2020 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)

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

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

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

Julio Laborde
julio@recital.ai
reciTAL

***Abstract*—A growing literature on ideology estimation through scaling methods in social networks restricts the scaling procedure to nodes that provide interpretability of the resulting feature space. On Twitter, for example, it is common to consider the sub-network of parliamentarians and their followers. While effective in inferring meaningful ideological features, this restriction limits interesting applications such as country-wide measurement of polarization and its evolution. We propose two methods to propagate ideological features beyond these sub-networks: based on homophily (linked users have similar ideology), and based on structural similarity (nodes with similar neighborhoods have similar ideologies). In our methods, we leverage the concept of ideological coherence of a neighborhood as a parameter for propagation. Using Twitter data, we produce an ideological scaling for 370K users, and analyze the two 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 [1], [2]. These methods exploit geometrical representations in a feature space. Multidimensional network scaling methods, for example, can be used to infer latent features that are determinant in the structure of networks. These new scaling methods have been especially successful in identifying latent features in social networks related to the ideology of users [3]. 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 [4]. This raises the question: How to compute features for portions of the network outside this *seed* sub-graph? In this setting, 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 is different from 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 proposes two methods for the propagation of scaled ideological features, from a seed set of nodes, to larger parts of 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 methods 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. We test these methods in an application case based on the ideology scaling on the Twitter network in France.

II. Related Work

This work concern mainly three domains of research.

1) Latent ideological features in social networks: This article is concerned by methods that leverage the interpretability of a subset of nodes to extract meaningful feature spaces (for example accounts of MPs). This stands in contrast with 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. [5]). A numerous family of methods learn latent features in social networks, but specifically in relation with the probability of existence of a link [6], [7], or the retrieval of missing features for some nodes [8]. The most relevant works for this article are those that use multidimensional scaling to extract ideological features [3], [9]. Other works seek to predict ideology for large networks but use ideological labels and not a –continuous– ideological feature space [10], [11].

2) Propagation in social networks: Propagation in network (e.g. rumors and misinformation [12]) is similar but

This work has been funded by the French National Agency for Research under grant ANR-19-CE38-0006: Geometry of Public Issues (GOPI).
IEEE/ACM ASONAM 2020, December 7-10, 2020

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 [13]. This connection is also related, for example, to types of users in rumor propagation cascades in social networks [14]. Propagation in social networks also has connections with label propagation graphs [15]. But this differs from continuous features spaces in that a label can be equated to a discrete classification.

3) Node coherence: Node coherence is a concept that can be encountered in domains such as neuroscience [16], artificial intelligence [17], or communications networks [18], and often addresses the degree of dissimilarity of signals coming from different sources. The most relevant domain of works along this line, is found in node coherence in opinion dynamics [19]. In this domain, the concept of opinion coherence is sometimes considered in dynamic models [20], where the goal is the simulation in artificial societies rather than inference in social network data.

III. Preliminaries

A. The universe network

Let us consider a large, directed *universe network* $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, for some non-empty set $\mathcal{V}$ (social network accounts) with directed edges $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ (relations of *following* between them). Information flows *downstream*, contrary to the direction of the edges.

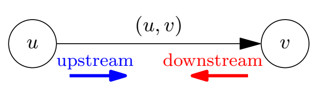

Fig. 1. Convention used for *upstream* and *downstream* directions on directed social networks, where edges indicate following/friendship.

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 *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 $V \cap n_d(V) = \varnothing$, for $d \in \{\text{up}, \text{down}\}$. Given a direction d, we will denote its opposite direction with a bar, as $\overline{d}$.

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 the known N-dimensional features by $e(v) \in \mathbb{R}^N$, and its i-th component by $e_i(v)$. When unavailable, N-dimensional features may be attributed through feature propagation.We denote estimated features for a node $v \in \mathcal{V}$ by $\hat{e}(v) \in \mathbb{R}^N$. If both, true and estimated features (through propagation methods) can be known, 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, denoted by $E(\hat{e}(v), e(v)) = \|\hat{e}(v) - e(v)\|_p$. For a set of nodes V, the mean estimation error is computed as $E(V) = (1/|V|) \sum_{v \in V} E(\hat{e}(v), e(v))$.

C. Node coherence

Given $V \subset \mathcal{V}$ for which features are known or estimated, we denote by $I(V)$ the incoherence of V:

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

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

Whenever we have a set of nodes $V \subset \mathcal{V}$, we can consider the up- or downstream neighborhoods of nodes that are also coherent. We formalize this notion, for both directions, as the *ε-coherent upstream* and *downstream neighborhood* of set V: $n_d^{\varepsilon}(V) = \{u \in n_d(V) : I(n_{\overline{d}}(u) \cap V) \leq \varepsilon\}$ for $d \in \{\text{up}, \text{down}\}$, and $\overline{d}$ the direction opposite to d.

