, hereafter called *pivots*. As with the previous method, Method B generates two sequences of sets of nodes: one of coherent nodes $\{V_i\}_{i \geq 0}$, for which we estimate the ideology, and one of incoherent nodes $\{\bar{V}_i\}_{i \geq 0}$ that we now avoid using as pivots.

Let us consider a seed set $V_0 \subset \mathcal{V}$ for which features are known and a direction $d \in \{\text{up}, \text{down}\}$. Starting at V_0 , a sequence of projected ε -coherent sets of nodes V_0, V_1, V_2, \dots is computed as $V_{i+1} = V_i \cup \Delta V_i$ for $i = 0, 1, 2, \dots$, for which we consider the set P_i^ε of ε -coherent pivots:

$$P_i^\varepsilon = n_d^\varepsilon(V_i) \setminus \bar{V}_i, \quad (6)$$

where $\bar{V}_{i+1} = \bar{V}_i \cup \Delta \bar{V}_i$ for $i = 0, 1, 2, \dots$, with $\bar{V}_0 = \emptyset$, and

$$\Delta \bar{V}_i = \{v \in n_d(V_i) : v \notin (n_d^\varepsilon(V_i) \cup \bar{V}_i)\}. \quad (7)$$

As with the previous family of Method A, the sets \bar{V}_i store the nodes deemed incoherent and that cannot be used, but now as pivots. The sets P_i^ε are used at each iteration to compute additions ΔV_i according to coherent structural similarity:

$$\Delta V_i = \{v \in n_d^\varepsilon(P_i^\varepsilon) : v \notin V_i\}. \quad (8)$$

Again by definition, sequences $\{V_i\}_{i \geq 0}$ and $\{\bar{V}_i\}_{i \geq 0}$ are incremental by disjoint additions. In contrast with the first family of Method A from Section IV-A, now it is the coherence of the pivot nodes in sets P_i^ε that is assured.

Whenever new neighbors ΔV_i are discovered at the i -th iteration, their features are estimated as

$$\hat{e}(v) = \frac{1}{|C_i(v, P_i^\varepsilon, V_i)|} \sum_{u \in C_i(v, P_i^\varepsilon, V_i)} \hat{e}(u), \quad \text{for } v \in \Delta V_i, \quad (9)$$

where $C_i(v, P_i^\varepsilon, V_i)$ is the set of co-neighbors of v in V_i through pivot P_i^ε :

$$C_i(v, P_i^\varepsilon, V_i) = V_i \cap n_{\bar{d}}(n_d(v) \cap P_i^\varepsilon). \quad (10)$$




Fig. 4. Schematic representation of the computation of ΔV_i and $\Delta \bar{V}_i$ using pivot P_i^ε on the i -th step of an upstream (co-followers-based) sequence of projected ε -coherent neighborhoods using **Method B**. In the upstream variant, users from V_i and ΔV_i co-follow users from pivot set P_i^ε .

V. FRENCH POLITICAL TWITTER DATASET

In order to analyze the feature propagation methods described in the previous section on real data, here we present a pertinent part of the Twitter network² and an ideological scaling producing an interpretable ideological feature space.

A. French MPs and their followers

Our data collection starts with the set \mathcal{P} of the 831 (out of 925) French MPs present on Twitter³, belonging to 10 main different parties or groups/alliances of parties. We then proceeded to collect all the followers of the accounts of MPs in \mathcal{P} . This collection was conducted on May 2019, and resulted in the constitution of the set $\mathcal{F} = n_{\text{down}}(\mathcal{P})$ of followers of \mathcal{P} , which amounts to $|\mathcal{F}| = 4.487.430$ unique Twitter accounts.

B. An ideological inference procedure via multidimensional scaling

We follow the methodology described by Barbera et al. [13] for inferring ideological features through multidimensional scaling of the sub-graph of the MPs and their followers. After removing from this sub-graph the followers that follow less than 3 MPs, and then removing users that had a repeated set of followed MPs (to assure full rank of the adjacency matrix), we obtained a set $\tilde{\mathcal{F}}$ of 368.831 accounts. We represent this sub-graph as a $\{0, 1\}^{|\tilde{\mathcal{F}}| \times |\mathcal{P}|}$ adjacency matrix, where a relation of following is encoded with the value 1, and its absence with the value 0. Next, we produce a reduced-dimensionality representation of these 368.831 observations using a Correspondence Analysis (CA) [35]. The first 2 principal components of the reduced-dimensionality space, PC1 and PC2, explain 0,97% and 1,07% of the inertia. This might seem as little explanation of the variability in the observed ways in which users from $\tilde{\mathcal{F}}$ follow MPs. However, the experience in several countries has revealed that the first components offer useful interpretations in the Twitter MPs' network. Examples include the UK, Spain, Italy, and Netherlands [12]). Similar results have also been found in France [36], [37]. In accordance with analyses made for the cases of other countries, those of France suggest feature interpretation where the first component is identified with left-right ideology.

