# Community detection in bipartite signed networks is highly dependent on parameter choice

Elena Candellone<sup>a,b</sup>, Erik-Jan van Kesteren<sup>a</sup>, Sofia Chelmi<sup>a</sup>, and Javier Garcia-Bernardo<sup>a,b</sup>

<sup>a</sup>Department of Methodology & Statistics, Utrecht University, Utrecht, The Netherlands <sup>b</sup>Centre for Complex Systems Studies, Utrecht University, Utrecht, The Netherlands

# ABSTRACT

Decision-making processes often involve voting. Human interactions with exogenous entities such as legislations or products can be effectively modeled as two-mode (bipartite) signed networks—where people can either vote positively, negatively, or abstain from voting on the entities. Detecting communities in such networks could help us understand underlying properties: for example ideological camps or consumer preferences. While community detection is an established practice separately for bipartite and signed networks, it remains largely unexplored in the case of bipartite signed networks. In this paper, we systematically evaluate the efficacy of community detection methods on bipartite signed networks using a synthetic benchmark and real-world datasets. Our findings reveal that when no communities are present in the data, these methods often recover spurious communities. When communities are present, the algorithms exhibit promising performance, although their performance is highly susceptible to parameter choice. This indicates that researchers using community detection methods in the context of bipartite signed networks should not take the communities found at face value: it is essential to assess the robustness of parameter choices or perform domain-specific external validation.

### Introduction

Social groups often make decisions by voting. From politicians passing legislation, to online marketplaces where users review products, to social media users voting on popular content, these interactions reveal complex patterns, such as political affiliations, consumer preferences, or ideological camps. In such systems, individuals do not interact directly with each other, but with different entities (e.g., bills, products, articles) by voting for, against, or abstaining. These interactions can be modeled as bipartite signed networks, where two different types of nodes (bipartite) are connected through both positive and negative relationships (signed). A key insight that can be recovered from those networks is the presence of communities<sup>1</sup>—groups of people with similar voting patterns. These indicate alliances between politicians, hidden groups trying to promote products or articles, or ideological groups in the presence of polarization.<sup>1</sup> While many methods exist to analyze community structures in bipartite and signed networks separately<sup>2–6</sup>, it is not clear whether they can handle the complexity of bipartite signed networks.

Community detection, a key research topic in network science<sup>1</sup>, often aims to uncover structural patterns by identifying groups with dense intra-community connections and sparse inter-community connections. In the case of signed networks, it is not enough to have dense intra-community connections, but these connections should be positive, while negative connections should occur mainly between communities. Existing methods<sup>7–9</sup> primarily focus on these principles of intra-community agreement and inter-community disagreement. These methods incorporate the information of negative relations in areas where it is crucial, such as finding communities in a network of friendships and enmities, to networks of alliances and wars<sup>10–12</sup>. However, the structure of agreement and disagreement in bipartite networks presents a unique challenge: interactions occur between different sets of entities, while most methods assume a unique type.

Here, we investigate the suitability of community detection methods to capture different groups on bipartite signed networks. To address this question, we developed a benchmark (pictured in Figure 1) to evaluate community detection methods devised for signed networks. Our approach involves generating synthetic bipartite signed networks, where users (who may be in one or two communities in different scenarios) vote on articles. We then apply two well-established algorithms and systematically evaluate their ability to recover user-level communities. We apply our methodology on two datasets: voting records from the US House of Representatives (1990-2022)<sup>13</sup> and user interactions on the Menéame platform<sup>14</sup>, a Reddit-like news aggregator.

Our findings reveal that community detection methods often identify spurious communities not anticipated by the generative process, and the results are highly sensitive to parameter choices.

This paper is organized as follows. Section Background contextualizes our contribution within the existing literature. The

<sup>&</sup>lt;sup>1</sup>Instead of finding clusters of people with similar voting patterns, community detection methods could also be applied to find groups of entities (e.g., legislation, products or online content) that are voted similarly. In this paper, we focus only on clustering people.

Methods section is subdivided into a description of the mathematical framework of bipartite signed networks (subsection bipartite signed Networks: Mathematical Framework), a detailed explanation of our benchmark (subsection Generating synthetic bipartite signed networks), and a contextualization of the community detection methods tested in this paper (subsection Community Detection Methods) along with the clustering evaluation metrics (subsection Evaluating Community Similarity). The Results section demonstrates the effectiveness of the community detection methods on the synthetic and real-world datasets employed.



**Figure 1.** Synthetic bipartite signed networks. We combine synthetic scenarios and insights from data to generate synthetic networks. Given a synthetic scenario, i.e., ideologies of users and stories modeled from probability distributions  $p_U$  and  $p_S$ , which can be either unimodal or bimodal (step 1), we sample users' and stories' ideologies,  $x_U$  and  $x_S$ , from those distributions (step 2). For each pair of user and story, the user then either votes positive, negative, or abstained from voting depending on the difference between the user and story ideologies (step 3), where the vote depends on two voting thresholds,  $t_+$  and  $t_-$ , set to match the voting probability in real datasets. Given the bipartite network, we then project it into a unipartite network and apply the community detection methods for a wide range of parameter choices (step 4).

# Background

Community detection is an extensively researched problem in the field of network science<sup>1</sup>. When dealing with unsigned networks, most techniques for detecting communities aim to maximize edge density within a community while minimizing it between communities. This interdisciplinary goal finds applications in various fields, such as biology (e.g., protein-protein interaction networks in cancer metastasis<sup>15</sup>), social media studies (e.g., political blogs' ideological affiliations<sup>16</sup>), and the science of science (e.g., scientists' collaboration network<sup>17</sup>). However, signed networks provide additional information—the sign of the edge—that can be relevant for determining and interpreting communities in the network<sup>6, 18, 19</sup>.

Research on signed networks is closely tied to the influential work of Heider<sup>20</sup> on Structural Balance Theory (SBT), later formalized by Harary and Cartwright<sup>21</sup>. SBT is based on the principles of "the enemy of my enemy is my friend" and "the friend of my friend is my friend". In network-theoretical terms, SBT posits that triangles with zero or two negative edges are balanced, while triangles with one negative edge are unbalanced. This concept extends to cycles (closed paths) of length  $\geq$  3, stating that a cycle is balanced if the product of the edge signs is positive (i.e., there is an even number of negative interactions in a cycle). A network is considered strongly balanced if all cycles are balanced<sup>21</sup>.

SBT is regarded as one of the most relevant underlying mechanisms for the formation and evolution of signed networks, with various extensions that account for longer closed walks<sup>18,22–24</sup> used to understand situations like alliances among countries<sup>25,26</sup>

and polarization on social media<sup>27</sup>. Methods for community detection in signed networks typically do not maximize structural balance directly but instead minimize a related metric, *frustration*. Aref and Wilson<sup>28</sup> defined the *frustration count* as the sum of the number of negative edges within a community and the number of positive edges between communities<sup>28</sup>. As expressed by Doreian and Mrvar<sup>9</sup>, there is a parallelism between structural balance and frustration minimization. Two theorems by Cartwright and Harary<sup>21</sup> and by Davis<sup>29</sup> state that it is possible to find optimal partitions in a balanced network and that strong balance corresponds to null frustration (i.e., no frustrated edges).

While SBT and frustration provide a way to define communities in signed networks, applying them to bipartite signed networks has a key challenge: SBT and frustration are deeply linked to the presence of triangles, while bipartite networks have no cycles of odd length<sup>30</sup>. Bipartite networks exhibit disassortative properties, with two sets of nodes forming edges only across different sets. Often, the two sets of nodes have different types, such as movies and actors<sup>31</sup> or buyers and products on an online marketplace<sup>32</sup>. A potential solution is to project the networks: create a one-mode network where two nodes interact if they are connected to the same entity in the bipartite network. For example, this involves connecting politicians who voted on the same bill.

A body of literature within political science has traditionally dealt with a specific type of signed bipartite networks: co-voting networks of the United States Senate and House of Representatives<sup>33,34</sup>. The goal of that body of literature is to estimate legislators' ideology, for which they use ideal point models, where the probability of voting positive versus voting negative or abstaining is modeled as a function of a latent difference in ideology between the legislator and the bill<sup>35–39</sup>. The US co-voting networks have also been used to measure ideology<sup>40,41</sup>, calculate community structure<sup>42,43</sup>, points in which voting dynamics change<sup>44</sup> or to predict the sign of the votes<sup>32,45</sup>. However, given the dense nature of voting interactions (i.e., most legislators would vote to all the roll calls), these methods consider only a binary choice between agreement and disagreement—i.e., disagreement encapsulates both negative and absence of interactions. Our work complements this body of literature by extending the study of voting patterns to sparse (offline) voting records, allowing us to understand how network density affects community detection.

