

Centrality with Diversity

Liang Lyu
Duke University
Durham, NC, USA
lianglyu1998@gmail.com

Kamesh Munagala
Duke University
Durham, NC, USA
kamesh@cs.duke.edu

Brandon Fain
Duke University
Durham, NC, USA
btfain@cs.duke.edu

Kangning Wang
Duke University
Durham, NC, USA
knwang@cs.duke.edu

ABSTRACT

Graph centrality measures use the structure of a network to quantify central or “important” nodes, with applications in web search, social media analysis, and graphical data mining generally. Traditional centrality measures such as the well known PAGERANK interpret a directed edge as a vote in favor of the importance of the linked node. We study the case where nodes may belong to diverse communities or interests and investigate centrality measures that can identify nodes that are simultaneously important to many such diverse communities. We propose a family of diverse centrality measures formed as fixed point solutions to a generalized nonlinear eigenvalue problem. Our measure can be efficiently computed on large graphs by iterated best response and we study its normative properties on both random graph models and real-world data. We find that we are consistently and efficiently able to identify the most important diverse nodes of a graph, that is, those that are simultaneously central to multiple communities.

CCS CONCEPTS

• **Information systems** → Page and site ranking; • **Mathematics of computing** → Graph algorithms.

KEYWORDS

Graph centrality, PageRank, Fairness

ACM Reference Format:

Liang Lyu, Brandon Fain, Kamesh Munagala, and Kangning Wang. 2021. Centrality with Diversity. In *Proceedings of the Fourteenth ACM International Conference on Web Search and Data Mining (WSDM '21), March 8–12, 2021, Virtual Event, Israel*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3437963.3441789>

1 INTRODUCTION

A fundamental question in network analysis concerns *graph centrality*: Which nodes are “important” to understanding the overall

network structure? Particular measures of centrality vary from node degrees to citation indices in publication networks [19] to the celebrated PAGERANK originally developed for web search [30]. In such work, it is typically assumed that the graph is composed of a single homogeneous community of nodes. We explore the question of centrality in a context of graphs with many heterogeneous communities and propose to answer the question:

Which nodes are simultaneously “important” to multiple diverse communities within the graph?

As a concrete example, consider the Twitter graph of members of the United States Congress (where the members are nodes and directed edges represent following relationships on Twitter). Traditional graph centrality measures attempt to quantify the importance of particular members within the social network of the Congress as a whole. It is well known that the United States Congress is (roughly) composed of two communities: Democrats (or liberals) and Republicans (or conservatives). We are interested in the following question: Who are the most important members in the Congress, both in terms of network structure and bipartisanship? We address this particular example in our real-world experiments in Section 6. As another example, suppose we study publication networks composed of several sub-disciplines and are interested in identifying papers or authors who are most important to several sub-disciplines simultaneously or who serve as interdisciplinary bridges between publication communities.

To address these questions, we consider a model where in addition to a directed graph, we have a score vector for each node corresponding to community affiliation. These score vectors are real-valued and denote the degree of affiliation with different communities. We say that a node is *polarized* or *biased* if it is mostly affiliated with a particular community, and *balanced* if it is comparably affiliated with all communities. We begin with the same basic intuition as other classic measures of graph centrality, that a directed edge from one node to another constitutes a “citation” or “vote in favor” of the target node. In the standard PAGERANK model, the centrality is a fixed point solution that sets the measure of a node u to the average of the measures of the nodes that cite u . Thus, a node could have a high PAGERANK even if it is only cited by members of its own community. On the other hand, in our model, a node should only receive a high DIVERSE CENTRALITY if it is cited by nodes of high DIVERSE CENTRALITY score from multiple communities. Informally, the DIVERSE CENTRALITY of a node u therefore

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

WSDM '21, March 8–12, 2021, Virtual Event, Israel

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8297-7/21/03...\$15.00

<https://doi.org/10.1145/3437963.3441789>

is the *minimum* over communities of the average DIVERSE CENTRALITY of the nodes of that community that cite u . We present the formal and more general definition in Section 2.

1.1 Contributions and Outline

Our first contribution in Section 2 is the proposal of a general family of centrality measures called DIVERSE CENTRALITY that capture the idea of centrality with diversity in graphs with community structure. DIVERSE CENTRALITY is the fixed point solution to a generalized nonlinear eigenvalue problem; we first show that such a solution is guaranteed to exist and provide an iterative algorithm for computing a solution.

In Section 3 we define our simulation setup for studying DIVERSE CENTRALITY on random graphs motivated by real-world networks, including Erdős-Rényi graphs [12], preferential attachment graphs [3], and a polarity attachment variant inspired by affiliation networks [28] that we introduce. In Section 4, we show empirically that DIVERSE CENTRALITY can be computed efficiently via iterated best response and that it converges to a unique equilibrium.

In Section 5, we present three criteria that a DIVERSE CENTRALITY measure should satisfy. (1) Centrality: The measure should be greater for nodes with greater PAGERANK. (2) Local Polarity: The measure should be higher for balanced nodes than polarized nodes. (3) Neighborhood Polarity: The measure should be higher for nodes with diverse citations. We show by simulations on the random graph models that in contrast to simple methods of re-weighting PAGERANK, DIVERSE CENTRALITY simultaneously satisfies all of these properties. An added advantage of DIVERSE CENTRALITY compared to simple re-weighting is that it propagates diverse influence in the sense that it amplifies scores of nodes that are pointed to by nodes that are also influential and balanced.

In Section 6 we study DIVERSE CENTRALITY on real-world graphs. We consider two graphs: a network of political blogs and the Twitter network of the United States Congress. Both examples have externally verifiable community structure. We show that in addition to the aforementioned properties, the DIVERSE CENTRALITY model favors nodes that are not only influential (in terms of PAGERANK), but that are also bipartisan in their neighborhood structure. Interestingly, we show that for the United States Congress, the top-scoring senators found by our algorithm correlate well with the senators that join the most number of bipartisan bills, while not all these senators are as highly ranked by standard PAGERANK.

