

Decentralized Core-periphery Structure in Social Networks Accelerates Cultural Innovation in Agent-based Modeling

Jesse Milzman

DEVCOM Army Research Laboratory
Adelphi, MD, United States
jesse.m.milzman.civ@army.mil

Cody Moser

University of California Merced
Merced, CA, United States
cmoser2@ucmerced.edu

ABSTRACT

Previous investigations into creative and innovation networks have suggested that intellectual, artistic, and technological innovations often occur with the interaction of core and peripheral actors, with both individuals and teams at intermediary positions well-situated to innovate. In this work, we investigate the effect of global core-periphery network structure on the speed and quality of cultural innovation. Drawing on differing notions of core-periphery structure from [22] and [2], we distinguish decentralized core-periphery, centralized core-periphery, and affinity network structure. We generate networks of these three classes from stochastic block models (SBMs), and use them to run an agent-based model (ABM) of collective cultural innovation, in which agents can only directly interact with their network neighbors. In order to discover the highest-scoring innovation, agents must discover and combine the highest innovations from two completely parallel technology trees. We find that decentralized core-periphery networks outperform both centralized core-periphery networks and affinity networks, in terms of mean crossover time for this final innovation. We hypothesize that decentralized core-periphery structure provides a more fruitful environment for collective problem-solving, by allowing for the relative shielding of periphery nodes from the optimal innovations known by the core community at any given time. This prevents the disincentive for parallel explorations that emerges in a highly connected, centralized network. We then build upon the "Two Truths" hypothesis regarding community structure in spectral graph embeddings first articulated in [22]. The Two Truths hypothesis suggests that the adjacency spectral embedding (ASE) captures core-periphery community structure in a graph, while the Laplacian spectral embedding (LSE) captures affinity structure. For a given network generated from an SBM, we use ASE and LSE to resample new networks of similar structure, using either to parameterize a random dot product graph (RDPG) model. We find that, for core-periphery networks, ASE resampling best recreates networks with similar performance on the innovation SBM. Since the Two Truths hypothesis suggests that ASE captures core-periphery structure, this result further supports our hypothesis.

KEYWORDS

agent-based modeling; core-periphery networks; innovation networks, network science; adjacency spectral embedding; stochastic block models; cultural evolution; collective problem-solving

ACM Reference Format:

Jesse Milzman and Cody Moser. 2023. Decentralized Core-periphery Structure in Social Networks Accelerates Cultural Innovation in Agent-based Modeling. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 9 pages.

1 INTRODUCTION

From a broad perspective, innovation is understood as a form of collective problem-solving. For this and other reasons, the process of innovation is understood as a social process, as social collectives are capable of, and in some cases optimized for, retaining the knowledge of previous generations while building upon this knowledge for subsequent innovations, a phenomenon we refer to as "cumulative" culture [18, 19]. Human social networks tend to exhibit core-periphery structures, whereby a 'core' population is heavily inter-connected, and connected in turn to more 'peripheral' individuals and subcommunities [2]. Prior work on the structure of human networks has suggested that innovation emerges at the boundary between the core and periphery of creative networks [4, 6]. Individual innovators are often in an intermediate position with many core and peripheral connections, and successfully innovative teams tend to include both core and peripheral individuals [4].

However, to our knowledge, there has not been an in-depth examination of the relationship between the *global* core-periphery structure of a creative population and its innovation potential. Moreover, there are multiple candidate metrics purporting to quantify core-periphery structure within a given network, suggesting that the concept is ill-posed [23]. Most definitions of core-periphery structure, formal and informal, assume that core nodes ought to be heavily connected to peripheral nodes [2, 23]. Nonetheless, at least one recent work suggests a definition that precludes this feature [22], as it is not observed in the core-periphery structure observed in connectome networks [22], corresponding to the grey/white matter divide. The authors of [22] instead define core-periphery networks such that peripheral nodes are only loosely connected to either core or periphery, with comparable probabilities for the edges to either type. Thus, even in the simplest setting in which one ignores the possibility of multiple cores or intermediate layers, there is a conceptual distinction that must be made between the desired expected relationship between core and peripheral nodes.

We therefore utilize an agent-based model (ABM) of collective problem-solving known as the "Potions Game" in order to investigate the effects of core-periphery structure on the ability of a population of agents to locate optimal solutions in an innovation task [10]. In the model, agents are given a set of ingredients which can be combined with the ingredients of one of their neighbors to

produce new ones. These novelties are shared with their neighbors and can be combined with the previous set of ingredients to produce additional novelties. As new ingredients are discovered, agents are more likely to use them over old ones, but unbeknownst to the agents, the initial set of ingredients can be combined to produce new ingredients along two separate pathways. Because agents are unlikely to return to the original set of ingredients, this creates path dependency in the model, but if agents are able to find and exploit both pathways, they can combine their final ingredient to discover an ingredient which successfully ends the simulation.

By running our ABM on core-periphery networks, we may explore the distinctions between network structures along several dimensions, and their implications for collective problem-solving. As already mentioned, [2] and [22] suggest distinct notions of core-periphery structure, reflecting the degree to which the network core dominates the peripheral neighborhood. Thus, we split core-periphery structure into the two types suggested: one with a periphery which is not very strongly connected to itself, but is strongly connected to the core (a centralized periphery); and one with a periphery which is isolated both from itself and from the core (a decentralized periphery). Moreover, core-periphery networks in [22] are contrasted to “affinity” networks, in which both communities have more within than between similarity.

Finally, we discover a surprising crossover relationship between decentralized core-periphery structure, spectral graph embeddings, and performance on our innovation ABM. The adjacency spectral embedding (ASE) [26] and the Laplacian spectral embedding (LSE) [31] are common methods to embed graphs in Euclidean space for the purpose of community detection [15, 26, 31]. However, despite statistical results suggesting both are consistent estimators for theoretical models, they do not always recover the same communities. Priebe *et al.* observed that ASE seems to recover core-periphery communities, while LSE recovers affinity communities [22]. Keeping in mind this difference, we investigate the use of ASE and LSE to resample networks with a given structure (core-periphery or affinity), and compare the ABM performance of the resampled networks to the original.

2 RELATED WORK

Network scientists have collectively formulated a vast taxonomy of structural qualities to describe network topologies, including those of density/sparsity, efficiency [7], hierarchy [15, 17], affinity/modularity [9, 21], and core-periphery separation [2]. This last quality has been implicated in both individual and collective learning, via the analysis of creative and learning networks, respectively [1, 4, 5, 7].