C. Interpreting the ideological feature space

To provide an interpretation of PC1 and PC2 for our sets \mathcal{P} and $\tilde{\mathcal{F}}$, we consider exogenous variables for the parties of the MPs. We use the 2019 Chapel Hill Expert Survey (CHES) data [38]: an estimation on ideology and policy position of European parties, made by 421 political scientists specializing in political parties. Out of the 10 political parties identified

²In conformity with the General Data Protection Regulation 2016/679, the project called *Ideology scaling Twitter France*, whose data we exploit, has been declared the 19 Mars 2020 at the registry of data processing at the *Fondation Nationale de Sciences Politiques* (Sciences Po), and respects Twitter's data use policies.

³The list of French MPs present on social networks is provided by the French chambers of parliament at <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. 5. Distribution of the embedded ideological features PC1 and PC2 for the French MPs on Twitter per party. Parties are ordered according to the two criteria using the 2019 Chapel Hill Expert Survey data: left-right economics for feature PC1 (top), and attitudes toward European integration for feature PC2 (bottom).

for accounts in \mathcal{P} , 8 are also present in the CHES data. We computed the correlations between the positions of MPs on PC1 and PC2 of the feature space with all the 51 criteria were CHES data produces estimation for their political parties. This allowed to identify the two most relevant criteria related to our axes. Fig. 5 shows the ordering of these eight parties according to the two different criteria: 1) parties' economic views, from left to right, and 2) parties' attitudes towards European integration, from opposed to favorable. For each party in Fig. 5, we provide the embedded features of PC1 per MP for criterion 1 (left-right ideology), and the embedded features of PC2 per MP for criterion 2 (attitudes towards European integration). While it is not the objective of this article, Figure 5 illustrates the potential applicability of the multidimensional scaling procedure.

Under this interpretation, the French Twitter network of MPs and their followers (the sets \mathcal{P} and $\tilde{\mathcal{F}}$) may be represented in a 2-dimensional latent ideological feature space. In this feature space, the first dimension (PC1) provides an index for the concept of *left* and *right* positions, and the second dimension (PC2) provides an index for attitudes towards European integration, which is in line with observations that have been made in the case of other countries⁴. Fig. 6 illustrates the positions of the sets \mathcal{P} of parliamentarians and $\tilde{\mathcal{F}}$ of their followers in this bidimensional ideological feature space.

⁴Research results in several countries have also identified this second axis as related with issues dividing public along pro- and anti-establishment positions [39], [40].

VI. NUMERICAL EXPERIMENTS




Fig. 6. Position of \mathcal{P} French MPs in the first two embedded dimensions, PC1 (left to right ideology) and PC2 (pro- and against European integration), histogram of the positions of their followers $\tilde{\mathcal{F}}$, and positions of sampled accounts in set \mathcal{A} .

D. Datasets for experimentation

We are now concerned with the problem of using the set $\tilde{\mathcal{F}}$ with 368.831 accounts with known features to establish datasets for the testing and evaluation. We cannot set $V_0 = \tilde{\mathcal{F}}$ for testing and evaluation. This is because when producing the next set V_1 in the sequence (with either method) the new added nodes would not have known *true* features $e(v)$ (as computed with ideological scaling) to test the estimated ones $\hat{e}(v)$, for $v \in \Delta V_0$. To circumvent this difficulty, we take a subset $\mathcal{A} \subset \tilde{\mathcal{F}}$, and then collect its upstream and downstream neighborhoods to use in the described methods, allowing for some elements of these new neighborhoods to be also in $\tilde{\mathcal{F}}$. To account for the possible specificities of the different regions of the ideological feature, we sample 4.483^5 nodes uniformly in space from $\tilde{\mathcal{F}}$ to produce set \mathcal{A} (see Fig. 6). Had we not sampled with spatial uniformity, the sampling would have produced a majority of nodes near the origin of the ideological feature space (see the concentration of nodes near the origin in Fig. 6). Next, we collect the followees/friends of \mathcal{A} as $\mathcal{B} = n_{\text{up}}(\mathcal{A})$ (obtaining $|\mathcal{B}| = 1.304.812$), and the followers of set \mathcal{A} as $\mathcal{C} = n_{\text{down}}(\mathcal{A})$ (obtaining $|\mathcal{C}| = 5.528.716$). This sub-sampling and collection operations achieve sets for our test evaluations that are such that $|\mathcal{B} \cap \tilde{\mathcal{F}}| = 138.424$, and $|\mathcal{C} \cap \tilde{\mathcal{F}}| = 231.035$.