### Methods

#### **Bipartite Signed Networks: Mathematical Framework**

Consider a network G = (V, E), where V is the set of nodes of size N, and E is the set of edges of size M. In a bipartite network, nodes are divided into two disjoint sets, which we label U and S to represent users and stories (or legislators and bills)<sup>2</sup>. The edges connect nodes from one subset to the other only, i.e.,  $E \subseteq U \times S$ . A network can be uniquely represented by an adjacency matrix A of dimension  $N \times N$ , with each element given by the edge weight such that  $A_{i,j} = \omega_{i,j}$  if the edge  $(i, j) \in E$ , else  $A_{i,j} = 0$ . The edge weight, in the case of signed networks, can be either binary (negative or positive,  $\omega : E \to [-1,1]$ ) or real  $(\omega : E \to \mathbb{R})$ . In this paper, we consider the edge weights as simplified to the binary choice. As a special property of undirected bipartite networks, the adjacency matrix can be written as

$$A = \begin{pmatrix} 0 & B \\ B^T & 0 \end{pmatrix},$$

by ordering the nodes such that  $\{v_1, \ldots, v_{|U|}\} \in U$  and  $\{v_{|U|+1}, \ldots, v_N\} \in S$ . The matrix *B* is called the incidence matrix (or biadjacency matrix) and has dimension  $|U| \times |S|$ . The elements of *B* represent users' (in rows) votes on stories (columns).

From the incidence matrix, it is possible to obtain the bipartite projection of the network, denoted as  $P = B \cdot B^T$ . In the bipartite projection, nodes correspond to those from set U, and edges represent the presence of a common connection to nodes in set S. In the case of a weighted bipartite network, the edge weight in the bipartite projection is the sum of the edge weights

$$\omega_{kl} = \sum_{z \text{ s.t. } (k,z) \text{ and } (l,z) \in E} \omega_{k,z} \, \omega_{z,l}.$$
(1)

To provide intuition, if two politicians always agree on their votes, the edge between them has a positive weight, representing agreement. Conversely, if they always vote differently, their edge has a negative weight, representing disagreement. We project the bipartite signed network to analyze the performance of the community detection methods designed for unipartite networks, where all nodes are of the same type (e.g., we aim to find communities of users on social media or members of the US Congress). We use the adjacency matrix for the methods that can accommodate bipartite structures. In the following, we provide a detailed description of the generative mechanism to create synthetic bipartite signed networks.

 $<sup>^{2}</sup>V = U \cup S$  and  $U \cap S = \emptyset$ .

#### Generating synthetic bipartite signed networks

To assess the performance of current community detection methods, we generate synthetic networks with and without a community structure, with realistic voting mechanisms. We consider four scenarios that entail different choices for the *user* and *story ideologies*. In two scenarios, the ideology of the users is drawn from a Gaussian distribution. Therefore, we expect the community detection algorithms to find one community. In the other two scenarios, the ideology of the users is polarized (drawn from a mixture of two Gaussian distributions), and we expect the community detection algorithms to find two communities. The generative process is illustrated in Figure 1 and the following paragraphs.

1. **Sampling ideologies.** We assign a latent *ideology* score to each of |U| *users* and |S| *stories*, representing the two types of nodes in the bipartite signed network. While we focus our example on users voting on news stories for simplicity, this latent score could represent the ideology of politicians and bills, or the type of consumer preferences and product characteristics (e.g. gelato vs ice-cream). For each type of entity (*U* and *S*) we sample the ideology score from either one Gaussian distribution or a mixture of two Gaussian distributions with given averages ( $\mu_i$ ) and standard deviations ( $\sigma_i$ ). Specifically,

$$x_{U} \sim p_{U}(x_{U}) = \pi_{U} \cdot \mathcal{N}(\mu_{1}, \sigma_{1}^{2}) + (1 - \pi_{U}) \cdot \mathcal{N}(\mu_{2}, \sigma_{2}^{2}),$$
  

$$x_{S} \sim p_{S}(x_{S}) = \pi_{S} \cdot \mathcal{N}(\mu_{3}, \sigma_{3}^{2}) + (1 - \pi_{S}) \cdot \mathcal{N}(\mu_{4}, \sigma_{4}^{2}),$$
(2)

where  $p_U(x_U)$  and  $p_S(x_S)$  are respectively the probability distributions of *users' and stories' ideologies*. When the ideology is sampled from *one* Gaussian distribution,  $\pi_U = \pi_S = 1$ . When is sampled from a mixture of two, we set  $\pi_U = \pi_S = 0.5$ , which means that we assume that there is no imbalance in the sizes of two groups of *ideologies*. We set  $\sigma_1 = \sigma_2 = \sigma_3 = \sigma_4 = 0.1$  and  $\mu$  depending on scenario's choice. The four scenarios, illustrated in Figure 1A, are:

- Users and stories not polarized (U NP S NP): we consider both users and stories as not polarized. In other words, both *ideologies* are drawn from one Gaussian-distribution ( $\mu_1 = \mu_2 = \mu_3 = \mu_4 = 0$  in eq. 2). We expect to find all users in the same community.
- Users not polarized, stories polarized (U NP S P): Users are not polarized, and stories are polarized. In other words, users' ideologies are drawn from one Gaussian distribution ( $\mu_1 = \mu_2 = 0$ ) and stories' ideologies are drawn from a mixture of two Gaussian distributions ( $\mu_3 = -\mu_4 = 1$ ). We expect to find all users in one community.
- Users polarized, stories not polarized (U P S NP): Users are polarized, and stories are not polarized. In other words,  $\mu_1 = -\mu_2 = 1$  and  $\mu_3 = \mu_4 = 0$ . We expect to find all users in two communities.
- Users and stories polarized (U P S P): We consider both users and stories as polarized. In other words,  $\mu_1 = -\mu_2 = 1$  and  $\mu_3 = -\mu_4 = 1$ . We expect to find all users in two communities.

After this step, every node has an *ideology* assigned. Users will vote positive to a story if the difference of their ideologies is below a threshold, negative if it is above a threshold, and abstain otherwise.

2. **Computing the ideology difference distribution.** If two random variables are Gaussian-distributed, then their linear combination is also Gaussian-distributed. For a detailed proof, please refer to the Appendix Ideology Difference. The *ideology difference* between users and stories is distributed as a Gaussian Mixture with at most four peaks, e.g., in the case when all averages are different:

$$x_{U} - x_{S} \sim p_{\text{diff}}(x_{U} - x_{S}) = \pi_{U} \cdot \pi_{S} \cdot \mathscr{N}(\mu_{1} - \mu_{3}, \sigma_{1}^{2} + \sigma_{3}^{2}) + \pi_{U} \cdot (1 - \pi_{S}) \cdot \mathscr{N}(\mu_{1} - \mu_{4}, \sigma_{1}^{2} + \sigma_{4}^{2}) + (1 - \pi_{U}) \cdot \pi_{S} \cdot \mathscr{N}(\mu_{2} - \mu_{3}, \sigma_{2}^{2} + \sigma_{3}^{2}) + (1 - \pi_{U}) \cdot (1 - \pi_{S}) \cdot \mathscr{N}(\mu_{2} - \mu_{4}, \sigma_{2}^{2} + \sigma_{4}^{2}).$$
(3)

3. Extracting voting probabilities from data. The thresholds determining the type of vote (positive, abstain, or negative) are calculated so the voting probability in the simulation matches the voting probability in the real data (Section Datasets). In the simplest case, we define constant voting probabilities extracted from data. Voting probabilities  $v_+$  ( $v_-$ ) are computed as the average number of positive (negative) votes per user, normalized by the number of stories. We also consider a more sophisticated alternative (Appendix Degree-corrected synthetic networks), where we assign different voting probabilities to each user, given the distribution of voting probabilities in the real data's bipartite network.