One early attempt to formalize a model of core-periphery structure was that of Borgatti and Everett [2]. Grounding themselves in block modeling approaches common in social network science at the time, they identified ideal core-periphery networks as two-block graphs in which the nodes of one block (the ‘core’) are fully connected to every other node, while the nodes in the other block (‘periphery’) are connected only to the core nodes. They propose a statistical test for an *a priori* core-periphery partition, whose statistic is the Pearson coefficient between the entries of the adjacency matrix and those of the ideal block model for the same partition and

also propose an algorithm for assigning every node a continuous coreness value for weighted networks. Since [2], multiple works have developed new metrics [11, 23, 24, 29, 32], both for classifying nodes according to coreness, and detecting more complex structure such as multiple cores [12, 23] and multi-layered networks [11]. In this present effort, we will not be measuring coreness of nodes or quantifying core-periphery structure in given networks. Rather, we will generate networks at a desired level and type of core-periphery structure with an elementary random graph model and tie these networks to performance in our collective innovation task, as discussed in Sec. 3.1.

Ongoing work in the field of collective problem-solving has worked to identify network structures which optimize for the trade-off between the exploration of new solutions and the exploitation of current solutions [7]. Emerging from these studies is an understanding on the role that informational efficiency in the form of simple path lengths plays in solving both simple and complex problems; with a core trade-off being that for simple problems, efficient or highly centralized graphs are desirable, while for more complex problems, inefficient and less clustered graphs are more desirable [3, 13]. Understood in these dimensions, networks which are efficient are said to be fast, with the ability to drive rapid consensus and “exploit” the information available to them while networks which are inefficient are said to be slower and more “transiently diverse”, with the ability to engage in parallel exploration of different opportunities [7, 25]. Thus, networks in these tasks must work to strike the appropriate balance between optimal levels of exploration of the solution space and exploitation, *apropos* their task [16].

Many of the formal elements of our network analysis, including the generation of random networks from stochastic block models, the parameter distinction between core-periphery and affinity networks, and our spectral graph embedding analysis, draw from the work in [22]. The authors compare core-periphery and affinity network structures from a spectral clustering perspective, in order to explain the discrepancy between the communities assigned by ASE and LSE-based community detection methods. They describe a “Two Truths” phenomenon, in which ASE preferentially reveals and clusters core-periphery communities, while LSE does the same for affinity communities. In particularly striking results, they applied ASE and LSE clustering to a diffusion MRI connectome dataset, demonstrating that ASE captured grey/white matter core-periphery structure, while LSE captured left/right hemisphere affinity structure. To our knowledge, the implications of this Two Truths hypothesis for spectral graph resampling [14] have not been explored. There is also little work using spectral graph embeddings to analyze agent-based models, although they are commonplace in social network analysis.

3 PRELIMINARIES

In this section, we introduce the formal background and technical methods of our investigation. We use the terms ‘network’ and ‘graph’ interchangeably, in both cases always referring to undirected, unweighted graphs without loops. These are pairs $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} is the set of vertices or nodes, and \mathcal{E} is the set of edges. We will always identify vertices with their indices,

i.e. if $|\mathcal{V}| = N$, $\mathcal{V} = [N]$ where $[N] = \{1, \dots, N\}$. An edge is a couple $\{i, j\} \in \mathcal{E}$ for some $i \neq j \in \mathcal{V}$.

3.1 Stochastic Block Models of Affinity and Core-Periphery Networks

Similar to the influential work of Borgatti and Everett [2], we understand core-periphery structure as a network approximating a two-block model, with a heavily connected core and a less connected periphery. We have opted to follow the loosely formalized probabilistic interpretation of core-periphery structure from Priebe *et al.* [22], in which they present core-periphery structure in contradistinction to affinity community structure under certain parameter regimes for 2-block stochastic block models (SBMs). We expand this framework to include networks exhibiting the centralized core-periphery structure typical in the social networks literature, first formalized by [2]. As we will see, in this setting, the notions of core-periphery from [2] and [22] are in some sense mutually exclusive, although they both assume the existence of a ‘rich club’ core.

Stochastic block models (SBMs) are a popular model of random graphs (graph-valued random variables) with community structure, and are frequently used as the theoretical graph model underlying community detection algorithms within the network statistics literature, e.g. [15, 26]. They are defined as follows. Consider a collection of N nodes divided into $K \geq 2$ blocks (‘communities’), where each block k is of size n_k , so that $\sum n_k = N$. Let $\pi \in [K]^N$ be the block membership vector. The block matrix \mathbf{B} is a $K \times K$ symmetric matrix of probabilities that parametrizes the SBM (along with π), where $b_{k_1 k_2}$ provides the probability of an edge between a node in community k_1 and one in k_2 . This in turn forms an $N \times N$ edge probability matrix \mathbf{P} , composed of blocks $P_{k_1 k_2} = b_{k_1 k_2} \mathbb{1}_{k_1 \times k_2}$. The SBM is the random graph in which every edge (i, j) is sampled independently, with probability p_{ij} . We denote an SBM as a distribution on adjacency matrices: $A \sim \text{SBM}(\mathbf{B}, \pi)$ Such a P -parametrized random graph model, more generally, is typically referred to as an inhomogeneous Erdos-Renyi graph, denoted distributionally as $A \sim \text{ER}(P)$.

Priebe *et al.* [22] suggest a model for core-periphery networks as a 2-block SBM. A two-block SBM is parametrized by the block matrix:

$$\mathbf{B} = \begin{bmatrix} a & b \\ b & c \end{bmatrix} \quad (1)$$

They distinguish two generic cases of such SBMs. When $a, b \gg c$, the authors identify the SBM as exhibiting **affinity** structure, in which realizations of the network can usually be partitioned into two communities exhibiting dense intra-communal connectivity, with sparser connections between the two. In this case, neither community can be said to necessarily significantly dominate the other, especially if $a \approx c$ and $n_1 \approx n_2$. By contrast, the authors identified the case in which $a \gg b, c$ as exhibiting **core-periphery** structure. In this case, the first community exhibits dense intra-communal connectivity. By contrast, it is only weakly connected to the second community of peripheral nodes, which are sparsely connected among themselves. In this model of core-periphery structure, peripheral nodes, individually and communally, are isolated from the core as much as each other.

