

The Artificial Benchmark for Community Detection with Outliers and Overlapping Communities (**ABCD**+**o**²)

Jordan Barrett*, Ryan DeWolfe[†], Bogumił Kamiński[‡],
Paweł Prałat[§], Aaron Smith[¶], François Thériberge^{||}

June 4, 2025

Abstract

The **A**rtificial **B**enchmark for **C**ommunity **D**etection (**ABCD**) graph is a random graph model with community structure and power-law distribution for both degrees and community sizes. The model generates graphs similar to the well-known **LFR** model but it is faster, more interpretable, and can be investigated analytically. In this paper, we use the underlying ingredients of the **ABCD** model, and its generalization to include outliers (**ABCD**+**o**), and introduce another variant that allows for overlapping communities, **ABCD**+**o**².

Keywords— ABCD, community detection, overlapping communities, benchmark models

1 Introduction

One of the most important features of real-world networks is their community structure, as it reveals the internal organization of nodes. In social networks, communities may represent groups by interest; in citation networks, they correspond to related papers; in the Web graph, communities are

*Department of Mathematics and Statistics, Dalhousie University, Halifax, NS, Canada; e-mail: jbarrett@dal.ca

[†]Department of Mathematics, Toronto Metropolitan University, Toronto, ON, Canada; e-mail: ryan.dewolfe@torontomu.ca

[‡]Decision Analysis and Support Unit, SGH Warsaw School of Economics, Warsaw, Poland; e-mail: bkamins@sgh.waw.pl

[§]Department of Mathematics, Toronto Metropolitan University, Toronto, ON, Canada; e-mail: pralat@torontomu.ca

[¶]Department of Mathematics and Statistics, University of Ottawa, Ottawa, ON, Canada; e-mail: asmi28@uottawa.ca

^{||}Tutte Institute for Mathematics and Computing, Ottawa, ON, Canada; email: theberge@ieee.org

formed by pages on related topics, etc. Identifying communities in a network is therefore valuable, as it helps us understand the structure of the network.

Detecting communities is quite a challenging task. In fact, there is no definition of community that researchers and practitioners agree on. Still, it is widely accepted that a community should induce a graph that is denser than the global density of the network [17]. Numerous community detection algorithms have been developed over the years, using various techniques such as optimizing modularity, removing high-betweenness edges, detecting dense subgraphs, and statistical inference. We direct the interested reader to the survey [15] or one of the numerous books on network science [22].

Most community detection algorithms aim to find a partition of the set of nodes, that is, a collection of pairwise disjoint communities with the property that each node belongs to exactly one of them. This is a natural assumption for many scenarios. For example, most of the employees on LinkedIn work for a single employer. On the other hand, users of Instagram can belong to many social groups associated with their workplace, friends, sports, etc. Researchers might be part of many research groups. A large fraction of proteins belong to several protein complexes simultaneously. As a result, many real-world networks are better modelled as a collection of overlapping communities [30] and, moreover, many community detection problems should be tackled by finding overlapping sets of nodes rather than a partition.

In the context of overlapping communities, detection is even more challenging. For example, in non-overlapping community detection one can easily verify that a node i is highly connected to its proposed community C , whereas in overlapping community detection a node i might be assigned to many communities C_1, \dots, C_k , and the number of connections from i into each community may vary drastically.

To support the development and analysis of overlapping community detection algorithms, a large and diverse catalogue of networks with ground-truth communities is required for testing, tuning and training. Unfortunately, there are few such datasets with communities properly identified and labelled. As a result, there is a need for synthetic random graph models with community structure that resemble real-world networks to benchmark and tune unsupervised clustering algorithms. The popular **LFR** (**L**ancichinetti, **F**ortunato, **R**adicchi) model [27, 26] generates networks with communities and, at the same time, allows for heterogeneity in the distributions of both node degrees and of community sizes. Due to this structural freedom, the **LFR** model became a standard and extensively used model for generating artificial networks with ground-truth communities.