⁵Originally, we sampled 5.000 nodes, but the later collection of followers and followees for 517 of these nodes was not possible because these accounts were either closed between collections, or private and information could not be collected.

In this section we analyze some results of the methods proposed in Section IV for estimating ideological features on Twitter data. For both methods, we center the analysis around the set \mathcal{A} built for this purpose, and described in the previous section. We will be interested in analyzing, for different values of coherence ε , 1) the accuracy (the degree of error in estimating ideological features), and 2) the coverage (the size of the disjoint additions with estimated features) of the proposed methods. The main aspect we seek to investigate in this section is the relation between the coherence threshold parameter ε and the trade-off between estimation error and size of a disjoint addition ΔV_i at a step i . Throughout this section, we compute error E (cf. Equation (1)) and coherence I (cf. Equation (2)) using $p = 2$.

A. Directed sequences of ε -coherent neighborhoods

Because we have collected the upstream and downstream neighborhoods of set \mathcal{A} (sets \mathcal{B} and \mathcal{C}), we can readily compute the first step for Method A in both directions. Figure 7 reports the error $E(\Delta V_0)$ and the size of ΔV_0 for the first step ($i = 0$).




Fig. 7. Mean estimation error $E(\Delta V_0)$ and size of ΔV_0 for Method A in the upstream and downstream directions, using $V_0 = \mathcal{A}$.

While intuitive, the estimation of ideological features through directed propagation of Method A in either direction performs poorly. Imposing higher levels of coherence (lower thresholds ε for incoherence) does improve the mean estimation error of nodes in ΔV_0 . However, this improvement is marginal and, most importantly, limited. Upstream neighbors (friends/followees) of users in \mathcal{A} that are followed by highly coherent users, have their ideology estimated with a mean error no lower than $E(\Delta V_0) = 0.6$, which amounts to 12.5% of the diameter of $\tilde{\mathcal{F}}$ in the feature space. Similarly, downstream neighbors (followers) of users in \mathcal{A} that follow highly coherent users, have their ideology estimated with a mean error no lower than $E(\Delta V_0) = 0.575$ (12.0% of the diameter of $\tilde{\mathcal{F}}$). Had set ΔV_1 be computed (upstream or downstream), its ideological features would be estimated only from those already –poorly– estimated for set ΔV_0 (disjoint additions, assured by Equation (3), imply that $n_d(V_0) \cap \Delta V_1 = \emptyset$).

B. Sequences of projections using ε -coherent neighborhoods

The setting of Method B is similar to that of collaborative filtering (CF) [41] in Recommender Systems. In user-based CF, the similarity of users is computed according to structural similarity on a bipartite graph: users are similar if they have chosen similar items. Items can then be proposed to a user among those already chosen by similar users (structural similarity is related to preference). In the setting of our Method B, users are deemed similar if they follow, or are followed by (depending if the selected direction is upstream or downstream) similar users of a so-called pivot set for a predetermined coherence ε . Taking on known evaluation protocols for CF [42], we propose a method for assessing the accuracy and coverage of Method B. Given a positive integer K , we perform a K -fold bipartite cross-validation. We divide \mathcal{A} in K parts, taking one as $\mathcal{A}_{\text{test}}^K$ and the rest as $\mathcal{A}_{\text{train}}^K$. We set $V_0 = \mathcal{A}_{\text{train}}^K$, we compute P_0^ε using \mathcal{B} and \mathcal{C} for upstream and downstream directions, and then set $\mathcal{R} = \mathcal{A}_{\text{test}}^K \cap \Delta V_0$. To analyze the coverage, we examine the quantity $\text{cov}(\mathcal{R}) = |\mathcal{R}|/|\mathcal{A}_{\text{test}}^K|$, ranging from 0 (no coverage) to 1 (total coverage). We choose $K = 20$ and report the results for varying values of incoherence ε , providing, for the ensemble of the 20 folds, the median, the maximum, and the minimum value in Fig. 8.