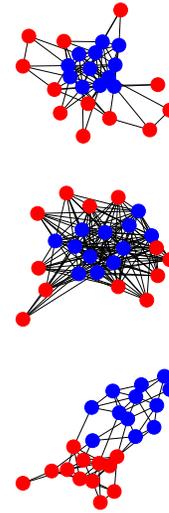


Figure 1: 2-Block Stochastic Block Model Networks. We generated three networks from model (1), with parameter sets corresponding to the three structural classes that we will be studying. All networks have $N = 24$ nodes, divided evenly between the blocks ($n_1 = n_2 = 12$). Blue nodes belong to the first block, while red nodes belong to the second. (A) A decentralized core-periphery network, generated with parameters $(a, b, c) = (0.8, 0.15, 0.05)$. As can be seen, the peripheral nodes have fewer neighbors than the core nodes, which include both core nodes and other peripheral nodes. (B) A centralized core-periphery network, generated with parameters $(a, b, c) = (0.8, 0.8, 0.05)$. (C) An affinity network, generated with parameters $(a, b, c) = (0.5, 0.5, 0.05)$.

This understanding of core-periphery structure from [22] is similar to that of [2] in that it understands core-periphery networks through a two-block model, composed of a dense core and sparse periphery. However, there is a major conceptual distinction in the two works that is deeper than the difference in mathematical framing. The approach in [2] identifies an ideal core-periphery network as one in which each peripheral node is connected to every core node. If we translated this into the more general probabilistic model (1), this would be the third case in which $a, b \gg c$, which was not considered in [22]. In the case when $a, b \gg c$, peripheral nodes will be heavily connected to the core, while isolated from each other. Within the context of our model, this would make it far more likely that peripheral nodes interact with and are influenced by the network core. By contrast, when $a \gg b, c$, the peripheral nodes will be relatively isolated from both communities as such, with only a handful of neighbors.

Thus, we distinguish these two forms of core-periphery structure, referring to that from [22] as **decentralized core-periphery (DCP)** structure, and that from [2] as **centralized core-periphery (CCP)** structure. We refer to ‘centralization’ to imply a notion of control on the part of the core, which is less evident in the core-periphery structure described by [22] than that in [2]. To summarize, there are three qualitatively distinct parameter regimes for (1) that we will be examining in this work:

- **Affinity:** $a, b \gg c$. In affinity networks, there are two distinguishable block communities with more intra-block edges than inter-block edges. The level of intra-block connectivity is comparable between the two communities.
- **Decentralized Core-Periphery (DCP):** $a \gg b, c$. There is a distinguishable core community, with dense intra-core connectivity. The rest of the nodes are peripheral, with only a few connections to other nodes, core or peripheral.
- **Centralized Core-Periphery (CCP):** $a, b \gg c$. There is a distinguishable core community, with dense intra-core connectivity *and* dense connectivity to peripheral nodes. Peripheral nodes are heavily connected to the core, and only loosely connected to each other.

We hypothesize that in our ABM, DCP networks will have a performance advantage compared to affinity networks, as peripheral nodes in CCP networks are able to freely “explore” possible alternative combinations in the Potions Game while the core exploits whichever innovation pathway they are on. Additionally, few nodes in an affinity network will be sufficiently isolated to explore alternatives with high probability. We similarly predict that in CCP structures, peripheral nodes conform more to the behavior of the core and therefore limit diversity in the network, making the path-dependent nature of the task difficult to overcome, leading to an advantage of DCP networks over CCP networks.

Hypothesis 1: Networks exhibiting decentralized core-periphery structure will outperform those with either affinity or centralized core-periphery structure in collective innovation tasks.

For all networks, we will set the two blocks to the same size ($n = n_1 = n_2 = N/2$). When comparing different parameter regimes for model (1), we want to maintain the same number of expected edges. The relative density/sparsity of networks has been implicated [7, 13] in collective problem-solving tasks, and we want to distinguish this effect from that of our competing core-periphery structures. For a given 2-SBM of size $N = 2n$, the expected number of edges is given by

$$\mathbb{E}|\mathcal{E}| = \binom{n}{2}(a + b + c) \quad (2)$$

Thus, to fix the total number of edges in the network, we need only fix the sum $a + b + c$. Where possible, we compare networks generated by a linear family of SBMs:

$$(a, b, c)(\theta) = (a_0, b_0, c_0) + \theta(\delta a, \delta b, \delta c), \quad 0 \leq \theta \leq \theta_{\max} \quad (3)$$

where we require $\delta a + \delta b + \delta c = 0$. The only exception to this will be the family (M3), where we must allow for greater network density in order to make a DCP-CCP comparison without compromising on the density of the core, i.e. in order to raise b without lowering a .

3.2 Simulation of Innovation Networks with Agent-Based Modeling

We adapted an ABM of cumulative innovation from [3, 18, 20], derived from the potion-mixing game studied by Derex and Boyd in human players [10]. The only parameter varying between simulations is the structure of the agent network.

- (1) **Model initialization:** A network is created with its respective parameters. For each node of the graph, an agent is initialized with a score of zero and an inventory comprised of six items: three from an “A trajectory” and three from a “B trajectory.” Each item in the inventory is comprised of three parts: the name/level of the item (e.g., a1, a2, a3, b1, b2, b3) and a score which each initial item and items discovered thereafter carries for itself (with scores of 6, 8, and 10 for the three initial items in each trajectory).
- (2) **Dyad selection:** At each step, each agent chooses a partner they are connected to on the network with a random probability. As neighbors are simply chosen with a random probability, it is possible for a focal neighbor to select an individual which is already interacting with them (e.g., if a network is initialized with just two agents, the two agents will simply select each other).
- (3) **Item selection:** In the model, new items are formed by triad combinations of old items. As triad combinations are made between dyads of agents, the focal agent randomly selects whether it will be trading either one item or two items with their partner. The focal agent and its partner then cycle through their respective inventories, assigning probabilities to each item in the array. This is obtained by summing the innovation scores of each item and dividing individual scores by each sum (e.g., the initial inventory innovation scores of 6, 8, 10, 6, 8, 10 will yield respective probabilities of .125, .167, .208, .125, .167, .208).
- (4) **Item combination:** Agents and their partners then select the number of items previously assigned to them in the last step, based on items’ calculated probabilities and without replacement, and combine them. The combination is saved as a list and compared to lists of valid combinations copied directly from Derex and Boyd [10] (SI Appendix, section 1). If an invalid combination is made, nothing happens. If a valid combination is made, then the dyad adds a new innovation (with its own respective score) to their inventories.
- (5) **Innovation diffusion:** If a new innovation is added to the agents’ inventories, both agents then check the inventories of all of their partners and spread it to neighbors which do not already possess it.
- (6) **Scoring:** Scores are then obtained for each agent based on the tier of discovery an agent has obtained: with the first tier yielding a score of 48, second tier 109, third tier 188, and the fourth tier (which requires a crossover from the A and B trajectory) being 358. The maximum score of an item in an agent’s list is determined to be their overall score.
- (7) **End and Crossover:** The simulation ends either when it has reached 1000 steps or when the network has achieved a “crossover event,” whereby the final inventions in both trajectories are themselves combined, indicating the network has discovered and united both paths of exploration.