Given our empirical observation that DIVERSE CENTRALITY finds central nodes that also serve as bridges between communities, in Section 6.1, we propose and investigate a natural generalization of BETWEENNESS CENTRALITY [14], which we call DIVERSE BETWEENNESS CENTRALITY. Though DIVERSE BETWEENNESS CENTRALITY might appear better suited to capture our desiderata, we show DIVERSE CENTRALITY dominates DIVERSE BETWEENNESS CENTRALITY in our experiments, and is indeed the preferred method for achieving centrality with diversity.

1.2 Related Work

The theory and practice of PAGERANK-based notions of graph centrality are significant topics in the literature on network analysis

and data mining [17, 21, 23, 24, 30]. Much of the literature is concerned with the efficient computation of PAGERANK on extremely large networks typical of the web graph. By contrast, we are primarily interested in understanding an alternate notion of centrality with diversity in a graph with community structure.

There is a long line of work on algorithms for community detection in social networks [5, 27, 29, 35, 36]. In contrast with this work, we are interested in graphs with known community structure or affiliation and how that influences the determination of a diverse centrality measure.

There is also a rich literature on modeling how communities and polarizing opinions arise in evolving networks [10, 15, 18, 34]. Polarization of opinions has become an increasingly important societal topic, since it is largely a corrosive process that leads to the splintering of society and the formation of opinion bubbles. Our paper can be viewed as taking a step towards mitigating the effect of polarization. Our centrality measure suggests a way to calculate influence in a “depolarized” manner by prioritizing nodes that are simultaneously important and bridge different viewpoints.

The DIVERSE CENTRALITY measure can be viewed as enforcing the concept of max-min (or Rawlsian) fairness on centrality measures. Related notions of fairness have been widely studied in resource allocation [7, 8, 33], classification [4, 11, 16, 26, 31], and so on. The main difference in our work is that we enforce such fairness locally at each neighborhood in the graph, and examine the effect of propagating such a notion through the network. We show that such a process leads to interesting insights even in real-world networks.

The influential work of Everett and Borgatti [13] (see also [25]) proposes a number of measures for the centrality of a given group of network nodes. These are defined by extending the notions of degree centrality and BETWEENNESS CENTRALITY to measure the extent to which group members serve as “bridge nodes” in the network. For instance, group betweenness is defined as the sum over all pairs of nodes of the fraction of the shortest paths between those nodes that use group members. Our focus is on the case where there are many groups and each node’s group membership is *fractional*. In Section 6, we propose a natural generalization of BETWEENNESS CENTRALITY to our setting; however, we show empirically that it is not as effective as DIVERSE CENTRALITY.

Finally, just as PAGERANK is a special case of the eigenvalue problem, the DIVERSE CENTRALITY measure is a special case of *non-linear eigenvalue* problems, which have a long and rich history in functional analysis and numerical methods; see [9] for a survey. In contrast to the standard eigenvalue problem that admits to efficient computation, there are no general efficient algorithms known for the non-linear generalizations. Nevertheless, we find that our notion of DIVERSE CENTRALITY is simple enough that algorithms similar to those for computing PAGERANK are effective in practice.

2 DIVERSE CENTRALITY

In this section we define DIVERSE CENTRALITY as a fixed point solution to a generalized nonlinear eigenvalue problem and prove that a solution always exists. We then present an iterative algorithm to find such an equilibrium.

Let (V, E) be a directed graph with n nodes $V = \{1, 2, \dots, n\}$. Let d_i be the outdegree of node i . In addition, we assume there

are K communities within the network. For each node i , there is an *affiliation vector* $\mathbf{q}^{(i)} \in \mathbb{R}_+^K$ with $\|\mathbf{q}^{(i)}\|_1 = 1$ characterizing the communities it belongs to.

To be concrete, consider again the example of the Twitter graph of members of the US Congress. There are $K = 2$ communities. $q_1^{(i)}$ and $q_2^{(i)} = 1 - q_1^{(i)}$ describe to what extent node i belongs to the Democratic and Republican communities. A node with $|q_1^{(i)} - q_2^{(i)}|$ close to 1 would be considered highly polarized or partisan; a value close to 0 would indicate balance or bipartisanship.

Now we can define DIVERSE CENTRALITY. To get the intuition, consider a single node i . We calculate the centrality of i according to each community in turn by taking the average centrality of nodes that cite i , weighted by their affiliation with that community, and adding a damping term based on i 's own affiliation with that community. Then we take the minimum (or more generally some concave function) of these scores over all of the communities so that i 's centrality is limited by the communities for which i is less central. This defines a nonlinear iterative procedure, to which DIVERSE CENTRALITY is the fixed point solution.

Definition 1. The *DIVERSE CENTRALITY* of a graph (V, E) with damping factor $p \in (0, 1)$ is a vector $\mathbf{s} \in \mathbb{R}_{\geq 0}^n$ such that $\|\mathbf{s}\|_1 = 1$, and for all $i \in V$,

$$\lambda s_i = f\left((1-p)\frac{\mathbf{q}^{(i)}}{n} + p \sum_{j:(j,i) \in E} \frac{s_j}{d_j} \mathbf{q}^{(j)}\right)$$

where $f: \mathbb{R}_{\geq 0}^K \rightarrow \mathbb{R}_{\geq 0}$ is a concave function and $\lambda \in \mathbb{R}_+$ is the normalization constant. We require f to map any strictly positive vector to a positive number, i.e. $x_i > 0 \forall i \in [K]$ implies $f(\mathbf{x}) > 0$.

Note that if $K = 1$ or $f(\mathbf{x}) = \|\mathbf{x}\|_1$, then the notion becomes the standard PAGERANK. We are interested in the case where $K \geq 2$ and f is concave to prioritize diversity. For example, we desire $f((0.5, 0.5))$ to be greater than $f((0.1, 0.9))$ or $f((0.9, 0.1))$. Some convenient choices of f include the minimum function and the geometric mean. In our experiments, we focus on the minimum.

The damping terms of $(1-p)\frac{\mathbf{q}^{(i)}}{n}$ are similar to those in standard PAGERANK – $\frac{1-p}{n}$ – but multiplied by $\mathbf{q}^{(i)}$ to account for the local polarity of the node. We use a damping factor of $p = 0.85$ for PAGERANK and DIVERSE CENTRALITY in our experiments.

2.1 Existence of Equilibrium

Now we establish the existence of DIVERSE CENTRALITY.