4. Defining voting thresholds. We determine voting thresholds, denoted as  $t_+$  and  $t_-$ , that meet the following conditions:

$$\int_{-t_{+}}^{t_{+}} p_{\text{diff}}(x) \, dx = v_{+} \quad \text{and} \quad \int_{t_{-}}^{+\infty} p_{\text{diff}}(x) \, dx = \frac{v_{-}}{2}.$$
(4)

where  $p_{\text{diff}}$  is the ideology difference distribution given in Equation 3,  $v_+$  and  $v_-$  are the voting probabilities obtained from data. In this way, we have an amount of positive and negative votes proportional to what is expected from real data.

5. Creating the bipartite signed network. Finally, for each *user* and *story* with *ideologies*  $x_U$  and  $x_S$ , the vote of user U to story S is determined by

$$v(x_U, x_S) = \begin{cases} +1 & \text{if } |x_U - x_S| < t_+, \\ -1 & \text{if } |x_U - x_S| > t_-, \\ 0 & \text{otherwise.} \end{cases}$$
(5)

After this step, we obtain a bipartite signed network, where the voting patterns respect both controlled, synthetic settings, and real data properties.

Finally, given the incidence matrix of the synthetic network, we obtain the bipartite projection where the nodes are the *users* and the edge weight is given by Equation 1.

The generative approach allows us to generate synthetic bipartite signed networks with voting probabilities matching real datasets, but where we modify the number of clusters expected (either one or two).

#### **Community Detection Methods**

We compare two widely used community detection methods: SPONGE and community-spinglass. Both methods infer communities by minimizing network frustration—i.e., assigning nodes into communities in a way that minimizes the number of positive edges and maximizes the number of negative edges between communities. Since this problem is NP-hard<sup>28</sup>; different approximations have been proposed. As an additional test, we also analyzed an algorithm that does not necessarily follow that principle—the Stochastic Block Model (SBM). For more details and results on the SBM, please refer to Appendix Weighted SBM with Edge Attributes).

#### SPONGE

The *SPONGE* method (Signed Positive Over Negative Generalized Eigenproblem), introduced by Cucuringu et al.<sup>7</sup> minimizes the number of "violations", which is equivalent to frustration minimization. These violations consist of positive edges between communities and negative edges within communities. The method achieves this objective by incorporating a regularization term for the size of the partitions. This regularization term avoids setting each node in its community.

The optimization problem is formulated as follows:

$$\min_{C_1,\dots,C_k} \sum_{i=1}^k \frac{x_{C_i}^T (L^+ + \tau^- D^-) x_{C_i}}{x_{C_i}^T (L^- + \tau^+ D^+) x_{C_i}},\tag{6}$$

where  $\{C_1, \ldots, C_k\}$  represents the communities,  $L^+$  and  $L^-$  are the Laplacian matrices of the network with only positive and negative edges, respectively, and  $D^+$  and  $D^-$  are diagonal matrices.  $x_{C_i}$  with  $C_i \in \{C_1, \ldots, C_k\}$  are vectors where  $(x_{C_i})_j = 1$  if node  $j \in C_i$ , and 0 otherwise. The latter terms  $(\tau D)$  correspond to the previously mentioned regularization for the partition size, while the Laplacian acts as a *frustration count*. This optimization problem is equivalent to a spectral problem, as discussed in detail in the original paper<sup>7</sup>. An open-source Python implementation of the *SPONGE* algorithm is available on GitHub<sup>46</sup>.

The number of clusters k is a required parameter, while default values for the regularization parameters are set to  $\tau^+ = \tau^- = 1$ . We conducted tests for different values of k, ranging from k = 1 (no communities) to k = 10.

#### community-spinglass

The *community-spinglass* method, initially proposed by Reichardt and Bornholdt for unsigned networks<sup>47</sup>, presents a community detection approach grounded in statistical mechanics. Traag and Bruggeman later extended this method to signed networks<sup>8</sup>.

The method is based on finding a Hamiltonian's ground state (i.e., minimal energy), defined as a function of the partition of nodes into communities  $\{\sigma\}$ . The final form of the Hamiltonian, detailed in the original paper by Traag and Bruggeman<sup>8</sup>, is given by:

$$H(\{\sigma\}) = -\sum_{ij} \left[ A_{ij} - \left(\gamma^+ p_{ij}^+ - \gamma^- p_{ij}^-\right) \right] \delta\left(\sigma_i, \sigma_j\right),\tag{7}$$

Here, *A* is the adjacency matrix (could be both the adjacency matrix previously defined as *A*, or, as in our case, the bipartite projection matrix *P*), and  $\gamma^+$  (or  $\gamma^-$ ) is a parameter rewarding the presence (or absence) of positive (or negative) links within a community. The term  $\delta(\sigma_i, \sigma_j)$  equals 1 when nodes *i* and *j* are part of the same community ( $\sigma_i == \sigma_j$ ), and thus counts only interactions found in the same community. If  $\gamma = 1$ , present positive and missing negative edges are given equal importance. The case  $\gamma^+ = \gamma^- = 0$  corresponds to frustration minimization. The parameters  $p_{ij}$  represent the expected value for the link between two nodes in a random null network. There is an open-source implementation of this algorithm in the Python version of the *igraph* package<sup>48</sup>. To determine the sensitivity of the results to the  $\gamma$  parameters (the one that recovers the real community structure in our synthetic scenarios), we conducted tests with combinations of  $\{\gamma^+, \gamma^-\} \in [0.5, 1, 2]$ .

#### **Evaluating Community Similarity**

To assess how well the different community detection methods capture the community structure, we use the *Rand Index*<sup>49</sup> implemented in the Python *scikit-learn* package<sup>50</sup>. Since we draw the user ideology from a mixture of Gaussian distributions, we know which users belong to the same community. Given this true community labeling and the estimated labeling using the community detection algorithms, the Rand Index is defined as:

$$\mathrm{RI} = \frac{\mathrm{TP} + \mathrm{TN}}{\binom{n}{2}},\tag{8}$$

where TP (TN) is the number of pairs of elements found in the same (different) community in the true and estimated community labeling. The denominator accounts for all possible combinations of pairs between the nodes. The Rand Index is bounded between 0 (indicating that every pair of users is clustered in the wrong community) and 1 (indicating a perfect matching). We focus mostly on the Rand Index since it produces interpretable estimates even when only one community is present in the data. We also consider other evaluation metrics, but they fail to account for situations where only one community is expected (see Appendix Clustering evaluation metrics).

#### Datasets

In addition to evaluating the community detection methods on simulated network data, we apply and compare them on two real-world datasets: the Menéame votes and the US House of Representatives votes. These examples represent two different interaction modalities: sparse interactions, typical of the online world, and dense interactions, typical of the offline world.

#### Sparse voting: Menéame social platform

We gathered data from the Spanish social media platform known as Menéame<sup>14</sup>. In this paper, each post is referred to as a *story* and includes a hyperlink to a news article and a summary provided by the user sharing the story. Users can engage with the story by liking, or disliking them. For each story, positive votes were encoded as +1 and negative votes as -1. We constructed an incidence matrix  $B_M$  where the rows represent users, and the columns represent the stories on the homepage. We collected data from 11,628 users and 48,561 stories, covering the period from the 27th of November 2022 to the 17th of July 2023. The bipartite projection network comprises 11,628 nodes and 9,834,491 edges. The voting probabilities are  $v_+^M = 1.30741\%$  and  $v_-^M = 0.15602\%$  respectively for liking and disliking. We use these probabilities to create synthetic networks (Section Generating synthetic bipartite signed networks) with 1,000 users and 4,176 stories. The number of stories is set to match the proportion of stories/users found in the data.

#### Dense voting: House of Representatives

We obtained data from the Clerk of the House of Representatives website<sup>13</sup>. For each bill, we retrieved the following information: the name of the Representative, their party, and their vote. Each Representative could cast a vote of *Aye* (previously *Yea*), which was encoded as a positive vote with a weight of +1, *Nay* (previously *No*), encoded as a negative vote with a weight of -1, or be categorized as *Present* or *Not Voting*, both encoded as a null vote with a weight of 0. From this, we created an incidence matrix  $B_H$  of the bipartite network, where the rows are the Representatives and the columns are the Bills. We collected data from 1990 to 2022, encompassing votes from 1801 members of the House of Representatives presents basic descriptive statistics for each snapshot network. The average number of nodes per snapshot is N = 441 with a standard deviation of SD = 4. The average number of edges is M = 96726 with a standard deviation of SD = 1653. We computed voting probabilities for each snapshot, and the average positive voting probability is  $v_{+}^{H} = 60.09\%$  with standard deviation SD = 4.82%, while the average negative voting probability  $v_{-}^{H} = 33.14\%$  with standard deviation SD = 5.14%.