3.3 Spectral Graph Embeddings, Resampling, and the Two Truths Hypothesis

Central limit theorems exist suggesting that both ASE and LSE provide asymptotically consistent estimates of the latent positions

of random dot product graphs (RDPGs), of which SBMs are a special case [27, 28]. We want to ask whether these two forms of spectral embedding capture a meaningful distinction between core-periphery and affinity networks, especially with respect to their performance in our innovation task.

Priebe *et al.* put forward a two-truths hypothesis comparing the performance of ASE+GMM and LSE+GMM as community detection tools [22]. They hypothesize that the latent positions estimated by ASE tend to emphasize core-periphery community structure, grouping core and periphery nodes near each other in Euclidean space. By contrast, the LSE seems to emphasize affinity communities. In the context of their work, they were interested in community labels, with ASE and LSE used to discover different but equally true community partitions in connectome networks with known ground truths. However, as mentioned above, asymptotic results suggest that the embedded positions themselves estimate latent positions for RDPGs.

Thus, given a network, we may use ASE and LSE to parameterize an RDPG model, which in turn will allow us to resample new networks with similar structure to the original. From the Two Truths hypothesis, we expect the ASE-RDPG model to better capture DCP structure in the original network, while LSE-RDPG will reproduce affinity structure.

We start by defining the ASE and LSE. Let a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be given, with $N = |\mathcal{V}|$. The adjacency \mathbf{A} matrix is an $N \times N$ symmetric matrix with Boolean values, where $a_{ij} = 1$ iff $\{i, j\} \in \mathcal{E}$, with zeros along the diagonal. The normalized Laplacian matrix is given as $\mathbf{L} = \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$, where \mathbf{D} is the degree matrix, i.e. diagonal where $d_{ii} = |\{e \in \mathcal{E} \mid i \in e\}|$. Both \mathbf{A} and \mathbf{L} are positive semidefinite, and have exclusively nonnegative eigenvalues. For either matrix \mathbf{M} , we will have the eigen decomposition:

$$\mathbf{M} = \mathbf{U} \mathbf{S} \mathbf{U}^T \quad (4)$$

where \mathbf{S} is the diagonal matrix of the (ordered) eigenvalues of \mathbf{M} , and the columns of \mathbf{U} are the corresponding eigenvectors. For a given embedding dimension $d < \text{rank}(\mathbf{M})$, we may instead compute the rank k approximation:

$$\mathbf{M} \approx \mathbf{U}_d \mathbf{S}_d \mathbf{U}_d^T \quad (5)$$

where \mathbf{S}_d is the $d \times d$ diagonal matrix of the largest d eigenvalues of \mathbf{M} , and \mathbf{U}_d is the $N \times d$ matrix of corresponding orthonormal eigenvectors. There are reasons both methodological and computational for using a low-rank decomposition, which we do not discuss here [30]. These eigenvectors span a d -dimensional space into which we will embed our nodes. To choose the embedding dimension d , we use the standard profile likelihood approach from [33]. We then form the matrix $\hat{\mathbf{X}} = \mathbf{U}_d \sqrt{\mathbf{S}_d}$, with rows \hat{X}_i interpreted as the embedded node positions in \mathbb{R}^d . This full embedding process is implemented in the open source `graspologic` package [8].

Both the ASE and LSE can be used to generate new networks of similar structure by treating the embedded node positions as latent positions within a random dot product graph (RDPG) model. The RDPG model generalizes SBMs. A d -dimensional RDPG model is parametrized by the $N \times d$ matrix of latent positions \mathbf{X} , with row X_i corresponding to the latent position of node i . The RDPG is an inhomogenous Erdos-Renyi graph with edge probability matrix $\mathbf{P} = \mathbf{X} \mathbf{X}^T$, i.e. the probability of edge $\{i, j\}$ is given by $p_{ij} = X_i \cdot X_j$.

With this in mind, we present ASE-RDPG and LSE-RDPG resampling distributions. These are the mappings from a given network \mathcal{G} to the RDPG taking $\hat{\mathbf{X}}$ as its latent positions, i.e. $\mathcal{G} \rightarrow \text{ER}(\hat{\mathbf{X}} \hat{\mathbf{X}}^T)$, with $\hat{\mathbf{X}}$ formed as above via ASE or LSE. However, given the relatively small size of our networks, we found it necessary to adjust our latent positions in order to generate fully connected networks. Suppose we are embedding a network with M edges. For a given edge probability matrix $\mathbf{P} = \hat{\mathbf{X}} \hat{\mathbf{X}}^T$, derived from either ASE or LSE, we may compute the expected number of edges as $\mathbb{E}e(\mathbf{P}) = \sum_{j>i} p_{ij}$. We then computed the adjusted edge probability matrix $\tilde{\mathbf{P}} = (M/\mathbb{E}e(\mathbf{P})) \mathbf{P}$, and used $\tilde{\mathbf{P}}$ to generate new networks in place of \mathbf{P} . Thus, using \mathbf{X}_A and \mathbf{X}_L to denote ASE and LSE spectral embeddings, we define our adjusted ASE-RDPG and LSE-RDPG distributions:

$$\text{ASE-RDPG}(\mathcal{G}) = \text{ER}(r_A \mathbf{X}_A \mathbf{X}_A^T) \quad (6)$$

$$\text{LSE-RDPG}(\mathcal{G}) = \text{ER}(r_L \mathbf{X}_L \mathbf{X}_L^T) \quad (7)$$

$$\text{where } r = \frac{|\mathcal{E}|}{\mathbb{E}e(\mathbf{X} \mathbf{X}^T)}$$

When generating new networks, we also discarded any that were not connected.