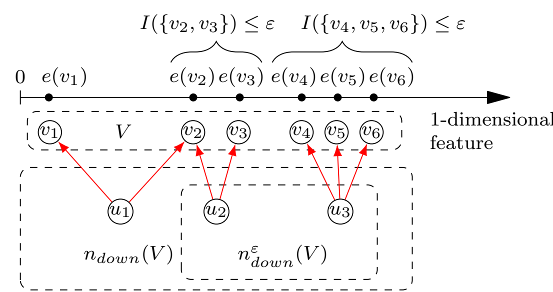

Fig. 2. Illustration of the construction of a ε-coherent downstream neighborhood $n_{\text{down}}^{\varepsilon}(V)$ of a set V embedded in a 1-dimensional feature space.

IV. Propagation Methods for Latent Features in Social Networks

We consider two approaches for propagating ideological features: Method A) directed sequences of ε-coherent neighborhoods, and Method B) sequences of projections using ε-coherent neighborhoods. Method A (rooted in the homophily hypothesis) assumes that the features attached to nodes are "diffusing" along the social network formed by follower/followee relationships. For Method B (rooted in the structural similarity hypothesis) ideological features of a node are estimated using those of other nodes that occupy a similar position in the network.

A. Directed sequences of ε-coherent neighborhoods

This method generates two sequences of sets of nodes: coherent nodes $\{V_i\}_{i \geq 0}$, for which we estimate the ideology, and incoherent nodes $\{\overline{V}_i\}_{i \geq 0}$, 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 $V_0, V_1, V_2, \ldots$ is computed as $V_{i+1} = V_i \cup \Delta V_i$

for $i = 0, 1, 2, \ldots$, with $\Delta V_i = \left\{ v \in n_{\bar{d}}^{\varepsilon}(V_i) : v \notin \left(V_i \cup \overline{V}_i \right) \right\}$ where $\overline{V}_{i+1} = \overline{V}_i \cup \Delta\overline{V}_i$ for $i = 0, 1, 2, \ldots$, with $\overline{V}_0 = \varnothing$ and $\Delta\overline{V}_i = \left\{ v \in n_d(V_i) : v \notin \left(\overline{V}_i \cup \Delta V_i \right) \right\}$. By definition, the sequence of ε-coherent neighbors $\{V_i\}_{i \geq 0}$ is incremental ($V_i \subset V_{i+1}$) by disjoint additions ($V_i \cap \Delta V_i = \varnothing$). Sequence $\{\overline{V}_i\}_{i \geq 0}$ is also increasing ($\overline{V}_i \subset \overline{V}_{i+1}$) by disjoint additions ($\overline{V}_i \cap \Delta\overline{V}_i = \varnothing$), but of incoherent neighbors. Whenever new ε-coherent neighbors ΔV_i are discovered at the i-th step, the features of their nodes are estimated as

$$\hat{e}(v) = \frac{1}{\left| n_{\bar{d}}(v) \cap V_i \right|} \sum_{u \in n_{\bar{d}}(v) \cap V_i} \hat{e}(u), \quad \text{for } v \in \Delta V_i, \tag{2}$$

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

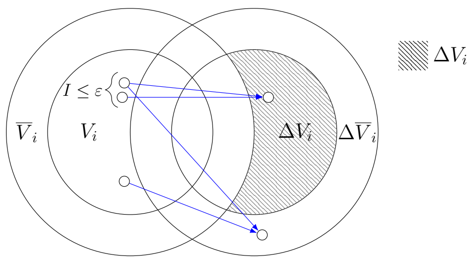

Fig. 3. Schematic representation of the computation of ΔV_i and $\Delta\overline{V}_i$ on the i-th step of an upstream directed sequence of ε-coherent neighborhoods using **Method A**. In the upstream variant, users from V_i follow users from ΔV_i.