A model similar to **LFR**, the **A**rtificial **B**enchmark for **C**ommunity **D**etection (**ABCD**) [21], was recently introduced and implemented¹, along with a faster and multithreaded implementation² (**ABCDe**) [19]. Undirected variants of **LFR** and **ABCD** produce graphs with comparable properties, but **ABCD** (and especially **ABCDe**) is faster than **LFR** and can be easily tuned to allow the user to make a smooth transition between the two extremes: pure (disjoint) communities and random graphs with no community structure. Moreover, **ABCD** is easier to analyze theoretically—for example, in [20] various theoretical asymptotic properties of the are investigated, including the modularity function that, despite some known issues such as the “resolution limit” reported in [16], is an important graph property of networks in the context of community detection. In [6], some interesting and desired self-similar behaviour of the **ABCD** model is discovered; namely, that the degree distribution of ground-truth communities is asymptotically the same as

¹<https://github.com/bkamins/ABCDGraphGenerator.jl/>

²<https://github.com/tolcz/ABCDeGraphGenerator.jl/>

the degree distribution of the whole graph (appropriately normalized based on their sizes). Finally, the building blocks in the model are flexible and may be adjusted to satisfy different needs. Indeed, the original **ABCD** model was adjusted to include outliers (**ABCD+o**) [23] and was generalized to hypergraphs (**h-ABCD**) [24]³. For these reasons **ABCD** is gaining recognition as a benchmark for community detection algorithms. For example, in [2] the authors use both **ABCD** and **LFR** graphs to compare 30 community detection algorithms and, in their work, highlight that “*being directly comparable to **LFR**, **ABCD** offers additional benefits, including higher scalability and better control for adjusting an analogous mixing parameter.*”

In this paper we introduce the **ABCD+o**² model: a generalization of the **ABCD+o** model that allows for overlapping communities. The **LFR** model has been extended in a similar way [26], and in this model the nodes are assigned to communities based on the construction of a random bipartite graph between nodes and communities resulting in (a) a small amount of overlap between almost every pair of communities, and (b) rarely any pair of communities with a large overlap. In **ABCD+o**², we instead generate overlapping communities based on a hidden, low-dimensional geometric layer which tends to yield fewer and larger overlaps. This geometric approach to community structure is justified by the fact that latent geometric spaces are believed to shape complex networks (e.g., social media networks shaped by users’ opinions, education, knowledge, interests, etc.). These latent spaces have been successfully employed for many years to model and explain network properties such as self-similarity [33], homophily and aversion [18]. For more details, we direct the reader to the survey [9] or the book [32]. In addition to the geometric layer, the ancillary benefits of the **ABCD** model (an intuitive noise parameter, a fast implementation, and theoretical analysis) are still present, making the **ABCD+o**² model an attractive option for benchmarking community detection algorithms.

The remainder of the paper is organized as follows. In Section 2 we present the **ABCD+o**² model, with a full description of generating a graph in Section 2.5. Next, in Section 3 we show some properties of the model and compare these properties to those of real graphs with known overlapping communities. In Section 4 we demonstrate an application of the model by benchmarking community detection algorithms and comparing their quality under different levels of noise and overlap. Finally, some concluding remarks are given in Section 5.

2 **ABCD+o**² — **ABCD** with Overlapping Communities and Outliers

As mentioned in the introduction, the original **ABCD** model was extended to include outliers resulting in the **ABCD+o** model. For our current needs, we extend **ABCD+o** further to allow for non-outlier nodes to belong to multiple communities, resulting in the **ABCD+o**² model, **ABCD** with overlapping communities and outliers.

2.1 Notation

For a given $n \in \mathbb{N} := \{1, 2, \dots\}$, we use $[n]$ to denote the set consisting of the first n natural numbers, that is, $[n] := \{1, 2, \dots, n\}$.

³<https://github.com/bkamins/ABCDHypergraphGenerator.jl>

We use standard probability notation throughout the paper. For a random variable X , write $\mathbb{P}(X = k)$ for the probability that $X = k$, and write $\mathbb{E}[X]$ for the expected value of X . For a distribution \mathcal{D} , write $X \sim \mathcal{D}$ to mean X is sampled from the distribution \mathcal{D} .