In Hypothesis 1 of Sec. 3.1, we predicted that DCP networks will outperform affinity (and CCP) networks in innovation tasks. Building on the Two Truths hypothesis from [22], we hold that ASE-RDPG resampling will generate new networks with similar DCP structure to the original network, and likewise LSE-RDPG will recreate affinity structure. Putting the positive claims of these two together, we offer the following hypothesis:

Hypothesis 2: Given a base network, ASE-RDPG resampling will recreate (decentralized) core-periphery structure in new networks, which will in turn better approximate the performance of the original SBM-generated network in our task, relative to LSE-RDPG resampling.

In order to determine which spectral resampling method better approximates the ABM performance of the base networks, we will compare distributions of discovery times from the original and resampled networks, using the Earth Mover's Distance (EMD). For discrete, real-valued distributions $\hat{f}(x_i)$ and $\hat{g}(y_j)$, the EMD is defined as:

$$d_{\text{EMD}}(\hat{f}, \hat{g}) = \inf_{\gamma \in \Gamma(\hat{f}, \hat{g})} \sum_{i,j} \gamma(x_i, y_j) |x_i - y_j| \quad (8)$$

where $\Gamma(\hat{f}, \hat{g})$ is the collection of all product distributions $\gamma(x_i, y_j)$ with \hat{f} and \hat{g} as marginal distributions.

4 RESULTS

4.1 Decentralized Core-Periphery Structure Accelerates Innovation Crossover Relative to Affinity Structure.

We first wanted to examine whether decentralized core-periphery structure promotes innovation crossover. In our first experiment, we generated networks along a parameter edge-level curve:

$$(a, b, c)(\theta) = (a_0, b_0, c_0) + \theta(-1, 0, 1), \quad 0 \leq \theta \leq \theta_{\max} \quad (M1)$$

We will choose (a_0, b_0, c_0) as a non-edge case in the decentralized core-periphery region. As θ increases, the discrepancy between intra-block connectivities decreases, until they assume an equitable affinity structure. In our first experiment, we set $a_0 = 0.75, c_0 = 0.05$, and consider four values for inter-block probability $b = b_0 = 0.01, 0.05, 0.1, \text{ or } 0.15$. For each parameter regime, we generated 500 networks from the SBM, and ran the simulation on each network until the final innovation was discovered, recording the number of time steps needed for each.

We present the results of this simulation experiment in Fig. 2. We show the distribution of discovery times as we increase θ , transitioning from decentralized core-periphery to affinity structure, for each choice of $b = b_0$. Consider first the case $b_0 = c_0 = 0.05$, where $(a, b, c)_0 = (0.75, 0.05, 0.05)$ is a relatively balanced decentralized core-periphery SBM, with equal core-periphery and intra-periphery edge probabilities. As can be seen, the mean crossover time is lowest for the decentralized core-periphery networks generated at $\theta = 0$, supporting Hypothesis 1. As θ increases, the mean crossover time likewise increases. Interestingly, the median is more stable across the parameter curve. The distribution of crossover times is non-Gaussian, and has a long right tail. The tail ‘shrinks’ as $\theta \rightarrow 0^+$, suggesting that decentralized core-periphery structure may accelerate innovation in expectation by reducing the likelihood of stagnation.

Now, consider this relationship between θ (DCP-to-affinity) and mean crossover time in light of the progression of distributions for all four values of b in Fig. 2. We see that this relationship is dependent upon the low value we chose for the inter-block parameter b . The performance of core-periphery networks *increases* if we decrease b to 0.01, further isolating the periphery from the core, while discovery times degrade rather quickly if we double or triple b . This suggests that the accelerated innovation we observe in simulation for core-periphery networks is dependent upon the network inefficiency that separates peripheral nodes from most of the core.

4.2 Centralized Core-Periphery Structure Does Not Accelerate Innovation Crossover, Relative to Affinity Structure.

The results presented in Fig. 2 suggest that the advantage in discovery times observed in simulation for our decentralized core-periphery innovation networks is dependent on their decentralization, i.e. the low inter-block connectivity parametrized in our SBM by b . This would cast doubt on any advantage in innovation speed for centralized core-periphery networks, in the sense of [2].

We considered this explicitly with two θ -parametrized families of SBMs. We thus turn to examining the performance of SBMs exhibiting CCP structure in our simulation. Similar to before, we consider the transition from CCP to affinity structure in our SBM:

$$(a, b, c) = (a_0, b_0, c_0) + \theta (0, -1, 1) \tag{M2}$$

We present the results of our simulations for (M2) in Fig. 3. Networks exhibiting centralized core-periphery structure seem to realize no advantage in innovation tasks, relative to affinity networks composed two internally dense, loosely connected communities. To the contrary, the affinity networks actually outperform CCP

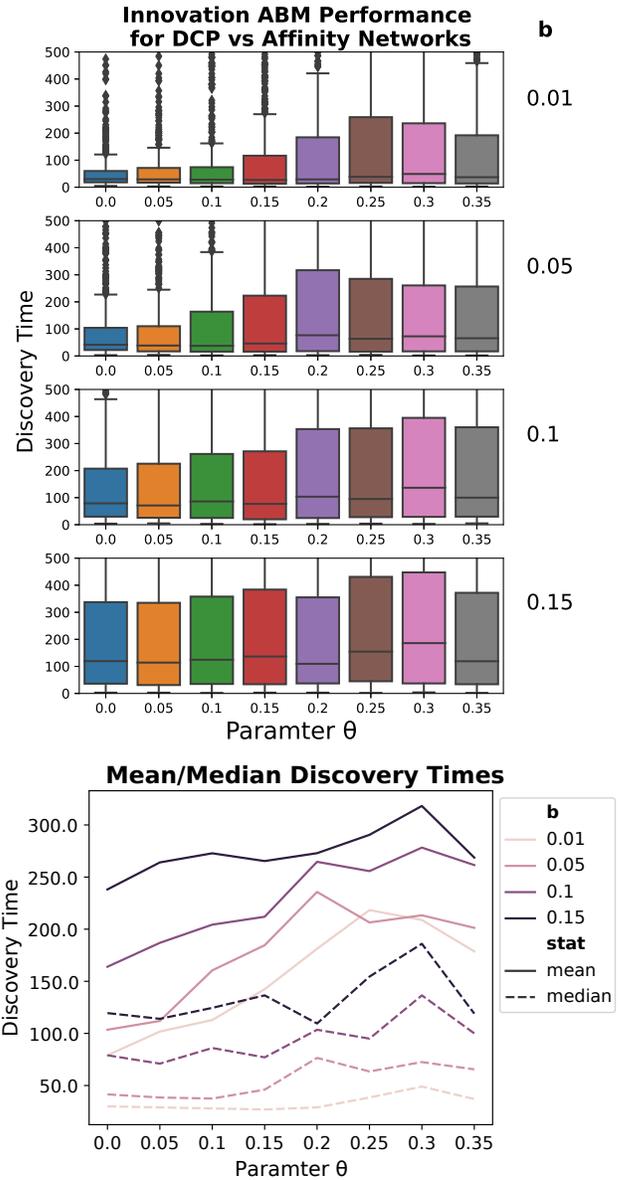


Figure 2: (TOP) The distributions of crossover discovery times (in discrete time steps) for networks generated from the family of SBMs (M1), parameterized by θ . For each row, $a_0 = 0.75, c_0 = 0.05$, and core-periphery connectivity b is held constant according to the value to the right of the plot. Networks exhibit DCP structure for small θ , and assume an affinity structure as θ grows. **(BOTTOM)** We separately visualize the means and medians of the distributions from the boxplots.