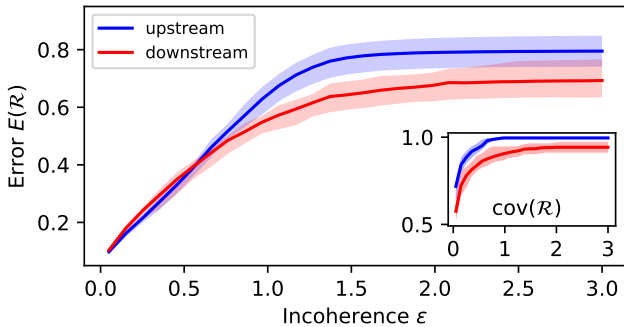


Fig. 8. Estimation error and coverage of set $\mathcal{R} = \mathcal{A}_{\text{test}}^K \cap \Delta V_0$ and P_0 for the Method B in a K -fold cross-validation retrieval of ideologies of \mathcal{A} ($K = 20$). The distribution of errors and coverage over the folds is reported with the median, the minimum, and the maximum values .

Results obtained with Method B, reported in Fig. 8, are more satisfactory than those obtained with Method A. Thresholds for incoherence below (approximately) $\varepsilon = 1$ allow already for less error in estimation in comparison to Method A. In fact, under our setting, estimation errors for Method B can be made as low as (approximately) 0.1 in the feature space by sacrificing coverage of new nodes added in the following step of the sequence. With a small value of threshold for ε (lower than .1) the estimation error is around 0.1 in distance in the feature space (2.1% of the diameter of $\tilde{\mathcal{F}}$). This trade-off in coverage does not come at great expense, as suggested by our proposed metric $\text{cov}(\mathcal{R})$: at least 50% of nodes left in the test set $\mathcal{A}_{\text{test}}^K$ can be recovered, independently of the threshold for incoherence. The amount of nodes in the set of

pivots P_0^ε is that of ΔV_0 for Method A, and follows intuition in that, when forcing less incoherence, less nodes are available for the search of co-neighbors (cf. Equation (10)).

While it was possible to evaluate the quality of Method A precisely, the same evaluation would have required the collection of the downstream neighbors of \mathcal{B} (i.e. $n_{\text{down}}(n_{\text{up}}(\mathcal{A}))$) and the upstream neighbors of \mathcal{C} (i.e. $n_{\text{up}}(n_{\text{down}}(\mathcal{A}))$), which would amount to a mixed-direction snowball-sampling of the Twitter network from \mathcal{A} . Thus the need for a different evaluation protocol for both methods. While there is a slight risk of over-estimating the coverage capacity of Method B using our testing protocol (due to the fact that \mathcal{R} is restricted to \mathcal{A}) it has the great advantage of allowing for experimentation with large sets: the (up- and downstream) neighbors of \mathcal{A} , i.e., the total population for feature propagation, account for nearly 6.5M users.

VII. CONCLUSIONS

We proposed two methods (A and B) for the propagation of latent ideological features on Twitter. These methods use propagation for the estimation of ideology of users, represented in a continuous n-dimensional features space. Method A is based on the homophily hypothesis: if users are connected (one follows the other) their ideology is similar. Method B is based on the structural similarity hypothesis: if users are connected to the same neighborhoods their ideology is similar. In these methods, we proposed the consideration and formalization of these notions when depending on ideological coherence of neighbors, which we model with a single parameter.

To analyze our methods, we collected Twitter data using the accounts of French MPs. We were able to produce a 2-dimensional ideological embedding for a sub-graph of nearly 370K users. With the help of the Chapel Hill Expert Survey, we validate the interpretation of the two emerging dimensions as related to left-right ideology, and attitudes toward European integration. Collecting again neighborhoods of some of these users we used our methods to propagate these ideological features with predefined degrees of accuracy within a potential population of nearly 6.5M users. To total number of users for which ideology is propagated, depends on the error to be accepted, which is determined by a coherence threshold ε . The protocols established for evaluation allow us to analyze the relation between estimation accuracy for ideology in propagation, and the coverage of the method (i.e., the number of users to which we can propagate). The trade-off between these two competing and desirable properties was modulated by our coherence parameter. Analyzing Methods A and B, we find that, when coherence is considered, the ideology of a user on Twitter is better estimated using other users that are structurally similar (Method B), than using other users that might be directly connected (Method A). These results lend support to the structural similarity hypothesis over the homophily hypothesis. Even more: while the difference is small, the ideology of a user is better estimated using users that follow the same accounts, than using users that are followed the same Twitter accounts.