THEOREM 1. *The DIVERSE CENTRALITY of a directed graph always exists, that is, there exists \mathbf{s} that satisfies Definition 1.*

PROOF. We use Brouwer's fixed-point theorem, which states that any continuous mapping from a compact and convex set to itself must have a fixed point, to show an equilibrium of DIVERSE CENTRALITY always exists. Define function G that maps the simplex $[0, 1]^n \cap \{\mathbf{s} : \|\mathbf{s}\|_1 = 1\}$ to itself such that $G_i(\mathbf{s}) = \frac{t_i}{\sum_{j=1}^n t_j}$, where

$$t_i = f\left((1-p)\frac{\mathbf{q}^{(i)}}{n} + p \sum_{j:(j,i) \in E} \frac{s_j}{d_j} \mathbf{q}^{(j)}\right), \forall i \in [n].$$

Intuitively, the function G is almost f except that it also maps the output of f back to the simplex proportionally. Note that by Definition 1, f is concave and hence continuous; therefore, G is also continuous as it is the ratio between two continuous positive functions (each t_j is positive so the denominator is positive). Furthermore, $G_i(\mathbf{s}) > 0$ for all $i \in [n]$. The function G is thus a continuous mapping from the compact and convex set $[0, 1]^n \cap \{\mathbf{s} : \|\mathbf{s}\|_1 = 1\}$ to itself. By Brouwer's fixed-point theorem, G has a fixed point \mathbf{s} , i.e. $G(\mathbf{s}) = \mathbf{s}$. This \mathbf{s} satisfies our definition of DIVERSE CENTRALITY that s_i is proportional to t_i . \square

2.2 Algorithm for DIVERSE CENTRALITY

We present Algorithm 1 to compute DIVERSE CENTRALITY. It starts from a random weight vector $\mathbf{s}^{(0)}$ and iteratively updates itself using the formula in Definition 1. It terminates when an iteration does not change the weight vector much: When the L^1 -distance between the weight vectors of two iterations is smaller than the given precision ϵ .

Algorithm 1: DIVERSE CENTRALITY

Input : $(V, E), p, f(\cdot), \epsilon, \{\mathbf{q}^{(i)}\}_{i \in [n]}$

- 1 Choose any $\mathbf{s}^{(0)} \in \mathbb{R}_{\geq 0}^n$ with $\|\mathbf{s}^{(0)}\|_1 = 1$, and set $k \leftarrow 0$;
- 2 **do**
- 3 $k \leftarrow k + 1$;
- 4 $t_i \leftarrow f\left((1-p)\frac{\mathbf{q}^{(i)}}{n} + p \sum_{j:(j,i) \in E} \frac{s_j^{(k-1)}}{d_j} \mathbf{q}^{(j)}\right), \forall i \in [n]$;
- 5 $s_i^{(k)} \leftarrow \frac{t_i}{\sum_{j=1}^n t_j}, \forall i \in [n]$;
- 6 **while** $\|\mathbf{s}^{(k)} - \mathbf{s}^{(k-1)}\|_1 > \epsilon$;
- 7 **return** $\mathbf{s}^{(k)}$;

Experimentally, we find that on most random graphs the algorithm converges to an equilibrium in a small number of iterations and the equilibrium is most likely unique. The experiments are described in detail in Section 4.

3 SIMULATION SETUP

In this section, we describe the simulation setup that we will use to study the convergence behavior of the algorithm, as well as its properties relative to simpler baselines. We perform this study using a collection of well-motivated random graph models.

For simplicity, all experiments in subsequent sections focus on the special case of $K = 2$. We will refer to the two communities as blue and red, with b_i and r_i scores to denote node i 's affiliation to each community. Formally, $b_i = q_1^{(i)}$ and $r_i = q_2^{(i)}$. We will sometimes refer to these values as ‘‘polarities’’ of a node; note that having either r_i or b_i close to 1 indicates polarization within that community. We use a damping factor of $p = 0.85$ for PAGERANK and DIVERSE CENTRALITY in all the experiments in Sections 4 and 6. Further, the function f in Definition 1 is chosen to be the minimum.

3.1 Random Graph Models

We use three classes of randomly generated undirected graphs, along with their modifications, in our experiments.

A FULLY RANDOM graph is generated using the Erdős–Rényi model [12]. The graph contains $n = 1000$ nodes, and each node’s r_i value is drawn from a uniform distribution on $(0, 1)$ with $b_i = 1 - r_i$. Each edge is included with probability $e = 0.2$.

A PREFERENTIAL ATTACHMENT graph is generated by a preferential attachment mechanism such as the Barabási–Albert model [3]. The graph begins with a clique of $m = 20$ nodes. New nodes are then added one at a time until a total of $n = 1000$ nodes, and each new node i is randomly connected to m existing nodes with probability proportional to their degrees. Formally, the m nodes are drawn from the distribution $p_j = \frac{d_j}{\sum_{k:k < i} d_k} \forall j < i$. All nodes in the graph have random polarities (i.e. r_i is drawn uniformly at random from $(0, 1)$). This generative model simulates social networks without any consideration of polarized community formation.

A POLARITY ATTACHMENT graph has nodes that are connected based on how similar their polarities are, resulting in polarized clusters. This is motivated by affiliation networks [28]. The graph contains $n = 1000$ nodes with random polarities, and each pair of nodes i and j is connected with a probability of $p_{ij} = 0.5 (r_i r_j + b_i b_j)$. Nodes with similar polarities are thus more likely to be connected. This generative model simulates the creation of polarized neighborhoods in the graph, such as what may occur when people are more likely to connect to others with similar opinions as their own.

4 CONVERGENCE AND UNIQUENESS

In this section, we compute DIVERSE CENTRALITY scores and PAGERANK scores on the random graphs introduced in Section 3.1. We compare the rate of convergence of the two models in order to assess the efficiency of Algorithm 1. We also examine the uniqueness of DIVERSE CENTRALITY scores found by Algorithm 1.

4.1 Rate of Convergence

Figure 1 shows the distribution of the number of iterations of PAGERANK and DIVERSE CENTRALITY on FULLY RANDOM, PREFERENTIAL ATTACHMENT and POLARITY ATTACHMENT graphs, when $\epsilon = 10^{-10}$. Each run uses a different randomly generated graph, and for both models the importance vector is initialized to $s_i = \frac{1}{n}$ for $1 \leq i \leq n$.