Power-law distributions will be used to generate both the degree sequence and community sizes so let us formally define it. For given parameters $\gamma \in (0, \infty)$, $\delta, \Delta \in \mathbb{N}$ with $\delta \leq \Delta$, we define a truncated power-law distribution $\mathcal{P}(\gamma, \delta, \Delta)$ as follows. For $X \sim \mathcal{P}(\gamma, \delta, \Delta)$ and for $k \in \mathbb{N}$ with $\delta \leq k \leq \Delta$,

$$\mathbb{P}(X = k) = \frac{\int_k^{k+1} x^{-\gamma} dx}{\int_\delta^{\Delta+1} x^{-\gamma} dx}. \quad (1)$$

2.2 The Configuration Model

The well-known configuration model is an important ingredient of all variants of the **ABCD** models, so let us formally define it here. Suppose that our goal is to create a graph on n nodes with a given degree distribution $\mathbf{d} := (d_i, i \in [n])$, where \mathbf{d} is a sequence of non-negative integers such that $m := \sum_{i \in [n]} d_i$ is even. We define a random multi-graph $\text{CM}(\mathbf{d})$ with a given degree sequence known as the **configuration model** (sometimes called the **pairing model**), which was first introduced by Bollobás [10]. (See [7, 36, 35] for related models and results.)

We start by labelling nodes as $[n]$ and, for each $i \in [n]$, endowing node i with d_i half-edges. We then iteratively choose two unpaired half-edges uniformly at random (from the set of pairs of remaining half-edges) and pair them together to form an edge. We iterate until all half-edges have been paired. This process yields a graph $G_n \sim \text{CM}(\mathbf{d})$ on n nodes, where G_n is allowed self-loops and multi-edges and thus G_n is a multi-graph.

2.3 Parameters of the **ABCD+o²** Model

The following parameters govern the **ABCD+o²** model.

Note that the ranges for parameters γ and β are suggestions chosen according to experimental values commonly observed in complex networks [3, 29]. In fact, users may inject any degree sequence and sequence of community sizes as inputs to the model.

2.4 Big Picture

The **ABCD+o²** model generates a random graph on n nodes with degree sequence $(d_i, i \in [n])$ and community size sequence $(s_i, i \in [L])$ following truncated power laws with exponents γ and, respectively, β . There are s_0 outliers and $\hat{n} = n - s_0$ non-outliers. The non-outliers span a family of L communities $(C_j, j \in [L])$ with each non-outlier belonging to at least one such community. These communities are generated, and will overlap (unless $\eta = 1$), according to a hidden d -dimensional reference layer in such a way that non-outliers will belong to η communities, on average. The non-outliers, with their respective degrees, populate $(C_j, j \in [L])$ randomly with the caveat that high degree nodes cannot enter small communities unless absolutely necessary. Moreover, one may insist on a correlation between the degree of a node and the number of communities it belongs to. This correlation is controlled by the parameter ρ .

Parameter $\xi \in [0, 1]$ dictates the amount of noise in the network. Each non-outlier node i has its degree d_i split into two parts: *community degree* Y_i and *background degree* Z_i (and thus $d_i = Y_i + Z_i$).

Parameter	Range	Description
n	\mathbb{N}	Number of nodes
s_0	\mathbb{N}	Number of outliers
η	$[1, \infty)$	Average number of communities a non-outlier node is part of
d	\mathbb{N}	Dimension of reference layer
ρ	$[-1, 1]$	Pearson correlation between the degree of nodes and the number of communities they belong to
γ	$(2, 3)$	Power-law degree distribution with exponent γ
δ	\mathbb{N}	Min degree as least δ
Δ	$\mathbb{N} \setminus [\delta - 1]$	Max degree at most Δ
β	$(1, 2)$	Power-law community size distribution with exponent β
s	$\mathbb{N} \setminus [\delta]$	Min community size at least s
S	$\mathbb{N} \setminus [s - 1]$	Max community size at most S
ξ	$[0, 1]$	Level of noise

Table 1: Parameters of the **ABCD+o²** model.