The formalism used to treat coherence in social networks, and the methods proposed for propagation, can be extended to any other network where directed edges model the fact a user can receive information from another. On Facebook, for example, the fact that user a and user b are *friends*, could be represented by two directed edges between both users, and in different directions. In other networks, such as Instagram or YouTube (using the *subscription* relation) the formalism and methods can be directly applied.

ACKNOWLEDGMENTS

We thank Justin Clarke and Benjamin Ooghe-Tabanou for their help in the collection of Twitter data.

REFERENCES

- [1] P. Goyal and E. Ferrara, "Graph embedding techniques, applications, and performance: A survey," *Knowledge-Based Systems*, vol. 151, pp. 78–94, 2018.
- [2] H. Cai, V. W. Zheng, and K. C.-C. Chang, "A comprehensive survey of graph embedding: Problems, techniques, and applications," *IEEE Transactions on Knowledge and Data Engineering*, vol. 30, no. 9, pp. 1616–1637, 2018.
- [3] D. Wang, P. Cui, and W. Zhu, "Structural deep network embedding," in *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 2016, pp. 1225–1234.
- [4] S. White and P. Smyth, "A spectral clustering approach to finding communities in graphs," in *Proceedings of the 2005 SIAM international conference on data mining*, 2005, pp. 274–285.
- [5] K. Pearson, "Liii. on lines and planes of closest fit to systems of points in space," *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, vol. 2, no. 11, pp. 559–572, 1901.
- [6] J. Tang, M. Qu, M. Wang, M. Zhang, J. Yan, and Q. Mei, "Line: Large-scale information network embedding," in *Proceedings of the 24th international conference on world wide web*, 2015, pp. 1067–1077.
- [7] J. C. Platt, "Fast embedding of sparse similarity graphs," in *Advances in neural information processing systems*, 2004, pp. 571–578.
- [8] K. Imai, J. Lo, J. Olmsted *et al.*, "Fast estimation of ideal points with massive data," *American Political Science Review*, vol. 110, no. 4, pp. 631–656, 2016.
- [9] K. T. Poole and H. Rosenthal, "A spatial model for legislative roll call analysis," *American Journal of Political Science*, pp. 357–384, 1985.
- [10] —, "Patterns of congressional voting," *American journal of political science*, pp. 228–278, 1991.
- [11] R. Bond and S. Messing, "Quantifying social media's political space: Estimating ideology from publicly revealed preferences on facebook," *American Political Science Review*, vol. 109, no. 1, pp. 62–78, 2015.
- [12] P. Barberá, "Birds of the same feather tweet together: Bayesian ideal point estimation using twitter data," *Political analysis*, vol. 23, no. 1, pp. 76–91, 2015.
- [13] 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*, vol. 26, no. 10, pp. 1531–1542, 2015.
- [14] P. Barberá and G. Rivero, "Understanding the political representativeness of twitter users," *Social Science Computer Review*, vol. 33, no. 6, pp. 712–729, 2015.
- [15] S. Yan, D. Xu, B. Zhang, H.-J. Zhang, Q. Yang, and S. Lin, "Graph embedding and extensions: A general framework for dimensionality reduction," *IEEE transactions on pattern analysis and machine intelligence*, vol. 29, no. 1, pp. 40–51, 2006.
- [16] C. Heaukulani and Z. Ghahramani, "Dynamic probabilistic models for latent feature propagation in social networks," in *International Conference on Machine Learning*, 2013, pp. 275–283.
- [17] E. Colleoni, A. Rozza, and A. Arvidsson, "Echo chamber or public sphere? predicting political orientation and measuring political homophily in twitter using big data," *Journal of communication*, vol. 64, no. 2, pp. 317–332, 2014.
- [18] I. Himmelboim, K. D. Sweetser, S. F. Tinkham, K. Cameron, M. Danelo, and K. West, "Valence-based homophily on twitter: Network analysis of emotions and political talk in the 2012 presidential election," *New media & society*, vol. 18, no. 7, pp. 1382–1400, 2016.