The results suggest that not only does the DIVERSE CENTRALITY model always converge to an equilibrium in practice, but its number of iterations is typically within a constant factor of PAGERANK. The rate of convergence is the slowest on FULLY RANDOM graph with an average of 23.883 iterations, or 2.99 times of PAGERANK.

These experiments offer empirical evidence that the DIVERSE CENTRALITY model converges in practice, and its rate of convergence is competitive compared to traditional centrality algorithms such as PAGERANK.

4.2 Uniqueness of Equilibrium

We also examine whether Algorithm 1 always converges to a unique equilibrium independent of initialization. For this, we generate 600 runs each for the 3 graphs in Section 3.1. Each run uses a different random graph, and in each run Algorithm 1 is executed twice, one




Figure 1: Number of iterations of PAGERANK and DIVERSE CENTRALITY on different graphs over 1000 runs.

with a constant initial vector s such that $s_i = \frac{1}{n} = 10^{-3} \forall 1 \leq i \leq n$, and the other with a random s . Each execution converges to a local equilibrium with $\epsilon = 10^{-10}$ (L^1 -norm), so we calculate and analyze the differences of each node’s scores between the two executions.

We also observe that the maximum difference of any node’s score between the constant initialization and random initialization executions observed is 3.544×10^{-11} (over all 1800 runs on different graphs). This is insignificant compared to the average node score of $\frac{1}{n} = 10^{-3}$. Therefore, these experiments provide strong empirical evidence that Algorithm 1 always converges to a unique global equilibrium, regardless of the initial vector s .

5 NORMATIVE PROPERTIES

In this section, we compare the quality of the solution found by DIVERSE CENTRALITY against simpler baselines. First, we propose three criteria that a centrality measure with diversity should satisfy:

Centrality. The measure should be higher for nodes that are more central according to standard graph centrality measures such as PAGERANK.

Local Polarity. The measure should be higher for nodes that belong to multiple communities in the sense of having balanced rather than polarized affiliation scores.

Neighborhood Polarity. The measure should be higher for nodes that bridge communities in the sense that they are cited by multiple communities.

5.1 Baseline Algorithms

We compare our DIVERSE CENTRALITY model against three baseline centrality measures listed below. Note that by definition, the PAGERANK model captures centrality, while REWEIGHT NODE BIAS and REWEIGHT NEIGHBOR BIAS capture local polarity and neighborhood polarity respectively.

- The PAGERANK model, with damping factor $p = 0.85$.
- REWEIGHT NODE BIAS: Re-weight the score vector of PAGERANK proportional to $w_i = \min\{r_i, b_i\}$. This means balanced nodes having greater w_i will be considered more important.

- **REWEIGHT NEIGHBOR BIAS:** Re-weight the score vector of PAGERANK proportional to $w_i = \min \left\{ \frac{R_i}{R_i+B_i}, \frac{B_i}{R_i+B_i} \right\}$, where

$$R_i = \sum_{\substack{j:(j,i) \in E \\ j \neq i}} r_j + \sum_{\substack{j:(i,j) \in E \\ j \neq i}} r_j, \quad (1)$$

$$B_i = \sum_{\substack{j:(j,i) \in E \\ j \neq i}} b_j + \sum_{\substack{j:(i,j) \in E \\ j \neq i}} b_j. \quad (2)$$

This means nodes in balanced neighborhoods which have a greater w_i will be considered more important.

We show that DIVERSE CENTRALITY captures all three criteria – centrality, local polarity, and neighborhood polarity – concurrently, while PAGERANK, REWEIGHT NODE BIAS and REWEIGHT NEIGHBOR BIAS cannot. Section 5.2 analyzes the effects of fixing neighborhood polarity and centrality while changing local polarity. Section 5.3 analyzes the effects of fixing local polarity and centrality while changing neighborhood polarity. Section 5.4 analyzes the effects of fixing local and neighborhood polarities while changing centrality.

We further note that DIVERSE CENTRALITY differs from the baselines in a more fundamental way. As mentioned before, it propagates influence in the sense that it amplifies the score of nodes that are themselves pointed to by nodes that are central and have balanced polarity. This feature is not shared by simple modifications of PAGERANK based on node or neighborhood polarity.

5.2 Local Polarity

Consider a bidirectional graph CHANGE LOCAL POLARITY with 2000 nodes. The graph is first generated using the FULLY RANDOM model; then, 600 nodes are chosen uniformly at random and divided into three sets:

- V_1 with 150 nodes, each with polarity $r_i = 0.99, b_i = 0.01$.
- V_2 with 300 nodes, each with polarity $r_i = b_i = 0.5$.
- V_3 with 150 nodes, each with polarity $r_i = 0.01, b_i = 0.99$.

Since the FULLY RANDOM graph does not consider polarities during random generation, V_1, V_2 and V_3 will have close to balanced neighborhoods. The random selection of nodes ensures that centrality of these nodes will span the entire range of centrality values as measured by PAGERANK.

Figure 2 shows the average scores of nodes in V_1, V_2 and V_3 given by DIVERSE CENTRALITY. The spectrum of PAGERANK values is first divided into 7 buckets;¹ nodes within each bucket are further grouped by their local polarity (i.e. whether they are in V_1, V_2 or V_3), and for each given bucket, the average DIVERSE CENTRALITY scores of nodes from V_1, V_2 and V_3 are shown with the red, grey and blue bars respectively. Welch’s t -tests are performed within each bucket, comparing all polarized nodes against all balanced nodes.

For both DIVERSE CENTRALITY and REWEIGHT NEIGHBOR BIAS, we then compare the differences in scores between balanced nodes and polarized nodes in Figure 3. Within each of the 7 buckets, for each model, the average scores of balanced nodes minus the average scores of polarized nodes is plotted, with t -tests performed similarly to Figure 2.

¹The 7 buckets are obtained by first dividing the entire range of PAGERANK values evenly into 15 intervals, and then merging the first 5 and last 5 intervals which contain too few nodes individually.




Figure 2: Average DIVERSE CENTRALITY scores of nodes in CHANGE LOCAL POLARITY graph given range of PAGERANK values and local polarity (864 runs).