For both datasets, we then obtained the (dis)agreement unipartite network by multiplying the incidence matrix by its transpose, resulting in  $P_H = B B^T$  (as defined in Subsection Bipartite Signed Networks: Mathematical Framework). This projection is used for the SPONGE and community-spinglass algorithms.

# **Results and Discussion**

### Sparse (social) networks

#### Evaluation on synthetic networks

We first tested the performance of the algorithms on a sparse network, similar to those observed in online social networks. Figure 2 shows the results using our preferred measure for performance (the Rand Index), while Figure A6 in the Appendix shows that the results hold for other similarity measures. When analyzing the scenarios in which the users' ideology was drawn from a mixture of two Gaussian distributions (and thus two communities are expected), we found that *community-spinglass* and *SPONGE* had a Rand index over 75%—i.e., more than 75% of user pairs are clustered in the right community (Figure 2A,C). The constant performance of *SPONGE* indicates that when enforcing a higher number of communities than expected, the algorithm identifies those communities by adding a few users to the remaining communities. *community-spinglass* performed better when the negative votes within a community were penalized more (i.e. when  $\gamma^-$  is low or medium)—i.e., with parameters matching the principle of frustration minimization.



**Figure 2.** Community Detection on Sparse Synthetic Networks. We tested the community detection methods on the four synthetic scenarios. Panels (A,C) show the results for the scenarios with two communities of users, while panels (B,D) show the results with one community. We used the Rand Index to evaluate the performance of the algorithms with different parameter choices. Higher values of the Rand Index indicate better alignment between expected and empirical communities. Panels (A-B) show the results for *community-spinglass*. We experimented with different combinations of the parameters  $\{\gamma^+, \gamma^-\} \in [0.5, 1, 2]$ . Lower values of  $\gamma^+$  indicate less importance given to positive ties in a community, whereas lower values of  $\gamma^-$  penalize the presence of negative links within a community. Note that  $\gamma^- = 0.5$  is generally the best parameter choice, as it finds the expected communities for the synthetic scenarios. Panels (C–D) show the results for *SPONGE*. We conducted tests for different values of the number of clusters k, ranging from k = 1 (no communities) to k = 10. Note that the algorithm correctly identifies the expected communities in scenarios with polarized users (panel C), while it generates spurious communities in cases where stories introduce a latent ideology to a crowd of neutral users (the U NP S P case in panel D).

The results were more diverse for the scenarios in which the users' ideology was drawn from one Gaussian distribution (and thus one community is expected, Figure 2B,D). Similarly to the previous case, the *community-spinglass* method correctly recognizes all expected scenarios when negative votes within a community are penalized heavily ( $\gamma^- = 0.5$ ), while the results were independent of the weight of positive links within a community ( $\gamma^+$ ). *SPONGE* performs relatively well when neither the users nor the stories are polarized—Rand index over 70% even for a high number of clusters, but not when the stories are polarized. In the latter case, increasing the fixed number of clusters drastically reduces the Rand Index value, likely due to the latent ideology of the stories that split the group of users into those that are closer to one type of story and those that are closer to the other (even when the users are not very different from each other).

Finally, we allowed different users to have different voting probabilities, finding similar results (see Figure A5 in Appendix Degree-corrected synthetic networks) except for the scenario where neither users nor stories are polarized. In this scenario, SPONGE recovers the correct results (one community of users), while *community-spinglass* is only able to recover the correct results when both negative votes within a community, and positive links within a community are weighted more ( $\gamma^- = 0.5$ ,  $\gamma^+ = 2$ ).

#### Evaluation on real networks: Menéame data

We applied the community detection methods on the Meneame co-voting network. We find that both *community-spinglass* and *SPONGE* find two or three large communities (depending on the parameter choice) and several small communities.

Table 1 presents the results for *community-spinglass*. We restricted  $\gamma^+$  and  $\gamma^-$  to [0.5, 1], the choices giving best results on synthetic networks. The algorithm identifies one larger community, followed by another of less than 1000 users when  $\gamma^+ = 0.5$  is selected. Setting  $\gamma^+ = 1$  results in three main communities. Therefore, by emphasizing the importance of negative ties while treating positive links as neutrally relevant, we observe a structural tripartition of this empirical network.

cor	nmunity id.	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
$\gamma^+$	$\gamma^{-}$																			
0.5	0.5	10629	978	8	5	1	1	1	1	1	1	1	1							
0.5	1.0	10756	858	8	2	1	1	1	1											
1.0	0.5	6166	2902	2323	121	39	28	17	10	6	3	3	2	2	1	1	1	1	1	1
1.0	1.0	6059	3002	2530	13	9	6	3	2	2	1	1								

**Table 1.** Community-Spinglass on Menéame Network. We tested the *community-spinglass* algorithm on the co-voting network extracted from the social media platform Menéame. We used the parameter choices found most suitable on synthetic networks (see Figure 2A–B). When  $\gamma^+ = 0.5$ , we observe the presence of a single prominent community, followed by a secondary one with less than 1000 users. In contrast, with  $\gamma^+ = 1$ , the users are divided into three main communities, with approximately 6000, 3000, and 2000 users.

Next, we applied the *SPONGE* algorithm on real data, as highlighted in Table 2. Our findings indicate that by increasing the fixed number of clusters, a larger community of approximately 7000 users can be identified. Additionally, we observed that a community initially consisting of over 4400 users tended to bifurcate. It resulted in two distinct communities, one with around 3400 users and the other with approximately 680 users.

Since the Menéame data does not contain information about the political leanings of the users, we compare the overlap between users for the parameter choices when both algorithms find either two or three communities (Appendix Comparing algorithms on real data, Figure A9). In the case of two communities, there is no overlap between both methods. Both communities of *SPONGE*, of 7,087 and 4,396 users, get mapped to the largest community of *community-spinglass*, of 10,629 users.

We find considerable overlap in the case of three communities. 88–89% of the users in the two largest communities of *community-spinglass* are mapped to the two largest communities of *SPONGE*, while all users in the smallest community of *community-spinglass* (683 users) are mapped to the smallest community of *SPONGE* (2,323 users, the rest being mapped to the largest community of *community-spinglass*). Our results highlight the sensitivity of the results to parameter choice, and stress the need for more robust assessment and external validation when dealing with real data.

community id. no. of clusters	0	1	2	3	4	5	6	7	8	9
1	11628									
2	11591	37								

community id. no. of clusters	0	1	2	3	4	5	6	7	8	9
3	11583	37	8							
4	7087	4496	37	8						
5	7172	4410	25	13	8					
6	7247	4325	25	13	10	8				
7	7436	3471	665	25	13	10	8			
8	7401	3453	708	25	13	10	10	8		
9	7412	3459	687	25	13	10	10	8	4	
10	7417	3450	683	25	13	10	10	8	8	4

**Table 2. SPONGE on Menéame Network.** We tested the *SPONGE* algorithm on the network of interactions on Menéame. When we set the number of communities (clusters) to be 1–3, the algorithm over 99% of all users in one community. When we set the number of clusters to be larger than 3, the algorithm identifies a community with approximately 7000 users and a secondary one containing approximately 4400 users. When the number of clusters increases above 7, the secondary community is further split into two communities of 3400 and 600 users.

#### Dense (political) networks

#### Evaluation on Synthetic Networks

Offline voting contexts—such as legislative bodies—are typically dense: there are very few abstentions. This lack of "missing data" could simplify the task of community detection. However, we find the opposite to be true. Compared with the sparse case, both *SPONGE* and *community-spinglass* perform worse in all synthetic scenarios (Figure 3 and Figure A7 for the other similarity measures).

Similar to the sparse case, we find that both algorithms perform well when both users and stories are polarized (case *UP SP* in Figure 3A,C). Unlike the sparse case, increasing the fixed number of clusters in *SPONGE* decreases the performance (Figure 3C). When users are polarized but stories are not polarized, 50–80% of the user pairs are correctly identified, depending on parameter choice.