B. Sequences of projections using ε-coherent neighborhoods

For Method B we compute sequences $\{V_i\}_{i \geq 0}$, for which we estimate the ideology, and $\{\overline{V}_i\}_{i \geq 0}$ that we avoid using, but now as *pivots*. Let us consider a seed set $V_0 \subset \mathcal{V}$ with known features. Starting at V_0, a sequence of projected ε-coherent sets of nodes $V_0, V_1, V_2, \ldots$ is computed as $V_{i+1} = V_i \cup \Delta V_i$ for $i = 0, 1, 2, \ldots$, for which we consider the set P_i^{ε} of ε-coherent pivots: $P_i^{\varepsilon} = n_d^{\varepsilon}(V_i) \backslash \overline{V}_i$, where $\overline{V}_{i+1} = \overline{V}_i \cup \Delta\overline{V}_i$ for $i = 0, 1, 2, \ldots$, with $\overline{V}_0 = \varnothing$, and $\Delta\overline{V}_i = \left\{ v \in n_d(V_i) : v \notin \left(n_d^{\varepsilon}(V_i) \cup \overline{V}_i \right) \right\}$.

As with the previous family of Method A, the sets $\overline{V}_i$ 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: $\Delta V_i = \left\{ v \in n_{\bar{d}}^{\varepsilon}(P_i^{\varepsilon}) : v \notin V_i \right\}$. Again by definition, sequences $\{V_i\}_{i \geq 0}$ and $\{\overline{V}_i\}_{i \geq 0}$ 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

$$\hat{e}(v) = \frac{1}{\left| C_i(v, P_i^{\varepsilon}, V_i) \right|} \sum_{u \in C_i(v, P_i^{\varepsilon}, V_i)} \hat{e}(u), \quad \text{for } v \in \Delta V_i, \tag{3}$$

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}:

$$C_i(v, P_i^{\varepsilon}, V_i) = V_i \cap n_{\bar{d}}\left(n_d(v) \cap P_i^{\varepsilon} \right). \tag{4}$$

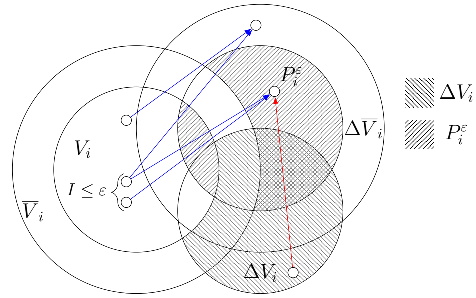

Fig. 4. Schematic representation of the computation of ΔV_i and $\Delta\overline{V}_i$ using pivot P_i^{ε} on the i-th step of an upstream (co-followers-based) sequence of projected ε-coherent neighborhoods using **Method B**. In the upstream variant, users from V_i and ΔV_i co-follow users from pivot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

V. French Political Twitter Dataset

To test our methods we present a pertinent part of the Twitter network [1] and its ideological scaling.

A. Multidimensional scaling for the French MPs and their followers

We consider the set $\mathcal{P}$ of the 831 (out of 925) French MPs present on Twitter[2] affiliated to 10 parties, and their followers. This collection was conducted on May 2019, and resulted in the set $\mathcal{F} = n_{\text{down}}(\mathcal{P})$, with $|\mathcal{F}| = 4.487.430$. After removing the followers that follow less than 3 MPs and users that had a repeated set of followed MPs we obtain the set $\tilde{\mathcal{F}}$ of 368.831 accounts by removing followers that follow less than 3 MPs, and users that had a repeated set of followed MPs. We represent this sub-graph as a $\{0, 1\}^{|\tilde{\mathcal{F}}| \times |\mathcal{P}|}$ adjacency matrix. Next, we produce a reduced-dimensionality representation using a Correspondence Analysis (CA) [21], and keeping the first 2 principal components: PC1 and PC2.

To provide and interpretation of PC1 and PC2 for our sets $\mathcal{P}$ and $\tilde{\mathcal{F}}$, we use the 2019 Chapel Hill Expert Survey (CHES) data [22]. Out of the 10 political parties identified for accounts in $\mathcal{P}$, 8 are also present in the CHES data. We identify the two most relevant CHES criteria related to our axes examining the correlations of the positions of MPs on PC1 and PC2 of the feature space with all the 51 criteria were CHES data. 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.

In this feature space, PC1 provides an index for the concept of *left* and *right* positions, and PC2 provides an index for attitudes towards European integration. Fig. 6 illustrates the positions of the sets $\mathcal{P}$ of parliamentarians and $\tilde{\mathcal{F}}$ of their followers in this bi-dimensional ideological feature space.

[1] These data have been declared, the 19 Mars 2020 at the registry of data processing at the *Fondation Nationale de Sciences Politiques* (Sciences Po), and respects the General Data Protection Regulation 2016/679 (GDPR) and Twitter's policies.