networks, in terms of lower mean and median discovery times. These results indicate that from the perspective of our collective

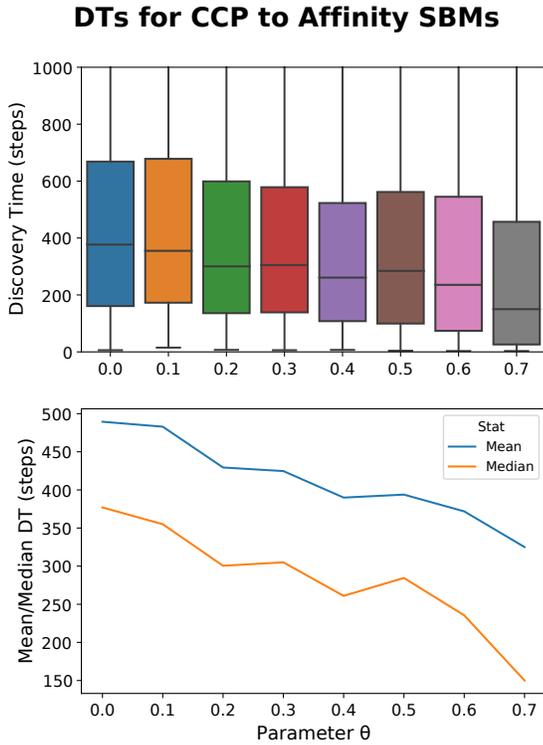


Figure 3: The distributions of discovery times for networks generated from the family (M2), with $(a_0, b_0, c_0) = (0.75, 0.75, 0.05)$. Networks exhibit CCP structure for $\theta = 0$, and assume an affinity structure as θ grows. As can be seen, mean and median discovery time decrease over the transition from CCP SBMs to affinity SBMs.

problem-solving, centralized and decentralized core-periphery networks are more different from each other than either is from a balanced affinity network, in terms of innovation discovery speed.

To conclude, our results in this section highlight longer discovery times exhibited by centralized core-periphery networks in comparison to affinity networks in the ABM. This suggests that the accelerated innovation observed in DCP networks requires both the decentralized ($a \gg b$) and core-periphery ($a \gg c$) elements of its structure. Moreover, our results highlight the qualitative distinction between the two types of core-periphery networks, which clearly impacts the diffusion and diversity of information within the network.

4.3 Decentralized Core-Periphery Structure Outperforms Centralized Core-Periphery Structure

In Sec. 4.1 and Sec. 4.2, we compared the innovation performance of DCP and CCP SBMs individually to affinity SBMs, finding DCP outperformed affinity networks, which are in turn outperformed CCP networks. We round out these comparisons by running a direct comparison between the two types of core-periphery networks:

DTs for DCP-to-CCP SBMs

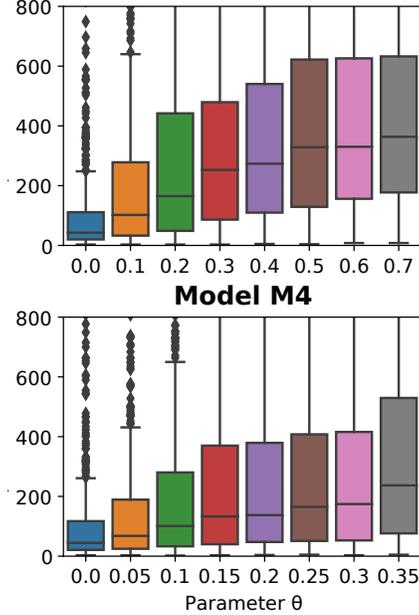


Figure 4: The distribution of discovery times for networks generated from (M3) and (M4), with $(a_0, b_0, c_0) = (0.75, 0.05, 0.05)$. Networks transition from DCP to CCP structure as θ increases. As can be seen, discovery time increases in both models as the core-periphery structure becomes centralized.

Statistical Distance of DTs from Spectral Bootstrapping

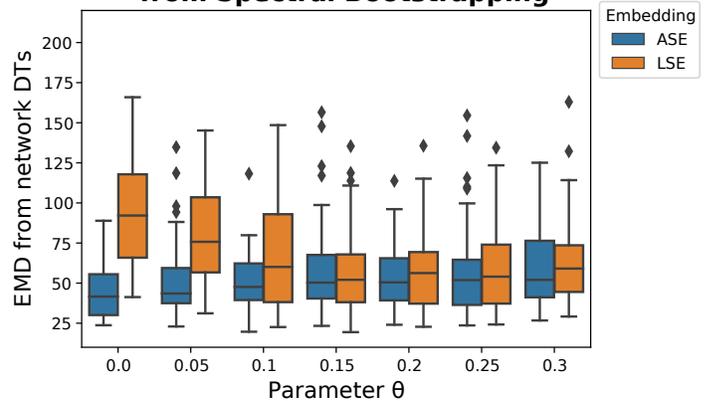


Figure 5: Statistical distances between the discovery time distributions of model (M1) networks (the original) and corresponding ASE/LSE-RDPG resampled networks. For $\theta = 0$ (DCP networks), we see that the DT distribution of ASE-RDPG networks more closely matches that of the original network, in line with the Two Truths hypothesis from [22].

decentralized and centralized. A more difficult task is that of determining the appropriate categorical comparison between DCP and CCP networks: we expect an SBM with $a, b \gg c$ to have more edges and shorter path distances than one with $a \gg b, c$, assuming we hold a and c constant. Thus, we will consider two families of SBMs, one preserving the expected number of edges at the expense of comparable core structure, and one allowing the CCP networks more edges in expectation in order to preserve comparable core density between CCP and DCP:

$$(a, b, c) = (a_0, b_0, c_0) + \theta (0, 1, 0) \quad (\text{M3})$$

$$(a, b, c) = (a_0, b_0, c_0) + \theta (-1, 1, 0) \quad (\text{M4})$$

The family (M3) breaks with our convention up to this point, increasing the expected number of edges linearly in θ in order to transition from DCP to CCP structure. The family (M4) preserves average density, but at the cost of a less dense core. Fortunately this mattered relatively little for the qualitative results.