Using a random rounding technique, we choose the community degrees and background degrees so that $\mathbb{E}[Y_i] = (1 - \xi)d_i$ and $\mathbb{E}[Z_i] = \xi d_i$. Note that the neighbours of outliers are sampled from the entire graph, ignoring the underlying community structure, meaning $Y_i = 0$ and $Z_i = d_i$ if i is an outlier.

Once nodes are assigned to communities and their degrees are split, the edges of each community are then independently generated by the configuration model on the corresponding community degree sequences, i.e., community C is generated by the configuration on the degree sequence $(Y_i, i \in C)$. Once communities are generated, the background graph is generated by the configuration model on the degree sequence $(Z_i, i \in [n])$. The final **ABCD+o²** model, after an additional clean-up phase to rewire self-loops and duplicate edges, is the union of the community graphs and the background graph.

2.5 The **ABCD+o²** Construction

The following 6-phase construction process generates the **ABCD+o²** synthetic networks.

Phase 1: creating the degree distribution.

This phase is the same as in the original **ABCD** model and its generalization, **ABCD+o**. By default, we sample $d_i \sim \mathcal{P}(\gamma, \delta, \Delta)$, independently for each $i \in [n]$, then re-label the samples so that $d_1 \geq \dots \geq d_n$. To ensure that $\sum_{i \in [n]} d_i$ is even, we decrease d_1 by 1 if necessary; we relabel as needed to ensure that $d_1 \geq \dots \geq d_n$. Alternatively, the degree sequence can be given explicitly as an input.

Phase 2: assigning nodes as outliers.

This phase is also the same as in the **ABCD+o** model. As mentioned in the big picture summary, the neighbours of outliers will be sampled from the entire graph, ignoring the underlying community structure. It feels that this part is straightforward, but a problem might occur when ξ is close to zero. In the extreme case, when $\xi = 0$, only the outliers have a non-zero degree in the background graph, and for a simple background graph to exist all outliers must have degree at most $s_0 - 1$. To handle this issue, we restrict which degrees can be assigned to outliers in the following way. We have that $\ell = \sum_{i \in [n]} \min(1, \xi d_i)$ is a lower bound for the expected number of nodes with positive degree in the background graph. Moreover, the s_0 outliers will have a positive degree in the background graph. Thus, we expect $\ell + (n - \ell)(s_0/n)$ nodes to have positive degree in the background graph, assuming outliers are selected uniformly at random. Therefore, we insist that a node i of degree d_i cannot become an outlier unless

$$d_i \leq \ell + s_0 - \ell s_0/n - 1. \quad (2)$$

In practice (when the number of nodes n is large, the number of outliers s_0 is relatively small, and the level of noise ξ is not zero), there are plenty of nodes with a non-zero degree in the background graph, and so there is no restriction needed for outliers. Nevertheless, we include the restriction to ensure no issues during Phase 6; a subset of s_0 nodes satisfying (2) are selected uniformly at random to become outliers.

Phase 3: creating overlapping communities.

By the end of Phase 2, we have a degree sequence $(d_i, i \in [n])$ and an assignment of degrees to outliers and non-outliers. It is important to keep in mind that, although communities are created in this phase, we do not delegate degrees to these communities until Phase 4. We instead construct a collection of *elements* which will be points in d -dimensional space, each assigned to some number of communities. We will construct $\hat{n} = n - s_0$ elements and they will belong to $\eta \geq 1$ communities on average. To be compatible with the original **ABCD** model, each element will belong to a single *primary* community, and these primary memberships will partition the elements. Then, we will grow each community by a factor of η , according to the geometric layer, so that the collective size of all communities is equal (in expectation) to $\eta \hat{n} = \eta(n - s_0)$. For each element v added to community C during the growth process, C will be referred to as a *secondary* community of v .