- [19] Z. Xiao, W. Song, H. Xu, Z. Ren, and Y. Sun, "Timme: Twitter ideology-detection via multi-task multi-relational embedding," in *Proceedings of the ACM SIGKDD international conference on Knowledge Discovery and Data Mining*, 2020.
- [20] Y. Bao, C. Yi, Y. Xue, and Y. Dong, "A new rumor propagation model and control strategy on social networks," in *2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, 2013, pp. 1472–1473.
- [21] F. D. Sahneh and C. Scoglio, "Competitive epidemic spreading over arbitrary multilayer networks," *Physical Review E*, vol. 89, no. 6, p. 062817, 2014.
- [22] D. Chakrabarti, Y. Wang, C. Wang, J. Leskovec, and C. Faloutsos, "Epidemic thresholds in real networks," *ACM Transactions on Information and System Security (TISSEC)*, vol. 10, no. 4, pp. 1–26, 2008.
- [23] S. Aral, L. Muchnik, and A. Sundararajan, "Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks," *Proceedings of the National Academy of Sciences*, vol. 106, no. 51, pp. 21 544–21 549, 2009.
- [24] M. McPherson, L. Smith-Lovin, and J. M. Cook, "Birds of a feather: Homophily in social networks," *Annual review of sociology*, vol. 27, no. 1, pp. 415–444, 2001.
- [25] A. Friggeri, L. Adamic, D. Eckles, and J. Cheng, "Rumor cascades," in *Eighth International AAAI Conference on Weblogs and Social Media*, 2014.
- [26] X. Zhu and Z. Ghahramani, "Learning from labeled and unlabeled data with label propagation," 2002.
- [27] S.-M. Cheng, W. C. Ao, P.-Y. Chen, and K.-C. Chen, "On modeling malware propagation in generalized social networks," *IEEE Communications Letters*, vol. 15, no. 1, pp. 25–27, 2010.
- [28] P. Thagard and K. Verbeurgt, "Coherence as constraint satisfaction," *Cognitive Science*, vol. 22, no. 1, pp. 1–24, 1998.
- [29] A. K. Mackworth, "Consistency in networks of relations," *Artificial intelligence*, vol. 8, no. 1, pp. 99–118, 1977.
- [30] P. Larsson and H. Rong, "Large-scale cooperative relaying network with optimal coherent combining under aggregate relay power constraints," in *WWRF 2004*, 2004.
- [31] B. Edmonds, "Modeling belief change in a population using explanatory coherence," *Advances in Complex Systems*, vol. 15, no. 06, p. 1250085, 2012.
- [32] N. Rodriguez, J. Bollen, and Y.-Y. Ahn, "Collective dynamics of belief evolution under cognitive coherence and social conformity," *PLoS one*, vol. 11, no. 11, p. e0165910, 2016.
- [33] V. Onesto, M. Romano, F. Gentile, and F. Amato, "Relating the small world coefficient to the entropy of 2d networks and applications in neuromorphic engineering," *Journal of Physics Communications*, vol. 3, no. 9, p. 095011, 2019.
- [34] L. D. Sailer, "Structural equivalence: Meaning and definition, computation and application," *Social Networks*, vol. 1, no. 1, pp. 73–90, 1978.
- [35] M. Greenacre, *Correspondence analysis in practice*, 2017.
- [36] F. Briatte and E. Gallic, "Recovering the french party space from twitter data," 2015.
- [37] J. P. Cointet, P. Ramaciotti Morales, D. Cardon, C. Froio, A. Mogoutov, B. Ooghe-Tabanou, and G. Plique, "What color are the yellow vests ? : Ideological scaling of facebook groups," *Statistique et Société*, 2020, upcoming article.
- [38] 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.
- [39] A. L. Schmidt, A. Peruzzi, A. Scala, M. Cinelli, P. Pomerantsev, A. Applebaum, S. Gaston, N. Fusi, Z. Peterson, G. Severgnini *et al.*, "Measuring social response to different journalistic techniques on facebook," *Humanities and Social Sciences Communications*, vol. 7, no. 1, pp. 1–7, 2020.
- [40] J. Chiche, B. Le Roux, P. Perrineau, and H. Rouanet, "L'espace politique des électeurs français à la fin des années 1990: nouveaux et anciens clivages, hétérogénéité des électors," *Revue française de science politique*, pp. 463–487, 2000.
- [41] J. B. Schafer, D. Frankowski, J. Herlocker, and S. Sen, "Collaborative filtering recommender systems," in *The adaptive web*, 2007, pp. 291–324.
- [42] J. Bobadilla, F. Ortega, A. Hernando, and J. Bernal, "A collaborative filtering approach to mitigate the new user cold start problem," *Knowledge-based systems*, vol. 26, pp. 225–238, 2012.