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

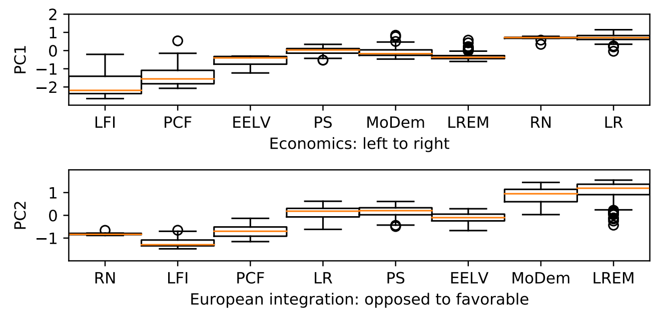

Fig. 5. Distribution of the embedded ideological features PC1 and PC2 for the French MPs on Twitter per party. Parties are ordered according to the two criteria using the 2019 Chapel Hill Expert Survey data: left-right economics for feature PC1 (top), and attitudes toward European integration for feature PC2 (bottom).

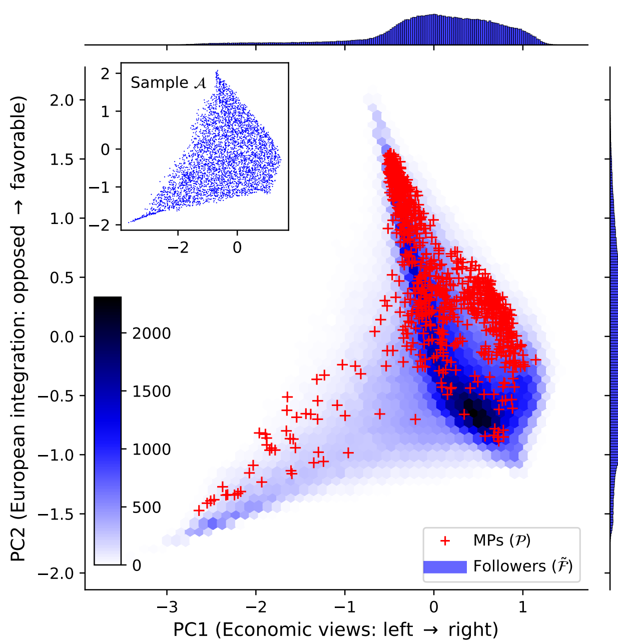

Fig. 6. Position of $\mathcal{P}$ French MPs in the first two embedded dimensions, PC1 (left to right ideology) and PC2 (pro- and against European integration), histogram of the positions of their followers $\tilde{\mathcal{F}}$, and positions of sampled accounts in set $\mathcal{A}$.

B. Datasets for experimentation

We are now concerned with the problem of using the set $\tilde{\mathcal{F}}$ to establish datasets for the testing of methods A and B. We cannot set $V_0 = \tilde{\mathcal{F}}$ for testing and evaluation, because the the nodes of the next set V_1 in the sequence would not have known *true* features $e(v)$. 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[3] nodes uniformly in space from $\tilde{\mathcal{F}}$ to produce set $\mathcal{A}$. 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 $|\mathcal{C} \cap \tilde{\mathcal{F}}| = 231.035$.

VI. Numerical Experiments

We center the analysis of both methods around the set $\mathcal{A}$. We will be interested in analyzing, for different values of coherence ε, 1) the accuracy 2) the coverage of the proposed methods, and 3) the trade-off between them.

A. Directed sequences of ε-coherent neighborhoods

Using upstream and downstream neighborhoods of set $\mathcal{A}$ (sets $\mathcal{B}$ and $\mathcal{C}$), we 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 ($i = 0$).

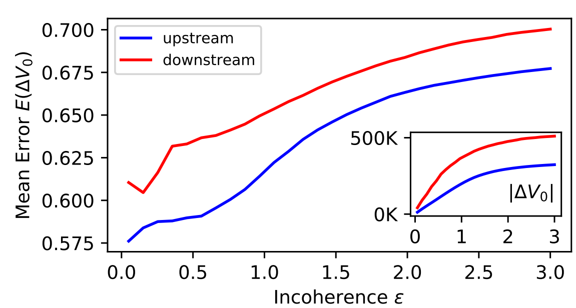

Fig. 7. Mean estimation error $E(\Delta V_0)$ and size of ΔV_0 for Method A in the upstream and downstream directions, using $V_0 = \mathcal{A}$.

While intuitive, Method A performs poorly in both directions. Imposing higher levels of coherence does improve the mean estimation error of nodes in ΔV_0. This improvement is marginal and, most importantly, limited. For example, upstream neighbors (friends) of users in $\mathcal{A}$ that are followed by highly coherent users, have mean estimated no lower than $E(\Delta V_0)$ =0.6 (12.5% of the diameter of $\tilde{\mathcal{F}}$).