We present the results of simulating both families of SBMs in Fig. 4. As can be seen, regardless of which parameter curve is taken, transitioning from DCP to CCP structure increases mean discovery times.

4.4 Performance of ASE/LSE-Resampled Networks Implicates (Decentralized) Core-Periphery Structure.

We have demonstrated that DCP networks outperform both affinity and CCP networks in our ABM, supporting Hypothesis 1. Recall that the distinction between DCP and affinity network structure, as previously articulated in [22], was motivated by a Two Truths phenomenon in spectral clustering, whereby both ASE- and LSE-based community detection found meaningful but distinct pairs of ground-truth communities in connectome networks. In what we have referred to as the Two Truths hypothesis, the authors suggested that ASE captures (decentralized) core-periphery structure, while LSE captures affinity structure. Thus, we posed Hypothesis 2, in which we speculated that networks generated by ASE-RDPG resampling would better reproduce the original ABM performance of a network exhibiting DCP structure, while those from LSE-RDPG resampling would do so for affinity networks.

To test this hypothesis, we again consider networks generated from the SBM family (M1). We used the base DCP parameter set $(a_0, b_0, c_0) = (0.75, 0.05, 0.15)$.¹ For each θ , we generated 50 networks. For each network, we ran 100 simulations of the ABM, forming an empirical distribution \hat{f} of crossover times. We then generated 100 networks each via ASE-RDPG and LSE-RDPG resampling, simulating our ABM once per resampled network. This provided us with empirical distributions \hat{g}_{ASE} and \hat{g}_{LSE} of discovery times. Note that the randomness of \hat{f} is purely that of the simulator, while both resampled \hat{g} have randomness associated with both the simulator and the network resampling itself. From all three distributions, we compute the statistical distance $d_{\text{EMD}}(\hat{f}, \hat{g})$ for

each embedding (Sec. 3.3). We collected these distances for all 50 base networks for each θ .

Figure 5 shows the distributions of the statistical distance between \hat{f} and \hat{g} , for each type of spectral embedding, grouped by θ . We see that for $\theta = 0$, i.e. the SBM with the strongest DCP structure, the average distance between DTs from the SBM-generated network and those of the ASE-RDPG resampled networks ($d_{\text{EMD}}(\hat{f}, \hat{g}_{\text{ASE}})$) was significantly less than the distance to the DT distribution of the LSE-RDPG networks ($d_{\text{EMD}}(\hat{f}, \hat{g}_{\text{LSE}})$). This suggests that, in line with the Two Truths hypothesis, ASE-RDPG is preferentially capturing and recreating the DCP structure of the original network, which in turn is responsible for more similar network performance in the ABM simulations. This supports Hypothesis 2.

However, we also see that when $\theta = 0.35$, i.e. $(a, b, c) = (0.45, 0.05, 0.45)$ and the SBM is generating affinity networks, ASE-RDPG and LSE-RDPG resampling recreate the DTs of the original network to a similar degree. That is to say, neither is preferentially producing networks closer in performance to the base SBM network. If we may take the Two Truths hypothesis for granted, our results in this section might suggest that the performance of DCP networks is due to DCP structure, but the performance of affinity networks is less reflective of affinity community structure as such. Alternatively, it is likely that the Two Truths phenomenon observed by [22] is a tendency, but the network features captured by each embedding are far more subtle than DCP versus affinity. In other words, the Two Truths hypothesis should be given less weight.

5 DISCUSSION

Our investigation highlights the role that core-periphery structure plays in a network’s ability to solve a collective problem-solving task. Furthermore, it has provided an operational understanding of the topological distinction between two forms of core-periphery structure, that of decentralized core-periphery (DCP) structure and that of centralized core-periphery (CCP) structure.

Utilizing an agent-based model in which populations must maintain a balance between exploring novel combinations and exploiting combinations they have, we showed that DCP networks, in which peripheral nodes are isolated from both a centralized core and one another, outperform affinity networks; that this same relationship is not found between affinity and CCP networks, in which peripheral nodes are isolated from one another but still maintain close links with the core; and that in a direct comparison, DCP networks outperform CCP networks.

In addition to this contribution to the literature on collective problem solving, we have investigated the implications of using different spectral graph embeddings for the resampling of networks in agent-based modeling. This investigation has provided an operational and dynamic significance to a graph-theoretic phenomenon, in this case the Two Truths hypothesis for spectral clustering [22]. By sampling generated networks and parameterizing based on their ASE, which is hypothesized to capture core-periphery structure, and LSE, which is hypothesized to capture affinity structure, and comparing them to performance in our original DCP networks, we found that our ASE generated networks performed most similarly to our DCP networks, while the LSE networks did not.

¹We initially attempted to use the same base DCP parameter set $(a_0, b_0, c_0) = (0.75, 0.05, 0.05)$. Interestingly, even with the adjustments described in Sec. 3.3 to scale the latent positions to the desired expected degree, we found that the ASE-RDPG rarely generated connected networks from a connected network.

ACKNOWLEDGMENTS

This work received support from the Institute for Humane Studies under grant no. IHS017131 awarded to CM.