The results of the two algorithms differ when users are expected to be found in a single community (Figure 3B,D). *Community-spinglass* algorithm performs well in the case where users and stories are not polarized (for  $\gamma^- = 0.5$ ), and performs poorly when stories are polarized but users are not, irrespectively of parameter choice. *SPONGE* performs extremely poorly for the entire parameter space (except for the trivial case where we force the algorithm to find only one cluster). Increasing the number of clusters results in users being uniformly distributed across all communities.

#### Evaluation on real networks: US House of Representatives data

We applied the community detection methods to the US House of Representatives data and considered the political affiliations of the representatives (Democrats, Republicans, and Independents) as the "true" communities.

We find that *community-spinglass* correctly identifies the political communities by choosing  $\gamma^- < 2$  (Figure 4A) for all years except 2001. The algorithm is unable to find the political affiliation for 2001 with parameters  $\{\gamma^+, \gamma^-\} = \{1, 0.5\}$  and  $\{\gamma^+, \gamma^-\} = \{2, 0.5\}$ . This could be due to the 9/11 attacks. In that year, a large fraction of the legislation<sup>51</sup> is attributed to the attacks and received broad support across partian lines. We find an increasing correlation over time between voting behavior and political affiliation<sup>36</sup>–i.e., we are increasingly more able to recover the correct political affiliation from the votes alone. This reflects the increase in partiasanship levels through the years observed by Andris et al.<sup>34</sup> among others.

In the case of *SPONGE* (Figure 4B), the political affiliation of the representatives is correctly identified when the number of clusters, k, is set to 2–4, with a gradual decrease as k increases. Similar to the *community-spinglass* case, we evidence an increasing ability to recover political affiliations through the years.

The good recovery of political affiliations for a wide parameter range in both *community-spinglass* and *SPONGE* indicate that both US representatives ans US bills are polarized. This is perhaps unsurprising, since bills are themselves introduced by US representatives.

# Conclusion

In this paper, we systematically assessed the performance of community detection algorithms designed from unipartite signed networks on bipartite signed networks, using both synthetic and real-world data. We analyzed sparse networks resembling users' behavior on online social platforms, and dense networks resembling US politicians' votes on bills.



**Figure 3.** Community Detection on Dense Synthetic Networks. We tested the community detection methods on the four synthetic scenarios. Panels (A, C) show the results for the scenarios with two communities of users, while panels (B, D) show the results with one community. We used the Rand Index to evaluate the performance of the algorithms with different parameter choices. Higher values of the Rand Index indicate better alignment between expected and empirical communities. Panels (A, B) show the results for *community-spinglass*. We experimented with different combinations of the parameters  $\{\gamma^+, \gamma^-\} \in [0.5, 1, 2]$ . Lower values of  $\gamma^+$  indicate less importance given to positive ties in a community, whereas lower values of  $\gamma^-$  penalize the presence of negative links within a community. Note that scenarios where users and stories are either polarized or neutral (i.e., **U NP S NP** and **U P S P**) are correctly identified for low values of  $\gamma^-$ . Panels (C, D) show the results for *SPONGE*. We conducted tests for different values of the number of clusters *k*, ranging from *k* = 1 (no communities) to *k* = 10. The algorithm correctly identifies the expected communities in scenarios with polarized users (panel C), while it generates spurious communities when users are not polarized (panel D).

For sparse networks, the algorithms' performance depended on the underlying network structure. Both algorithms were successful at recovering communities in the scenarios where two communities of users existed but often identified several communities when only one community of users existed. The results were generally highly dependent on parameter tuning, emphasizing the importance of careful selection in real-world applications.

In the context of dense synthetic networks, both algorithms performed well when both user and story ideologies were polarized, but generated spurious communities in all other scenarios. This was particularly the case for *SPONGE* in the scenarios when only one community of users existed. Both methods well performed on the US House of Representatives data, finding an increasing level of political polarization over time.

In conclusion, our study contributes to the understanding of community detection in bipartite signed networks. The presented benchmark and systematic evaluation provide a foundation for future research in developing robust algorithms for various real-world scenarios. Future studies using community detection in bipartite signed networks can use our framework and open-source code to generate synthetic scenarios with realistic voting probabilities. These scenarios will enable researchers to



**Figure 4. Community Detection on US House of Representatives Networks.** We tested the community detection methods on 33 co-voting networks resulting from the (dis)agreement of members of the US House of Representatives on several bills, divided per year. We considered the subdivision into political parties (i.e., Democratic, Republican, or Independents) as the "true" communities. Panel (A) shows the results for *community-spinglass*. We found that parameter choices with  $\gamma^- < 2$  consistently capture the subdivision of Representatives into political parties. Panel (B) shows the results for *SPONGE*. Confidence intervals represent the standard deviation range among different runs. We conducted tests for different values of the number of clusters *k*, ranging from k = 1 (no communities) to k = 10. For k = 1, we consistently find that the true subdivision is not recovered. However, for subsequent values of *k*, we observe a peak for values  $k \in [2,4]$  in most of the networks, with a gradual decrease in the method's efficacy as *k* increases. Note that the Rand Index increased for more recent years.

select optimal parameters and evaluate the potential sensitivity of the results to parameter choice.

Two lessons can be learned from our results. First, the sensitivity of results to parameter choices underscores the importance of robustness tests. Using different parameter choices implies defining communities in different ways. For example, giving more weight to negative ties will create communities without dissent. Giving extra importance to positive ties may help to recover small organized groups. We emphasize the importance of using robustness tests to validate the recovered communities. In particular, external validation (e.g., political affiliation or clustering of user comments in online media) should be used when feasible.

The second lesson to be learned is the importance of negative ties. For both the dense and sparse network cases, giving extra importance to negative ties in *community-spinglass* resulted in enhanced inference of political affiliations in both sparse and dense networks. This indicates that negative ties are more indicative of communities than positive ties in voting networks.

The analysis presented has several limitations that open up additional avenues for future research. The first limitation is that we focus on three algorithms for community detection in signed networks frequently used in the literature. Future work could use our benchmark to compare an expanded number of algorithms.

The second limitation is the projection of bipartite networks into unipartite networks, which creates information loss. At the same time, we test a method that is able to work directly with the bipartite networks (Stochastic Block Model, see Section Weighted SBM with Edge Attributes, Figure A8), this algorithm is unable to recover the community structure for the scenarios when two communities are expected. We highlight the need for new methods tailored to bipartite signed networks.

### Author contributions statement

Conceptualization: All. Data curation: JGB, EC. Formal analysis: EC, SC. Methodology: JGB, EC, EK. Supervision: JGB, EK. Roles/Writing - original draft: EC. Writing - review editing: JGB, EK

# Data and Code Availability

The data and code associated with this research are publicly available on GitHub. Interested readers can access the repository at the following URL: <a href="https://github.com/elenacandellone/signed-bipartite-nets">https://github.com/elenacandellone/signed-bipartite-nets</a>. This repository includes the datasets used in the study, as well as the code implementations of the methods and algorithms discussed in the paper.