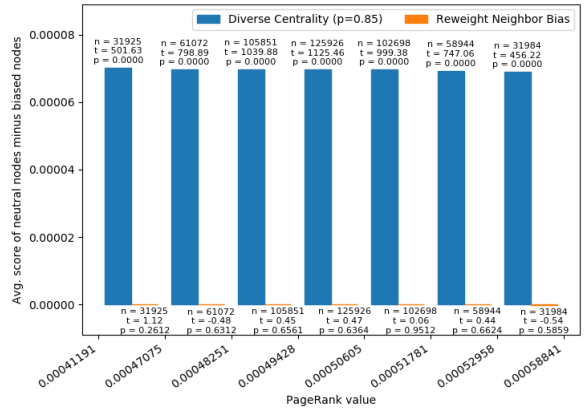


Figure 3: Average scores of balanced nodes minus polarized nodes in CHANGE LOCAL POLARITY graph, given range of PAGERANK values, for DIVERSE CENTRALITY and REWEIGHT NEIGHBOR BIAS models (864 runs).

Within each bucket of PAGERANK values, nodes have approximately constant neighborhood polarity and centrality, but different local polarity. Both Figures 2 and 3 show that the DIVERSE CENTRALITY model ranks balanced nodes in V_2 higher than polarized nodes in V_1 and V_3 , and the difference is statistically significant in all 7 buckets. Conversely, the REWEIGHT NEIGHBOR BIAS model gives nodes in each bucket approximately the same scores regardless of their local polarity, with none of the differences being significant.

This shows that our DIVERSE CENTRALITY model reflects the differences in local polarities of nodes even when centrality and neighborhood polarity are held constant, while the REWEIGHT NEIGHBOR BIAS model fails to capture local polarity.

5.3 Neighborhood Polarity

Consider a bidirectional graph CHANGE NEIGHBORHOOD POLARITY with 2000 nodes. The graph is first generated using the POLARITY




Figure 4: Average DIVERSE CENTRALITY scores of nodes in CHANGE NEIGHBORHOOD POLARITY graph given range of PAGERANK values and neighborhood polarity (600 runs).

ATTACHMENT model to allow polarized clusters to form. Then, a set of 600 nodes V_0 are chosen uniformly at random, and their polarities are changed to balanced with $r_i = b_i = 0.5$.

Although all nodes in V_0 have the same local polarity, their neighborhood polarity differs since neighborhood polarity in the POLARITY ATTACHMENT graph strongly correlates with their original polarities that were generated randomly.

Figure 4 shows the average scores of nodes in V_0 given by DIVERSE CENTRALITY. Similar to Section 5.2, the spectrum of PAGERANK values is divided into 7 buckets. Nodes within each bucket are further grouped by their neighborhood polarity: node i with $0.45 \leq R_i / (R_i + B_i) \leq 0.55$ is considered to be in a balanced neighborhood, where R_i and B_i are defined in Eq (1) and (2); otherwise, it is considered to be in a polarized neighborhood. Empirically, around 42% of nodes in each bin are classified as in balanced neighborhoods.

For each given bucket, the average DIVERSE CENTRALITY scores of nodes in red, balanced and blue neighborhoods (over 600 runs) are shown with the corresponding bars. Welch’s t -tests are performed within each bucket, comparing all nodes in balanced neighborhoods against nodes in polarized neighborhoods.

Within each bucket of PAGERANK values, nodes have the same local polarity and approximately constant centrality, but different neighborhood polarity. Figures 4 show that DIVERSE CENTRALITY ranks nodes in balanced neighborhoods higher than nodes in polarized neighborhoods, with statistical significance in all buckets.

For both DIVERSE CENTRALITY and REWEIGHT NODE BIAS models, we compare the differences in scores between nodes in balanced and polarized neighborhoods, similarly to Figure 3. Within all 7 buckets, the average DIVERSE CENTRALITY scores of nodes in balanced neighborhoods is significantly higher than nodes in polarized neighborhoods, with differences ranging from 6.890×10^{-5} to 7.011×10^{-5} . On the other hand, nodes in balanced neighborhoods typically have lower REWEIGHT NODE BIAS scores than those in polarized neighborhoods, with a much smaller magnitude of difference from 2.887×10^{-9} to 1.192×10^{-7} .

These results show that our DIVERSE CENTRALITY model captures the differences in neighborhood polarities of nodes even when

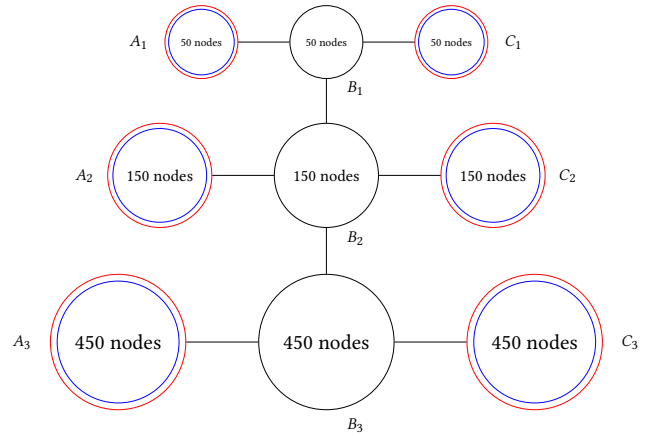


Figure 5: 9 CLUSTERS graph.




Figure 6: Average scores of B_1 , B_2 and B_3 in 9 CLUSTERS graph from PAGERANK and DIVERSE CENTRALITY (300 runs).

centrality and local polarity are held constant, and gives balanced neighborhoods a higher score than polarized neighborhoods. DIVERSE CENTRALITY thus captures our neighborhood polarity criterion, but REWEIGHT NODE BIAS does not.

5.4 Centrality

Given that our DIVERSE CENTRALITY algorithm has a similar formulation to standard PAGERANK, one would naturally expect DIVERSE CENTRALITY scores to correlate with PAGERANK. Specifically, in the case where all nodes in the graph are balanced with $r_i = b_i = 0.5$, Definition 1 reduces to PAGERANK with damping factor p . However, we would like to analyze the relationship with centrality in more general, nontrivial cases where nodes may have varying polarities, and observe the effects of changing centrality on a subset of nodes whose local and neighborhood polarities are fixed.