Similar to the degree sequence, the community size sequence $(s_j, j \in [L])$ follows a power-law with parameter β , minimum value s , and maximum value S . Hence, the sequence of primary communities $(\hat{s}_j, j \in [L])$ will need to follow a power-law with parameter β , minimum value $\hat{s} = \lceil s/\eta \rceil$, and maximum value $\hat{S} = \lfloor S/\eta \rfloor$. In addition, we require $\sum_{j \in [L]} \hat{s}_j = \hat{n}$. To satisfy both requirements, we sample $\hat{s}_j \sim \mathcal{P}(\beta, \hat{s}, \hat{S})$ independently until the sum of samples is at least \hat{n} . If, at this point, the sum is $\hat{n} + a$ with $a > 0$ then we perform one of two actions. If the last added sample has size at least $a + \hat{s}$, we reduce it by a . Otherwise (that is, if $c < a + \hat{s}$), we delete this sample, select $c - a$ old samples uniformly at random and increase each by 1. Lastly, let L be the random variable counting the number of communities, and relabel $(\hat{s}_j, j \in [L])$ so that $\hat{s}_1 \geq \dots \geq \hat{s}_L$. Similar to the degree sequence, the sequence of community sizes can instead be given explicitly as an input. Each primary community of size \hat{s}_j will grow to size $s_j = \lfloor \eta \hat{s}_j \rfloor$ where,

for $a \in \mathbb{Z}$ and $b \in [0, 1)$, the random variable $\lfloor a + b \rfloor$ is defined as

$$\lfloor a + b \rfloor := \begin{cases} a & \text{with probability } 1 - b, \text{ and} \\ a + 1 & \text{with probability } b. \end{cases}$$

Note that $\mathbb{E}[\lfloor a + b \rfloor] = a(1 - b) + (a + 1)b = a + b$. As a result,

$$\mathbb{E} \left[\sum_{j \in [L]} s_j \right] = \sum_{j \in [L]} \mathbb{E}[s_j] = \eta \sum_{j \in [L]} \hat{s}_j = \eta \hat{n} = \eta(n - s_0),$$

as desired.

We now create our d -dimensional reference layer that will guide the process of assigning elements to communities. One may think of this reference layer as various latent properties of objects associated with nodes (such as people's age, education, geographic location, beliefs, etc.) shaping communities (such as communities in social media). In this reference layer, each of the \hat{n} elements is generated as a vector in \mathbb{R}^d , sampled independently and uniformly at random from the ball of radius 1 centred at the origin $\mathbf{0} = (0, 0, \dots, 0)$. Let R be the set of \hat{n} elements and initialize $R_1 = R$. We assign elements to their primary communities, dealing with one primary community at a time. To form primary community \hat{C}_j , we first select the element in R_j furthest from $\mathbf{0}$. This element, together with its $\hat{s}_j - 1$ nearest neighbours in R_j , form the elements of \hat{C}_j . We then remove these elements from R_j to construct R_{j+1} and repeat the process until all elements have been assigned a primary community. At the end of the process, each element belongs to exactly one primary community and the collection $(\hat{C}_j, j \in [L])$ partitions the elements.

We now grow each primary community \hat{C}_j of size \hat{s}_j to form the full community C_j of size s_j . Initially, we set $C_j = \hat{C}_j$ for all $j \in [L]$. Each community will grow independently, meaning we can grow them in any order (or in parallel). As before, let R be the set of \hat{n} elements. For primary community \hat{C}_j , let $\mathbf{x}_j \in \mathbb{R}^d$ be the center of mass (mean) of \hat{C}_j . We investigate elements of R in order of increasing distance from \mathbf{x}_j . During the investigation, if an element $v \in R$ is not a member of \hat{C}_j , then we assign C_j as a secondary community of v . We stop investigating once the number of members in C_j (both primary and secondary) is s_j . Write η_v for the number of communities containing element $v \in R$.

At the end of the process, each element is assigned to exactly one primary community and some (possibly zero) secondary communities. In Figure 1 we show an example of the reference layer on $\hat{n} = 150$ elements and three communities. Each of the three primary communities in this example consists of 50 elements before growing by a factor of $\eta = 2.0$, attracting an additional 50 elements as its secondary members. From a computational perspective, to ensure that finding primary and secondary communities is efficient, we use k-d trees (short for k-dimensional tree) to perform spatial lookups of the sampled points [8].