B. Sequences of projections using ε-coherent neighborhoods

The setting of Method B is similar to that of collaborative filtering (CF) in Recommender Systems: the similarity of users is computed according to structural similarity, where users are similar if they have chosen similar items. In the setting of our Method B, users are deemed similar if they follow or are followed by similar users of a so-called pivot set for a predetermined coherence ε. Taking on known evaluation protocols for CF [23], we propose a method for assessing the accuracy and coverage of Method B. We perform a K-fold bipartite cross-validation. We divide $\mathcal{A}$ in K parts, taking one as $\mathcal{A}^K_{\text{test}}$ and the rest as $\mathcal{A}^K_{\text{train}}$. We set $V_0 = \mathcal{A}^K_{\text{train}}$, we compute

[3] Originally, we sampled 5.000, but collection of 517 of these nodes was not possible because accounts were *private*.

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.

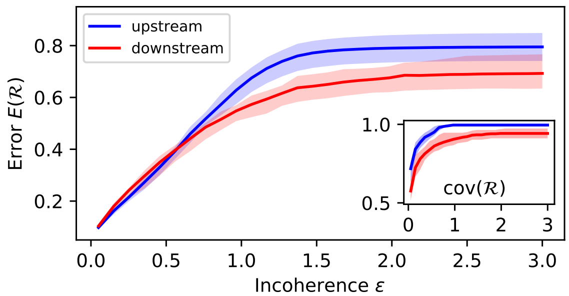

Fig. 8. Estimation error and coverage of set $\mathcal{R} = \mathcal{A}_{\text{test}}^K \cap \Delta V_0$ and P_0 for the Method B in a K-fold cross-validation retrieval of ideologies of $\mathcal{A}$ (K =20). The distribution of errors and coverage over the folds is reported with the median, the minimum,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

Results obtained with Method B are more satisfactory than those obtained with Method A. Thresholds for incoherence below $\varepsilon = 1$ (approximately) already allow for less error in comparison to Method A. In fact, under our setting, estimation errors for Method B can be made as low as 0.1 (approximately) in the feature space by sacrificing coverage. With a small value of threshold for ϵ (lower than 0.1) 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.

VII. Conclusions

We proposed two methods for the propagation of latent ideological features on Twitter: Method A (based on the homophily), and Method B (based on the structural similarity). 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 CHES data, we validated the interpretation of the two emerging dimensions. Collecting neighborhoods of some of these users we used our methods to propagate these ideological features within a potential population of nearly 6.5M users. We analyzed the relation between estimation accuracy for ideology in propagation, and the coverage of the method. 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 users that might be directly connected (Method A). These results lend support to the structural similarity hypothesis over the homophily hypothesis in the case of Twitter. Using these methods, we were able to propagate ideological features from small parts of a graph to larger portions while accounting for the estimation error.

References

[1] P. Goyal and E. Ferrara, “Graph embedding techniques, applications, and performance: A survey,” *Knowledge-Based Systems*, vol. 151.
[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.
[3] P. Barberá, “Birds of the same feather tweet together: Bayesian ideal point estimation using twitter data,” *Political analysis*, vol. 23, no. 1.
[4] P. Barberá and G. Rivero, “Understanding the political representativeness of twitter users,” *Social Science Computer Review*, vol. 33, no. 6.
[5] 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.
[6] C. Heaukulani and Z. Ghahramani, “Dynamic probabilistic models for latent feature propagation in social networks,” in *International Conference on Machine Learning*.
[7] 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*.
[8] 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*.
[9] 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.
[10] I. Himelboim, 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.
[11] 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.
[12] 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*.
[13] 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.
[14] A. Friggeri, L. Adamic, D. Eckles, and J. Cheng, “Rumor cascades,” in *Eighth International AAAI Conference on Weblogs and Social Media*, 2014.
[15] 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.
[16] P. Thagard and K. Verbeurgt, “Coherence as constraint satisfaction,” *Cognitive Science*, vol. 22, no. 1.
[17] A. K. Mackworth, “Consistency in networks of relations,” *Artificial intelligence*, vol. 8, no. 1.
[18] P. Larsson and H. Rong, “Large-scale cooperative relaying network with optimal coherent combining under aggregate relay power constraints,” in *WWRF 2004*, 2004.
[19] B. Edmonds, “Modeling belief change in a population using explanatory coherence,” *Advances in Complex Systems*, vol. 15, no. 06.
[20] 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.
[21] M. Greenacre, *Correspondence analysis in practice*, 2017.
[22] 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*.
[23] 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.