Consider a bidirectional graph 9 CLUSTERS whose general structure is shown in Figure 5. The graph contains nine clusters: nodes in the central clusters B_1 , B_2 and B_3 all have polarities $r_i = b_i = 0.5$, while nodes in other clusters have random polarities. Each pair of nodes within the same cluster have a probability of 0.5 of being

connected, while each pair of nodes across two clusters adjacent in Figure 5 have a probability of 0.1 of being connected.

All nodes in B_1 , B_2 and B_3 are balanced, and have largely balanced neighborhoods due to randomness.² However, the three clusters have different centrality with B_3 being the most central and B_1 the least.

Figure 6 shows the average scores given to each of B_1 , B_2 , B_3 by PAGERANK and DIVERSE CENTRALITY. The PAGERANK scores confirm the difference in centrality between the three clusters. More importantly, DIVERSE CENTRALITY also shows an increase in scores from B_1 to B_3 , similar to PAGERANK. This shows that DIVERSE CENTRALITY reflects the differences in centrality of nodes when both local and neighborhood polarities are held largely constant, giving important nodes a higher score.

6 REAL-WORLD GRAPHS

In this section we apply the DIVERSE CENTRALITY to some graphs generated from real-world networks, to show its performance on these datasets and compare the results with conventional wisdom regarding the networks. We choose two real-world graphs where centrality and polarization are relevant: the network of members of the United States Congress and a network of political blogs.

Before presenting our results, we first present a natural generalization of a different and equally classic centrality notion, BETWEENNESS CENTRALITY [14], to capture diversity. Subsequently, we will evaluate the performance of DIVERSE CENTRALITY using DIVERSE BETWEENNESS CENTRALITY along with PAGERANK as benchmarks.

6.1 DIVERSE BETWEENNESS CENTRALITY

The BETWEENNESS CENTRALITY of a vertex v quantifies the frequency it appears on the shortest path of two other vertices. Formally,

$$BC(v) := \sum_{s \neq v \neq t} \delta_{st}(v)$$

where $\delta_{st}(v)$ is the fraction of s - t shortest paths passing through v .

We naturally generalize it into DIVERSE BETWEENNESS CENTRALITY to incorporate diversity considerations. We define

$$DBC(v) := \sum_{s \neq v \neq t} \delta_{st}(v) |r_s - r_t|$$

to measure how often it appears on the shortest path of two other vertices, while giving larger weights if those two vertices are from different communities.

DIVERSE BETWEENNESS CENTRALITY can be computed in $O(nm)$ time on unweighted graphs, where n and m are the numbers of nodes and edges respectively, using a simple modification from Brandes' algorithm [6]. In practice, it is typically slower than DIVERSE CENTRALITY which takes $O(m)$ per iteration.

6.2 Congress Graph

We create a graph that models all members of the 116th United States Congress (2019–20) linked by their connections on Twitter. Each Congress member with an official Twitter account is represented

²Empirically, over 300 runs, $R_i / (R_i + B_i)$ values of nodes in B_1 , B_2 and B_3 have a mean of 0.499987 and standard deviation of 0.007689.

as a node. A directed edge from member X to member Y exists if X was following Y on Twitter as of May 30, 2020.³

There are $K = 2$ communities, and for node i , the affiliation vector $\mathbf{q}^{(i)} = [b_i, r_i]$ measures how much the Congress member's ideologies align with the Democrats and Republicans respectively. To generate this vector, we use the DW-NOMINATE procedure that provides a quantitative measure of their ideologies based on factors such as voting behavior [32]. For each member, this method produces a number between -1 (most liberal) and 1 (most conservative), which we scale linearly to fit our notion in Section 2.

When comparing DIVERSE CENTRALITY with PAGERANK and DIVERSE BETWEENNESS CENTRALITY on the Congress Graph, as shown in Figure 7, we observe a *bridging effect*:

The DIVERSE CENTRALITY model favors nodes that are not only bipartisan and influential (in terms of PAGERANK), but that also serve as a bridge to similarly bipartisan and influential nodes in other communities.

Specifically, while the graph in Figure 7 is clearly separated into two polarized clusters, each of them contain a sub-cluster shown on the bottom left and bottom right. The sub-clusters highly correlate to the Senate, with the main cluster largely being the House. In general, senators from both parties are well-connected to each other; each senator is connected to a large number of representatives from their own party and a smaller number from the other party; edges directly between representatives of different parties exist, but are less common. This means the Senate sub-clusters serve as bridging nodes between the two clusters, since two representatives from different parties are typically connected via one or more senators.

While both Figures 7a and 7b show the Senate sub-clusters ranked above average, the difference between the Senate and the House shown by DIVERSE CENTRALITY in Figure 7a is much more pronounced, as both sub-clusters receive much higher scores here. Figure 7c, on the other hand, shows DIVERSE BETWEENNESS CENTRALITY ranks the Senate similarly to the House. Instead, the measure appears to favor nodes with high outdegrees, as the highly ranked members follow most or all other members of Congress on Twitter. While these nodes can serve as a bridge across the two communities, they are not necessarily bipartisan and influential themselves.

Senators			
Node	DC	PR	DBC
Susan M. Collins	1	10	152
Lisa Murkowski	6	20	224
Kyrsten Sinema	9	23	2
Joe Manchin, III	5	64	198
Rob Portman	11	8	30

Table 1: Top 5 senators joining greatest proportion of bipartisan bills, with their ranks by DIVERSE CENTRALITY, PAGERANK, and DIVERSE BETWEENNESS CENTRALITY respectively.

³Two members in the graph did not follow any other members. In order to prevent their PAGERANK values from being artificially inflated, edges are added from each of the two sinks to all nodes in the graph.