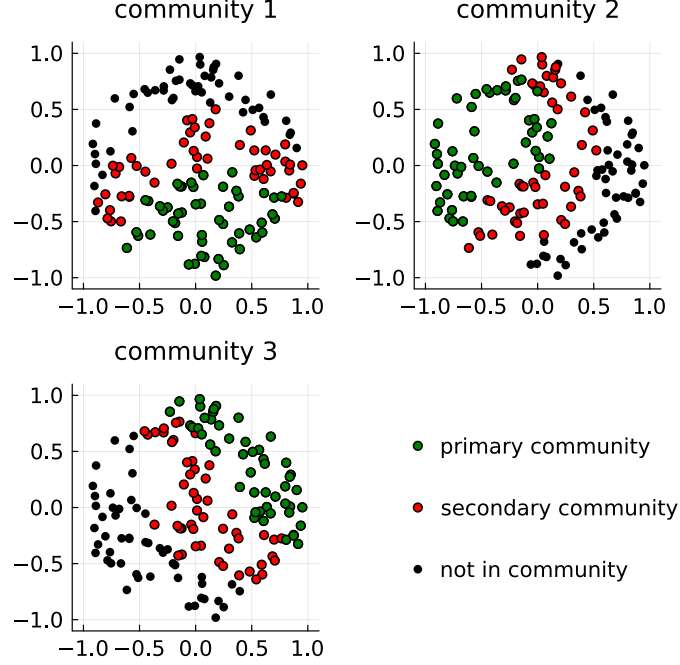


Figure 1: Example of the reference layer (in $d = 2$ dimension) on $\hat{n} = 150$ elements consisting of 3 overlapping communities with equal sizes and $\eta = 2.0$.

Phase 4: assigning degrees to nodes.

We now have a degree sequence $(d_i, i \in [n])$, an assignment of degrees to outliers and non-outliers, and a collection of overlapping communities containing elements, each element belonging to one primary community and some number of secondary communities. Let $\hat{\mathbf{d}}_{\hat{n}}$ be the subsequence (of length \hat{n}) of $(d_i, i \in [n])$ corresponding to the non-outliers. We are now ready to assign degrees in $\hat{\mathbf{d}}_{\hat{n}}$ to the set of elements R which will give us the nodes of the graph. Each element $v \in R$, once paired with node i of degree d_i , is expected to have ξd_i neighbours in the background graph and the remaining $(1 - \xi)d_i$ neighbours split evenly between its η_v communities. Note that, although each element has a distinct primary community, its degree will be split with no preference given to said primary community.

Similar to the potential problem with outliers, we want to avoid large degrees being assigned to small communities. If degree d_i is paired with element v then, for each community C_j containing v , s_j must be at least $(1 - \xi)d/\eta_v + 1$ to accommodate for the neighbours of v . This is in fact a lower bound as additional neighbours of v coming from the background graph might end up in C_j as well. To make additional room in the community graph, a small correction is introduced in both **ABCD** and **ABCD+o** guided by the parameter

$$\phi = 1 - \sum_{k \in [L]} \left(\frac{\hat{s}_k}{\hat{n}} \right)^2 \frac{\hat{n}\xi}{\hat{n}\xi + s_0}.$$

Typically, ϕ is very close to 1 and so this correction tends to be negligible in both theory and in practice. For consistency, we keep ϕ in the **ABCD+o**² model.

Before detailing the pairing process, let us explain how we deal with the correlation parameter ρ . Recall that our goal is to pair the \hat{n} non-outlier degrees with the \hat{n} elements such that the Pearson correlation between the degree and number of communities is approximately ρ . To this end, we perform multiple pairing attempts, with each attempt aiming to achieve a closer approximate the desired correlation. In each attempt, we fix a parameter $\alpha \in \mathbb{R}$ with negative values of α producing negative correlations and positive values producing positive correlations. The parameter α is tuned to achieve a correlation close to the desired value ρ . This is achieved with a binary search, starting with $\alpha_{\min} = -60$ and $\alpha_{\max} = 60$ and stopping once the desired precision is achieved or after the obtained correlation level stops improving. In some cases it is impossible to reach the correlation level required by the user, in which case the closest possible correlation is produced.