# References

- 1. Fortunato, S. Community detection in graphs. *Phys. Reports* **486**, 75–174, DOI: 10.1016/j.physrep.2009.11.002 (2010). 0906.0612.
- 2. Barber, M. J. Modularity and community detection in bipartite networks. *Phys. Rev. E* 76, 066102, DOI: 10.1103/ PhysRevE.76.066102 (2007).
- 3. Yen, T.-C. & Larremore, D. B. Community detection in bipartite networks with stochastic block models. *Phys. Rev. E* 102, 032309, DOI: 10.1103/PhysRevE.102.032309 (2020).
- 4. Beckett, S. J. Improved community detection in weighted bipartite networks. *R. Soc. open sci.* 3, 140536, DOI: 10.1098/rsos.140536 (2016).
- **5.** Doreian, P. & Mrvar, A. A partitioning approach to structural balance. *Soc. Networks* **18**, 149–168, DOI: 10.1016/ 0378-8733(95)00259-6 (1996).
- Tomasso, M., Rusnak, L. & Tešić, J. Advances in Scaling Community Discovery Methods for Signed Graph Networks. J. Complex Networks 10, cnac013, DOI: 10.1093/comnet/cnac013 (2022). 2110.07514.
- Cucuringu, M., Davies, P., Glielmo, A. & Tyagi, H. SPONGE: A generalized eigenproblem for clustering signed networks, DOI: 10.48550/arXiv.1904.08575 (2019). 1904.08575.
- Traag, V. A. & Bruggeman, J. Community detection in networks with positive and negative links. *Phys. Rev. E* 80, 036115, DOI: 10.1103/PhysRevE.80.036115 (2009). 0811.2329.
- 9. Doreian, P. & Mrvar, A. Partitioning signed social networks. *Soc. Networks* 31, 1–11, DOI: 10.1016/j.socnet.2008.08.001 (2009).
- **10.** Neal, Z. P. A sign of the times? Weak and strong polarization in the U.S. Congress, 1973–2016. *Soc. Networks* **60**, 103–112, DOI: 10.1016/j.socnet.2018.07.007 (2020).
- Leskovec, J., Huttenlocher, D. & Kleinberg, J. Signed networks in social media. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1361–1370, DOI: 10.1145/1753326.1753532 (ACM, Atlanta Georgia USA, 2010).
- 12. Kros, M., Jaspers, E. & van Zalk, M. Avoidance, antipathy, and aggression: A three-wave longitudinal network study on negative networks, status, and heteromisos. *Soc. Networks* 64, 122–133, DOI: 10.1016/j.socnet.2020.08.006 (2021).
- 13. Office of the Clerk, U.S. House of Representatives. https://clerk.house.gov/.
- 14. Menéame Portada de noticias de. https://www.meneame.net/.
- Jonsson, P. F., Cavanna, T., Zicha, D. & Bates, P. A. Cluster analysis of networks generated through homology: Automatic identification of important protein communities involved in cancer metastasis. *BMC Bioinforma*. 7, 2, DOI: 10.1186/1471-2105-7-2 (2006).
- Adamic, L. A. & Glance, N. The political blogosphere and the 2004 U.S. election: Divided they blog. In *Proceedings of the 3rd International Workshop on Link Discovery*, LinkKDD '05, 36–43, DOI: 10.1145/1134271.1134277 (Association for Computing Machinery, New York, NY, USA, 2005).
- Girvan, M. & Newman, M. Community structure in social and biological networks. *Proc. Natl. Acad. Sci. United States Am.* 99, 7821–7826, DOI: 10.1073/pnas.122653799 (2002).
- Talaga, S., Stella, M., Swanson, T. J. & Teixeira, A. S. Polarization and multiscale structural balance in signed networks. *Commun Phys* 6, 1–15, DOI: 10.1038/s42005-023-01467-8 (2023).
- Bonchi, F., Galimberti, E., Gionis, A., Ordozgoiti, B. & Ruffo, G. Discovering Polarized Communities in Signed Networks. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, CIKM '19, 961–970, DOI: 10.1145/3357384.3357977 (Association for Computing Machinery, New York, NY, USA, 2019).
- **20.** Heider, F. *The Psychology of Interpersonal Relations*. The Psychology of Interpersonal Relations. (John Wiley & Sons Inc, Hoboken, NJ, US, 1958).

- Cartwright, D. & Harary, F. Structural balance: A generalization of Heider's theory. *Psychol. Rev.* 63, 277–293, DOI: 10.1037/h0046049 (1956).
- 22. Estrada, E. & Benzi, M. Walk-based measure of balance in signed networks: Detecting lack of balance in social networks. *Phys. Rev. E* **90**, 042802, DOI: 10.1103/PhysRevE.90.042802 (2014).
- 23. Aref, S. & Wilson, M. C. Measuring Partial Balance in Signed Networks. J. Complex Networks 6, 566–595, DOI: 10.1093/comnet/cnx044 (2018). 1509.04037.
- 24. Estrada, E. Rethinking structural balance in signed social networks. *Discret. Appl. Math.* 268, 70–90, DOI: 10.1016/j.dam. 2019.04.019 (2019).
- Doreian, P. & Mrvar, A. Structural Balance and Signed International Relations. J. Soc. Struct. 16, 1–49, DOI: 10.21307/ joss-2019-012 (2015).
- **26.** Diaz-Diaz, F., Bartesaghi, P. & Estrada, E. Local balance of signed networks: Definition and application to reveal historical events in international relations. https://arxiv.org/abs/2303.03774v2 (2023).
- 27. Fraxanet, E., Pellert, M., Schweighofer, S., Gómez, V. & Garcia, D. Unpacking polarization: Antagonism and Alignment in Signed Networks of Online Interaction. https://arxiv.org/abs/2307.06571v2 (2023).
- 28. Aref, S. & Wilson, M. C. Balance and Frustration in Signed Networks. J. Complex Networks 7, 163–189, DOI: 10.1093/comnet/cny015 (2019). 1712.04628.
- **29.** Davis, J. A. Clustering and Structural Balance in Graphs. *Hum. Relations* **20**, 181–187, DOI: 10.1177/001872676702000206 (1967).
- **30.** Asratian, A. S. *Bipartite Graphs and Their Applications* (Cambridge, U.K.; New York : Cambridge University Press, 1998).
- **31.** Watts, D. J. & Strogatz, S. H. Collective dynamics of 'small-world' networks. *Nature* **393**, 440–442, DOI: 10.1038/30918 (1998).
- 32. Derr, T., Johnson, C., Chang, Y. & Tang, J. Balance in Signed Bipartite Networks (2019). 1909.06073.
- Lee, J., Li, G. & Wilson, J. D. Varying-coefficient models for dynamic networks. *Comput. Stat. & Data Analysis* 152, 107052, DOI: 10.1016/j.csda.2020.107052 (2020).
- **34.** Andris, C. *et al.* The Rise of Partisanship and Super-Cooperators in the U.S. House of Representatives. *PLoS ONE* **10**, e0123507, DOI: 10.1371/journal.pone.0123507 (2015).
- 35. Enelow, J. M. & Hinich, M. J. The Spatial Theory of Voting: An Introduction (CUP Archive, 1984).
- **36.** Poole, K. T. & Rosenthal, H. A Spatial Model for Legislative Roll Call Analysis. *Am. J. Polit. Sci.* **29**, 357–384, DOI: 10.2307/2111172 (1985). 2111172.
- **37.** Poole, K. T. *Spatial Models of Parliamentary Voting*. Analytical Methods for Social Research (Cambridge University Press, Cambridge ; New York, 2005).
- 38. Clinton, J., Jackman, S. & Rivers, D. The Statistical Analysis of Roll Call Data. Am. Polit. Sci. Rev. 98, 355–370, DOI: 10.1017/S0003055404001194 (2004).
- **39.** Yu, X. & Rodríguez, A. Spatial voting models in circular spaces: A case study of the U.S. House of Representatives. *The Annals Appl. Stat.* **15**, 1897–1922, DOI: 10.1214/21-AOAS1454 (2021).
- 40. Moody, J. & Mucha, P. J. Portrait of Political Party Polarization. Netw. sci. 1, 119–121, DOI: 10.1017/nws.2012.3 (2013).
- **41.** Waugh, A. S., Pei, L., Fowler, J. H., Mucha, P. J. & Porter, M. A. Party Polarization in Congress: A Network Science Approach. *arXiv* DOI: 10.48550/ARXIV.0907.3509 (2009).
- **42.** Zhang, Y. *et al.* Community Structure in Congressional Cosponsorship Networks. *Phys. A: Stat. Mech. its Appl.* **387**, 1705–1712, DOI: 10.1016/j.physa.2007.11.004 (2008). 0708.1191.
- 43. Mucha, P. J., Richardson, T., Macon, K., Porter, M. A. & Onnela, J.-P. Community Structure in Time-Dependent, Multiscale, and Multiplex Networks. *Science* 328, 876–878, DOI: 10.1126/science.1184819 (2010). 0911.1824.
- 44. Wilson, J. D., Stevens, N. T. & Woodall, W. H. Modeling and detecting change in temporal networks via a dynamic degree corrected stochastic block model (2016). 1605.04049.
- 45. Huang, J., Shen, H., Cao, Q., Tao, S. & Cheng, X. Signed Bipartite Graph Neural Networks (2021). 2108.09638.
- 46. Institute, A. T. SigNet package. https://github.com/alan-turing-institute/SigNet (2023).