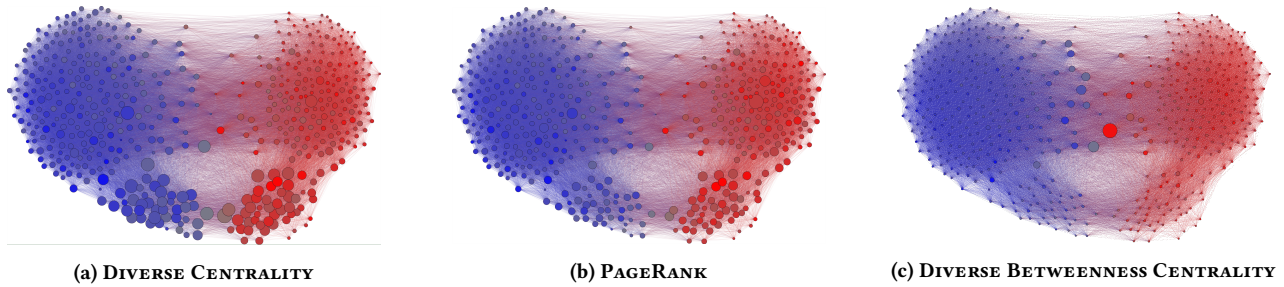


Figure 7: Congress Graph. Layout generated by the ForceAtlas2 algorithm [20]. Nodes are colored based on polarities ($q^{(i)}$).




Figure 8: Number of cut edges among top k nodes in Congress Graph ranked by DIVERSE CENTRALITY, PAGERANK, and DIVERSE BETWEENNESS CENTRALITY scores.




Figure 9: Number of cut edges between top k nodes in Political Blogs Graph ranked by DIVERSE CENTRALITY, PAGERANK, and DIVERSE BETWEENNESS CENTRALITY scores.

To verify that the DIVERSE CENTRALITY indeed finds the most important bipartisan members of the Congress, we consider the top five senators with the highest percentage of bills cosponsored that were introduced by the other party, an independent metric that serves as a proxy of bipartisanship of members [1]. As shown in Table 1, 4 out of these 5 senators are ranked significantly higher by DIVERSE CENTRALITY (DC) than either PAGERANK (PR) or DIVERSE BETWEENNESS CENTRALITY (DBC). Furthermore, these five senators are among top 11 in the entire graph ranked by DIVERSE CENTRALITY; This indeed shows that DIVERSE CENTRALITY prioritizes bipartisan and influential nodes in the graph.

We can quantify this bridging effect by analyzing intersections of top nodes and the cut between two clusters. Spectral clustering is first performed on the graph (converted to an undirected graph) to identify the two polarized clusters [22]. Then, for each ranking algorithm and for various values of k , we count the number of edges in the subgraph induced by the top k nodes that are also in the cut. Figure 8 shows the results for all these centrality models.

The DIVERSE CENTRALITY model has a greater number of cut edges for any given value of k . This shows that DIVERSE CENTRALITY has a greater focus than PAGERANK on bridging nodes that

connect the two polarized clusters, instead of central nodes within each community which may not be well connected to each other. DIVERSE CENTRALITY beats DIVERSE BETWEENNESS CENTRALITY in this metric too, even though the latter emphasizes this bridging effect in its definition.

6.3 Political Blogs Graph

We also ran the algorithms on a graph that models blogs on U.S. politics in 2005 [2]. The dataset is a directed graph that forms edges among blogs based on hyperlinks.⁴

We again assume $K = 2$ communities in the graph. Since the dataset does not contain fine-grained polarity data of the blogs, for our adaptation, all liberal blogs are assigned $r_i = 0.01, b_i = 0.99$, and all conservative blogs have $r_i = 0.99, b_i = 0.01$. These values are approximately discrete, denoting community memberships rather than extent of affiliation to each community.

We can again analyze the number of cut edges among top nodes to quantify this bridging effect, as shown in Figure 9. For any given value of k , DIVERSE CENTRALITY gives the greatest number of cut

⁴Four nodes are not connected the rest of the graph, so we removed them and other disconnected nodes to ensure spectral clustering finds a good cut.

edges across the two clusters than PAGERANK, similar to the Congress Graph. Therefore, on this graph, DIVERSE CENTRALITY still favors nodes on the frontier of the two clusters bridging them together. This example also shows that DIVERSE CENTRALITY still performs well when nodes have extreme affiliation vectors (in the sense of being nearly 0 or 1).

7 CONCLUSION

In this paper, we defined a new measure of centrality, DIVERSE CENTRALITY, that gives importance to nodes that bridge different communities. We propose several elementary criteria for such a centrality measure with diversity and show via simulations on realistic random networks that DIVERSE CENTRALITY satisfies all the proposed criteria, while simple modifications to standard PAGERANK do not. We further show how important bipartisan senators are naturally unearthed by our measure when run on a social network constructed from the United States Congress.

Our work leads to several open questions. First, we do not have a complete theoretical understanding of DIVERSE CENTRALITY. Is the fixed point unique under reasonable assumptions on the network? Can it be proven that our iterative algorithm for computing the fixed point always converges? Further, in several real-world settings, there could be several different types of communities, corresponding for instance to politics, demographics, religion, location, etc. An individual could belong to one or more communities for each of these dimensions. How can we generalize the notion of DIVERSE CENTRALITY to handle this setting, and how does it impact the performance of the resulting algorithms on real data?

ACKNOWLEDGMENTS

This work is supported by NSF grants CCF-1637397, ONR award N00014-19-1-2268; and DARPA award FA8650-18-C-7880.