REFERENCES

- [1] Danielle S Bassett, Nicholas F Wymbs, M Puck Rombach, Mason A Porter, Peter J Mucha, and Scott T Grafton. 2013. Task-based core-periphery organization of human brain dynamics. *PLoS Computational Biology* 9, 9 (2013), e1003171.
- [2] Stephen P Borgatti and Martin G Everett. 2000. Models of core/periphery structures. *Social Networks* 21, 4 (2000), 375–395.
- [3] Mauricio Cantor, Michael Chimento, Simeon Q Smelee, Peng He, Danai Papa-georgiou, Lucy M Aplin, and Damien R Farine. 2021. Social network architecture and the tempo of cumulative cultural evolution. *Proceedings of the Royal Society B* 288, 1946 (2021), 20203107.
- [4] Gino Cattani and Simone Ferriani. 2008. A core/periphery perspective on individual creative performance: Social networks and cinematic achievements in the Hollywood film industry. *Organization Science* 19, 6 (2008), 824–844.
- [5] Gino Cattani, Simone Ferriani, and Mariachiara Colucci. 2015. Creativity in Social Networks. In *The Oxford handbook of creative industries*. Oxford University Press Oxford, 75–95.
- [6] Damon Centola. 2021. Influencers, Backfire Effects, and the Power of the Periphery. *Personal Networks: Classic Readings and New Directions in Egocentric Analysis* 51 (2021), 73.
- [7] Damon Centola. 2022. The network science of collective intelligence. *Trends in Cognitive Sciences* (2022).
- [8] Jaewon Chung, Benjamin D Pedigo, Eric W Bridgeford, Bijan K Varjavand, Hayden S Helm, and Joshua T Vogelstein. 2019. GraSPy: Graph Statistics in Python. *J. Mach. Learn. Res.* 20, 158 (2019), 1–7. <https://github.com/microsoft/graspologic/>
- [9] Jeff Clune, Jean-Baptiste Mouret, and Hod Lipson. 2013. The evolutionary origins of modularity. *Proceedings of the Royal Society b: Biological sciences* 280, 1755 (2013), 20122863.
- [10] Maxime Derex and Robert Boyd. 2016. Partial connectivity increases cultural accumulation within groups. *Proceedings of the National Academy of Sciences* 113, 11 (2016), 2982–2987.
- [11] Ryan J Gallagher, Jean-Gabriel Young, and Brooke Foucault Welles. 2021. A clarified typology of core-periphery structure in networks. *Science advances* 7, 12 (2021), eabc9800.
- [12] Sadamori Kojaku and Naoki Masuda. 2017. Finding multiple core-periphery pairs in networks. *Physical Review E* 96, 5 (2017), 052313.
- [13] David Lazer and Allan Friedman. 2007. The network structure of exploration and exploitation. *Administrative science quarterly* 52, 4 (2007), 667–694.
- [14] Keith Levin and Elizaveta Levina. 2019. Bootstrapping networks with latent space structure. *arXiv preprint arXiv:1907.10821* (2019).
- [15] Vince Lyzinski, Minh Tang, Avanti Athreya, Youngser Park, and Carey E Priebe. 2016. Community detection and classification in hierarchical stochastic block-models. *IEEE Transactions on Network Science and Engineering* 4, 1 (2016), 13–26.
- [16] James G March. 1991. Exploration and exploitation in organizational learning. *Organization Science* 2, 1 (1991), 71–87.
- [17] Henok Mengistu, Joost Huizinga, Jean-Baptiste Mouret, and Jeff Clune. 2016. The evolutionary origins of hierarchy. *PLoS computational biology* 12, 6 (2016), e1004829.
- [18] Andrea B Migliano, Federico Battiston, Sylvain Viguier, Abigail E Page, Mark Dyble, Rodolph Schlaepfer, Daniel Smith, Leonora Astete, Marilyn Ngales, Jesus Gomez-Gardenes, et al. 2020. Hunter-gatherer multilevel sociality accelerates cumulative cultural evolution. *Science advances* 6, 9 (2020), eaax5913.
- [19] Andrea B Migliano, Abigail E Page, Jesus Gómez-Gardeñes, Gul Deniz Salali, Sylvain Viguier, Mark Dyble, James Thompson, Nikhill Chaudhary, Daniel Smith, Janis Strods, et al. 2017. Characterization of hunter-gatherer networks and implications for cumulative culture. *Nature Human Behaviour* 1, 2 (2017), 1–6.
- [20] Cody Moser and Paul E Smaldino. 2023. Innovation-Facilitating Networks Create Inequality. (2023).
- [21] Mark EJ Newman. 2006. Modularity and community structure in networks. *Proceedings of the national academy of sciences* 103, 23 (2006), 8577–8582.
- [22] Carey E Priebe, Youngser Park, Joshua T Vogelstein, John M Conroy, Vince Lyzinski, Minh Tang, Avanti Athreya, Joshua Cape, and Eric Bridgeford. 2019. On a two-truths phenomenon in spectral graph clustering. *Proceedings of the National Academy of Sciences* 116, 13 (2019), 5995–6000.
- [23] M Puck Rombach, Mason A Porter, James H Fowler, and Peter J Mucha. 2014. Core-periphery structure in networks. *SIAM Journal on Applied mathematics* 74, 1 (2014), 167–190.
- [24] Puck Rombach, Mason A Porter, James H Fowler, and Peter J Mucha. 2017. Core-periphery structure in networks (revisited). *SIAM review* 59, 3 (2017), 619–646.
- [25] Paul E Smaldino, Cody Moser, Alejandro Pérez Velilla, and Mikkel Werling. 2022. Maintaining transient diversity is a general principle for improving collective problem solving. (2022).
- [26] Daniel L Sussman, Minh Tang, Donniel E Fishkind, and Carey E Priebe. 2012. A consistent adjacency spectral embedding for stochastic blockmodel graphs. *J. Amer. Statist. Assoc.* 107, 499 (2012), 1119–1128.
- [27] Daniel L Sussman, Minh Tang, and Carey E Priebe. 2013. Consistent latent position estimation and vertex classification for random dot product graphs. *IEEE transactions on pattern analysis and machine intelligence* 36, 1 (2013), 48–57.
- [28] Minh Tang and Carey E Priebe. 2018. Limit theorems for eigenvectors of the normalized Laplacian for random graphs. *The Annals of Statistics* 46, 5 (2018), 2360–2415.
- [29] Francesco Tudisco and Desmond J Higham. 2019. A nonlinear spectral method for core-periphery detection in networks. *SIAM Journal on Mathematics of Data Science* 1, 2 (2019), 269–292.
- [30] Madeleine Udell and Alex Townsend. 2019. Why are big data matrices approximately low rank? *SIAM Journal on Mathematics of Data Science* 1, 1 (2019), 144–160.
- [31] Ulrike Von Luxburg. 2007. A tutorial on spectral clustering. *Statistics and computing* 17, 4 (2007), 395–416.
- [32] Eric Yanchenko and Srijan Sengupta. 2022. Core-periphery structure in networks: a statistical exposition. *arXiv preprint arXiv:2202.04455* (2022).
- [33] Mu Zhu and Ali Ghodsi. 2006. Automatic dimensionality selection from the scree plot via the use of profile likelihood. *Computational Statistics & Data Analysis* 51, 2 (2006), 918–930.