- **47.** Reichardt, J. & Bornholdt, S. Statistical Mechanics of Community Detection. *Phys. Rev. E* **74**, 016110, DOI: 10.1103/ PhysRevE.74.016110 (2006). cond-mat/0603718.
- **48.** Csardi, G. & Nepusz, T. The Igraph Software Package for Complex Network Research. *InterJournal* **Complex Systems**, 1695 (2005).
- **49.** Rand, W. M. Objective Criteria for the Evaluation of Clustering Methods. J. Am. Stat. Assoc. **66**, 846–850, DOI: 10.1080/01621459.1971.10482356 (1971).
- 50. Pedregosa, F. et al. Scikit-learn: Machine learning in Python. J. Mach. Learn. Res. 12, 2825–2830 (2011).
- Mccarty, N. The Policy Effects of Political Polarization. In Pierson, P. & Skocpol, T. (eds.) *The Transformation of American Politics*, 223–255, DOI: 10.1515/9781400837502-013 (Princeton University Press, 2007).
- 52. Hubert, L. & Arabie, P. Comparing partitions. J. Classif. 2, 193–218, DOI: 10.1007/BF01908075 (1985).
- 53. Vinh, N. X., Epps, J. & Bailey, J. Information theoretic measures for clusterings comparison: Is a correction for chance necessary? In *Proceedings of the 26th Annual International Conference on Machine Learning*, 1073–1080, DOI: 10.1145/1553374.1553511 (ACM, Montreal Quebec Canada, 2009).
- 54. Rosenberg, A. & Hirschberg, J. V-Measure: A Conditional Entropy-Based External Cluster Evaluation Measure. In Eisner, J. (ed.) Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL), 410–420 (Association for Computational Linguistics, Prague, Czech Republic, 2007).
- **55.** Holland, P. W., Laskey, K. B. & Leinhardt, S. Stochastic blockmodels: First steps. *Soc. Networks* **5**, 109–137, DOI: 10.1016/0378-8733(83)90021-7 (1983).
- **56.** Peixoto, T. P. Bayesian Stochastic Blockmodeling. In *Advances in Network Clustering and Blockmodeling*, chap. 11, 289–332, DOI: 10.1002/9781119483298.ch11 (John Wiley & Sons, Ltd, 2019).
- 57. Peixoto, T. P. Nonparametric weighted stochastic block models. *Phys. Rev. E* 97, 012306, DOI: 10.1103/PhysRevE.97. 012306 (2018).
- 58. Peixoto, T. P. The graph-tool python library, DOI: 10.6084/M9.FIGSHARE.1164194 (2017).

# Appendix

## Ideology difference distribution

Section Generating synthetic bipartite signed networks shows the benchmark we devised to create synthetic networks with simulated voting behavior. To compute voting probabilities, we combined controlled scenarios with real data. In this section, we demonstrate that the distribution of the difference between two random variables is a Gaussian Mixture, under the assumption that the two random variables are also distributed as Gaussian Mixtures.

Take two random variables X and Y distributed as in Equations 2. Their characteristic functions are given by

$$\phi_{X}(t) = \mathbb{E}\left[e^{itX}\right] = \int_{\mathbb{R}} e^{itx} \left[\frac{\pi_{x}}{\sqrt{2\pi\sigma_{1}^{2}}}e^{-\frac{(x-\mu_{1})^{2}}{2\sigma_{1}^{2}}} + \frac{(1-\pi_{x})}{\sqrt{2\pi\sigma_{2}^{2}}}e^{-\frac{(x-\mu_{2})^{2}}{2\sigma_{2}^{2}}}\right] dx = \pi_{x} \cdot e^{-\frac{1}{2}t^{2}\sigma_{1}^{2} + it\mu_{1}} + (1-\pi_{x}) \cdot e^{-\frac{1}{2}t^{2}\sigma_{2}^{2} + it\mu_{2}},$$
  
$$\phi_{Y}(t) = \mathbb{E}\left[e^{itY}\right] = \int_{\mathbb{R}} e^{ity} \left[\frac{\pi_{y}}{\sqrt{2\pi\sigma_{3}^{2}}}e^{-\frac{(x-\mu_{3})^{2}}{2\sigma_{3}^{2}}} + \frac{(1-\pi_{y})}{\sqrt{2\pi\sigma_{4}^{2}}}e^{-\frac{(x-\mu_{4})^{2}}{2\sigma_{4}^{2}}}\right] dy = \pi_{y} \cdot e^{-\frac{1}{2}t^{2}\sigma_{3}^{2} + it\mu_{3}} + (1-\pi_{y}) \cdot e^{-\frac{1}{2}t^{2}\sigma_{4}^{2} + it\mu_{4}}.$$

Then, the random variable X - Y has the following characteristic function

$$\begin{split} \phi_{X-Y}(t) &= \mathbb{E}\left[e^{it(X-Y)}\right] = \mathbb{E}\left[e^{itX}\right] \mathbb{E}\left[e^{-itY}\right] = \\ &= \left[\pi_x \cdot e^{-\frac{1}{2}t^2\sigma_1^2 + it\mu_1} + (1-\pi_x) \cdot e^{-\frac{1}{2}t^2\sigma_2^2 + it\mu_2}\right] \cdot \left[\pi_y \cdot e^{-\frac{1}{2}t^2\sigma_3^2 - it\mu_3} + (1-\pi_y) \cdot e^{-\frac{1}{2}t^2\sigma_4^2 - it\mu_4}\right] = \\ &= \pi_x \cdot \pi_y \cdot \left(e^{-\frac{1}{2}t^2(\sigma_1^2 + \sigma_3^2) + it(\mu_1 - \mu_3)}\right) + \pi_x \cdot (1-\pi_y) \cdot \left(e^{-\frac{1}{2}t^2(\sigma_1^2 + \sigma_4^2) + it(\mu_1 - \mu_4)}\right) + \\ &+ (1-\pi_x) \cdot \pi_y \cdot \left(e^{-\frac{1}{2}t^2(\sigma_2^2 + \sigma_3^2) + it(\mu_2 - \mu_3)}\right) + (1-\pi_x) \cdot (1-\pi_y) \cdot \left(e^{-\frac{1}{2}t^2(\sigma_2^2 + \sigma_4^2) + it(\mu_2 - \mu_4)}\right), \end{split}$$

that lead to the conclusion that X - Y is distributed as of Equation 3.

### **Degree-corrected synthetic networks**

To generate more realistic synthetic bipartite signed networks, we can extend the procedure explained in Subsection Generating synthetic bipartite signed networks. Instead of using fixed voting probabilities for all users, we can employ a voting probability distribution. This implies that we can replace step 3 in the previous mechanism with a range of probabilities, as follows:

- 1. Extract from data the positive and negative voting distributions;
- 2. Assign randomly to each user a positive and negative "degree", keeping the same coupling as in the data;
- 3. Repeat the voting procedure as before, but when the user reaches their assigned "degree", stop assigning votes.

With this modification, we aim to preserve the skewness of the voting distribution. In other words, we sample a realistic voting behavior from the data. It is particularly meaningful in the case of sparse networks, where a few users vote for almost every story, while most users will barely vote a few times.

Figure A5 displays the results for sparse networks. Similar to the case of uniform voting probabilities, *community-spinglass* better captures the expected communities for lower values of  $\gamma^-$ , except for the scenario where both users and stories aren't polarized (**U NP S NP**). One possible explanation lies in the increased complexity of the model, as users exhibit similar voting behaviors compared to real data, making them less identifiable within a neutral, unique group. However, *SPONGE* outperforms all scenarios, identifying approximately 80% of the user pairs in the correct clusters, except for the scenario **U NP S P**, where increasing the number of clusters leads to a decrease in accuracy. We observe that, when introducing additional complexity to the model, community detection methods perform similarly or even worse than in the uniform voting case.



**Figure A5.** Community Detection on Menéame Synthetic Networks, degree-corrected version. We tested the community detection methods on the four synthetic scenarios, with voting probabilities sampled from sparse data's voting distributions. Panels (A)-(B) show the results for *community-spinglass*. Lower values of  $\gamma^+$  indicate less importance given to positive ties in a community, whereas lower values of  $\gamma^-$  penalize the presence of negative links within a community. Note that  $\gamma^- = 0.5$  is generally the best parameter choice, as it finds the expected communities for the synthetic scenarios, except for the case where users and stories aren't polarized. Panels (C)-(D) show the results for *SPONGE*. We observed that the algorithm correctly identifies the expected communities in scenarios with polarized users (panel C), while it generates spurious communities in cases where stories introduce latent ideologies (U NP S P case).