REFERENCES

- [1] 2019. <https://www.govtrack.us/congress/members/report-cards/2019/senate/cosponsored-other-party>
- [2] Lada A. Adamic and Natalie Glance. 2005. The political blogosphere and the 2004 U.S. election. In *LinkKDD '05: Proceedings of the 3rd international workshop on Link discovery (KDD '05)*. 36–43.
- [3] Réka Albert and Albert-László Barabási. 2002. Statistical mechanics of complex networks. *Rev. Mod. Phys.* 74 (Jan 2002), 47–97. Issue 1.
- [4] J. Angwin, J. Larson, S. Mattu, and L. Kirchner. 2016. Machine bias: There's software used across the country to predict future criminals. And it's biased against blacks. *ProPublica* (May 2016). <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- [5] Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefevre. 2008. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* 2008, 10 (oct 2008).
- [6] Ulrik Brandes. 2001. A faster algorithm for betweenness centrality. *The Journal of Mathematical Sociology* 25, 2 (2001), 163–177. <https://doi.org/10.1080/0022250X.2001.9990249> arXiv:<https://doi.org/10.1080/0022250X.2001.9990249>
- [7] E. Budish. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *J. Political Economy* 119, 6 (2011), 1061–1103.
- [8] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Trans. Econ. Comput.* 7, 3, Article 12 (Sept. 2019), 32 pages. <https://doi.org/10.1145/3355902>
- [9] Raffaele Chiappinelli. 2018. What Do You Mean by “Nonlinear Eigenvalue Problems”? *Axioms* 7 (06 2018), 39.
- [10] M. H. DeGroot. 1974. Reaching a consensus. *J. American Statistical Association* 69 (1974), 118–121.
- [11] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard S. Zemel. 2012. Fairness through awareness. In *Innovations in Theoretical Computer Science 2012, Cambridge, MA, USA, January 8-10, 2012*. 214–226.
- [12] P. Erdős and A. Rényi. 1959. On Random Graphs I. *Publicationes Mathematicae Debrecen* 6 (1959), 290.
- [13] M. G. Everett and S. P. Borgatti. 1999. The centrality of groups and classes. *The Journal of Mathematical Sociology* 23, 3 (1999), 181–201.
- [14] Linton C Freeman. 1977. A set of measures of centrality based on betweenness. *Sociometry* (1977), 35–41.
- [15] N. E. Friedkin and E. C. Johnsen. 1990. Social Influence and Opinions. *J. Mathematical Sociology* 15, 3-4 (1990), 193–205.
- [16] Tatsunori Hashimoto, Megha Srivastava, Hongseok Namkoong, and Percy Liang. 2018. Fairness Without Demographics in Repeated Loss Minimization. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*. 1929–1938.
- [17] Taher H. Haveliwala. 2002. Topic-Sensitive PageRank. In *Proceedings of the 11th International Conference on World Wide Web (Honolulu, Hawaii, USA) (WWW '02)*. Association for Computing Machinery, New York, NY, USA, 517–526.
- [18] R. Hegselmann and U. Krause. 2002. Opinion dynamics and bounded confidence models, analysis, and simulation. *Artificial Societies and Social Simulations* 5, 3 (2002).
- [19] J. E. Hirsch. 2005. An index to quantify an individual's scientific research output. *Proceedings of the National Academy of Sciences* 102, 46 (2005), 16569–16572.
- [20] Mathieu Jacomy, Tommaso Venturini, Sebastien Heymann, and Mathieu Bastian. 2014. ForceAtlas2, a Continuous Graph Layout Algorithm for Handy Network Visualization Designed for the Gephi Software. *PLOS ONE* 9, 6 (06 2014), 1–12.
- [21] Glen Jeh and Jennifer Widom. 2003. Scaling Personalized Web Search. In *Proceedings of the 12th International Conference on World Wide Web (Budapest, Hungary) (WWW '03)*. Association for Computing Machinery, New York, NY, USA, 271–279.
- [22] Jianbo Shi and J. Malik. 2000. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22, 8 (2000), 888–905.
- [23] Xue-Mei Jiang, Gui-Rong Xue, Wen-Guan Song, Hua-Jun Zeng, Zheng Chen, and Wei-Ying Ma. 2004. Exploiting PageRank at Different Block Level. In *Web Information Systems – WISE 2004*, Xiaofang Zhou, Stanley Su, Mike P. Papazoglou, Maria Elzbieta Orłowska, and Keith Jeffery (Eds.). 241–252.
- [24] Sepandar D. Kamvar, Taher H. Haveliwala, Christopher D. Manning, and Gene H. Golub. 2003. Extrapolation Methods for Accelerating PageRank Computations. In *Proceedings of the 12th International Conference on World Wide Web (WWW '03)*. Association for Computing Machinery, 261–270.
- [25] Przemyslaw Kazienko and Tomasz Kajdanowicz. 2012. Label-dependent node classification in the network. *Neurocomputing* 75 (01 2012), 199–209.
- [26] Jon M. Kleinberg, Sendhil Mullainathan, and Manish Raghavan. 2016. Inherent Trade-Offs in the Fair Determination of Risk Scores. *CoRR* abs/1609.05807 (2016). <http://arxiv.org/abs/1609.05807>
- [27] Isabel M. Kloumann, Johan Ugander, and Jon Kleinberg. 2017. Block models and personalized PageRank. *Proceedings of the National Academy of Sciences* 114, 1 (2017), 33–38.
- [28] Silvio Lattanzi and D. Sivakumar. 2009. Affiliation Networks. In *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing (STOC '09)*. Association for Computing Machinery, New York, NY, USA, 427–434.
- [29] Jure Leskovec, Kevin J. Lang, and Michael Mahoney. 2010. Empirical Comparison of Algorithms for Network Community Detection. In *Proceedings of the 19th International Conference on World Wide Web (WWW '10)*. 631–640.
- [30] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. *The PageRank Citation Ranking: Bringing Order to the Web*. Technical Report 1999-66. Stanford InfoLab. <http://ilpubs.stanford.edu:8090/422/> Previous number = SIDL-WP-1999-0120.
- [31] Geoff Pleiss, Manish Raghavan, Felix Wu, Jon Kleinberg, and Kilian Q Weinberger. 2017. On Fairness and Calibration. In *Advances in Neural Information Processing Systems (NIPS)*. 5680–5689.
- [32] Keith T. Poole and Howard Rosenthal. 1985. A Spatial Model for Legislative Roll Call Analysis. *American Journal of Political Science* 29, 2 (1985), 357–384.
- [33] H. R. Varian. 1976. Two problems in the theory of fairness. *Journal of Public Economics* 5, 3-4 (1976), 249–260.
- [34] G. Weisbuch, G. Deffuant, F. Amblard, and J-P. Nadal. 2002. Meet, discuss and Segregate! *Complexity* 7, 3 (2002), 55–63.
- [35] Jierui Xie, Stephen Kelley, and Boleslaw K. Szymanski. 2013. Overlapping Community Detection in Networks: The State-of-the-Art and Comparative Study. *ACM Comput. Surv.* 45, 4 (Aug. 2013).
- [36] J. Yang, J. McAuley, and J. Leskovec. 2013. Community Detection in Networks with Node Attributes. In *Proc. IEEE Intl. Conf. Data Mining*. 1151–1156.