#### **Clustering Evaluation Metrics**

Section Evaluating Community Similarity explained the metric used to evaluate the similarity between the expected communities, both in real data with ground truth and in synthetic networks, to the empirical ones found by the community detection methods. However, the Rand Index is not the only metric available for clustering evaluation. It is possible to correct the RI for random assignments, and it is called the *Adjusted Rand Index*<sup>52</sup>, defined as:

$$ARI = \frac{RI - E[RI]}{\max(RI) - E[RI]}.$$
(9)

Another metric implemented in the *scikit-learn* package is based on the information-theoretic concept of Mutual Information (MI)<sup>53</sup>. The MI is defined as

$$\mathrm{MI}(\boldsymbol{\sigma}_{\mathrm{T}}, \boldsymbol{\sigma}_{\mathrm{E}}) = \sum_{i=1}^{|\boldsymbol{\sigma}_{\mathrm{E}}|} \sum_{j=1}^{|\boldsymbol{\sigma}_{E}^{i} \cap \boldsymbol{\sigma}_{E}^{j}|} \log\left(\frac{N|\boldsymbol{\sigma}_{T}^{i} \cap \boldsymbol{\sigma}_{E}^{j}|}{|\boldsymbol{\sigma}_{T}^{i}||\boldsymbol{\sigma}_{E}^{j}|}\right),\tag{10}$$

and can be normalized, obtaining the so-called Normalized Mutual Information,

$$NMI(\sigma_{\mathbf{T}}, \sigma_{\mathbf{E}}) = \frac{MI(\sigma_{\mathbf{T}}, \sigma_{\mathbf{E}})}{mean(H(\sigma_{\mathbf{T}}), H(\sigma_{\mathbf{E}}))},$$
(11)

where  $\sigma_{\rm T}$  and  $\sigma_{\rm E}$  are respectively the true and empirical configurations. Another possible clustering evaluation measure is called the v-score<sup>54</sup>. It is defined as the harmonic mean of two functions, the homogeneity *h* and the completeness *c*, defined as

$$h = 1 - \frac{H(C|K)}{H(C)}$$
$$c = 1 - \frac{H(K|C)}{H(K)}.$$

Figure A6 complements the results presented in Figure 2 by including the other evaluation metrics explained in this section. We show that metrics such as the ARI, NMI, and v-score, while generally preferable due to their adjustments for randomness, may fail to evaluate cases where we expect to find a single community. We find a coherent pattern in the case of the US House of Representatives Networks (Figure A7), where we chose as expected communities the political affiliations (i.e., Democrats, Republicans, and Independent), therefore the four metrics highlight the same behavior as the Rand Index.



**Figure A6.** Community Detection on Menéame Synthetic Networks, comparing several evaluation metrics. We conducted tests on four synthetic scenarios to evaluate community detection methods. We used four evaluation metrics: Rand Index (RI) in panels (A)-(E), Adjusted Rand Index (ARI) in panels (B)-(F), Normalized Mutual Information (NMI) in panels (C)-(G), and v-score in panels (D)-(H). We compared the results of *community-spinglass* and *SPONGE*. We found that the last three metrics showed similar behavior in failing to identify cases where one community was expected.



**Figure A7.** Community Detection on US House of Representatives Networks, various evaluation metrics. We tested the community detection methods on the US House of Representatives data. We compare four evaluation metrics: Rand Index (RI) in panels (A)-(E), Adjusted Rand Index (ARI) in panels (B)-(F), Normalized Mutual Information (NMI) in panels (C)-(G), and v-score in panels (D)-(H). We show the results for *community-spinglass* and *SPONGE*. As we identified the "true" subdivision as a tripartition in political affiliations, we found a coherent pattern among metrics.

year	number of nodes	number of edges	Ι	year	number of nodes	number of edges
1990	435	94169		2007	444	98202
1991	441	96716		2008	445	97982
1992	435	94232		2009	447	99444
1993	443	97315		2010	447	99543
1994	442	97290		2011	440	96530
1995	443	97787		2012	439	95854
1996	439	95844		2013	443	97541
1997	439	95334		2014	438	95527
1998	439	95485		2015	445	98239
1999	439	95833		2016	438	95437
2000	436	94671		2017	443	97331
2001	444	98122		2018	442	96777
2002	437	94981		2019	449	100026
2003	436	94603		2020	443	97293
2004	438	94977		2021	441	96943
2005	438	95464		2022	449	99752
2006	437	95035				

### Descriptive statistics of co-voting networks of US House of Representatives

**Table A3.** Snapshot co-voting networks of US House of Representatives. We collected data from 1990 to 2022 of the House of Representatives votes. We then subdivided the dataset per year and we generated the bipartite network of Representatives and Bills into (dis)agreement unipartite networks where the nodes are the Representatives and the edges are the co-votings between two Representatives. The weight of those edges is given by Equation 1.

#### Weighted SBM with Edge Attributes

The Stochastic Block Model  $(SBM)^{55}$  is a generative process for creating networks that exhibit a partition into blocks. This model is highly flexible, allowing not only the generation of networks with a predefined block structure but also the inference of communities in existing networks<sup>56</sup>. Given our context of weighted signed networks, we implemented the weighted version of the SBM<sup>57</sup>, incorporating the sign as an additional edge attribute. We utilized the SBM version in the Python package *graph-tool*<sup>58</sup>. As suggested in the package documentation, we sampled the logarithm of the absolute value of edge weights from a normal distribution and the edge sign, rescaled to values between 0 and 1, from a Bernoulli distribution. In contrast with previous methods, frustration minimization is not the explicit mechanism to generate communities. Therefore, it would also be possible to find communities with maximal frustration, with only negative edges inside a community and positive between communities.

Figure A8 depicts the results obtained by applying SBM to both the bipartite network, without the projection (with the option to add set membership as a node attribute), and the unipartite, projected networks. In both cases, we employed a hierarchical SBM. However, in both cases, the algorithm does not correctly identify the expected communities. Nonetheless, this is not necessarily a limitation, as the algorithm is not constrained to a specific optimization problem like previous algorithms. Instead, at the lower hierarchical levels it detects numerous smaller groups that may represent different optimal scenarios compared to the expected ones.



**Figure A8.** Stochastic Block Model on Sparse Synthetic Networks. We applied the hierarchical SBM to four sparse synthetic scenarios. Panel (A) displays the results on the bipartite network, i.e., the user-story network without projecting to a user-user network. An additional node attribute, the set membership, is given as a parameter of the algorithm. Panel (B) illustrates the results of the bipartite projection. We observe that in both cases, after two hierarchical levels, all four scenarios collapse into a unique community. This results in the maximum Rand Index (RI) for the U NP cases and approximately 50% for the other two cases. However, examining the lower hierarchical levels reveals that the SBM identifies a high number of communities for all scenarios, in contrast to the other algorithms analyzed in this paper.

#### Comparing algorithms on real data

Results subsection Menéame data shows that, depending on the choice of the parameters, both methods recover either two or three main communities in the Menéame Network. To further investigate the accuracy of these results, we examine whether the primary communities identified by each algorithm were the same or not. Figure A9 displays two distinct parameter choices that generate either two or three main communities. We observe that in the former case, communities 0 and 1 contain 60% and 40% of the largest community respectively, as detected by *SPONGE*. In the latter case, both *SPONGE* and *community-spinglass* accurately detected communities 0 and 1 with an accuracy of around 88-89%. These results show that although the methods follow the same principle of frustration minimization, they produce varying results and do not consistently generate stable communities for all parameter choices.



**Figure A9.** Comparing methods results on Menéame data. We found that when using specific parameter choices, both community detection methods recover either two or three main communities. Panel (A) displays the results for parameter choices that yield two larger communities for both algorithms. It is normalized by row and smaller communities are not shown. We observed that 60% of community 0 in *community-spinglass* is shared with community 0 in *SPONGE*, with the remaining 40% found in community 1 of SPONGE. Panel (B) shows the results for parameter configurations presenting three main communities. Communities 0 and 1 from *community-spinglass* can be identified respectively as communities 0 and 1 from *SPONGE* with an accuracy of around 88-89%.