In each pairing attempt with parameter $\alpha \in \mathbb{R}$, we iteratively assign degrees of non-outliers to elements as follows. Recall that the degree sequence $\hat{\mathbf{d}}_{\hat{n}}$ is sorted with \hat{d}_1 being the maximum degree. Starting with $i = 1$, let R_i be the collection of unassigned elements at step i . At step i , choose element v with probability proportional to η_v^α from the set of elements in R_i satisfying

$$\hat{d}_i \leq \frac{\eta_v}{1 - \xi\phi} \cdot \min \left\{ |C_k| - 1 : j \in C_k \right\}, \quad (3)$$

where recall that η_v is the number of communities element v belongs to. The chosen element v gets paired with the i th degree in the subsequence $\hat{\mathbf{d}}_{\hat{n}}$, namely, the degree \hat{d}_i . If the set defined by the condition (3) is empty, then we sample from the elements in R_i for which the value of $(\eta_v/(1 - \xi\phi)) \cdot \min\{|C_k| - 1 : j \in C_k\}$ maximal.

Once the pairing is complete, we relabel as follows. The nodes are labelled as $[n]$, with node i having degree d_i . The communities $(C_j, j \in L)$ now contain non-outlier nodes corresponding to degree/element pairing. If d_i was assigned a non-outlier, its corresponding node i belongs to some number η_i of communities, one of which is its primary community. If d_i was assigned to an outlier, its corresponding node i is assigned to an auxiliary “community” labelled C_0 ; the set C_0 is merely a label for the outliers and such an outlier i has $\eta_i = 0$.

Phase 5: creating edges.

At this point there are n nodes with labels from $[n]$; $\hat{n} = n - s_0$ of them are non-outliers and the remaining ones are outliers. There is also a family of overlapping communities with each non-outlier node $i \in [n]$ belonging to $\eta_i \geq 1$ communities. Finally, each node $i \in [n]$ (either outlier or non-outlier) is assigned a degree d_i which we interpret as a set of d_i unpaired half-edges. The remaining 2 steps construct the edges.

For each non-outlier $i \in [n]$ we split the d_i half-edges of i into *community* half-edges and *background* half-edges. To this end, define $Y_i := \lfloor (1 - \xi)d_i \rfloor$ and $Z_i := d_i - Y_i$ (note that Y_i and Z_i are random variables with $\mathbb{E}[Y_i] = (1 - \xi)d_i$ and $\mathbb{E}[Z_i] = \xi d_i$) and, for all non-outliers $i \in [n]$, split the d_i half-edges of i into Y_i community half-edges and Z_i background half-edges. Community half-edges are further split into the η_i communities the non-outlier node i belongs to, as evenly as possible. Specifically, for the communities containing node i , $Y_i - \eta_i \lfloor Y_i / \eta_i \rfloor$ communities (chosen randomly) each receive $\lfloor Y_i / \eta_i \rfloor + 1$ half-edges and the remaining communities each receive $\lfloor Y_i / \eta_i \rfloor$ half-edges. On the other hand, if $i \in [n]$ is an outlier then we set $Z_i = d_i$.

Once the assignment of degrees is complete, for each $j \in [L]$, we independently construct the *community graph* $G_{n,j}$ as per the configuration model on node set C_j and the corresponding degree sequence. In the event that the sum of degrees in a community is odd, we pick a maximum degree

node i in said community and decrease its community degree by one while increasing its background graph degree by one. Finally, construct the *background graph* $G_{n,0}$ as per the configuration model on node set $[n]$ and degree sequence $(Z_i, i \in [n])$. Let $G_n = \bigcup_{0 \leq j \leq n} G_{n,j}$ be the union of all graphs generated in this phase.

Phase 6: rewiring self-loops and multi-edges.

Note that, although we are calling $G_{n,0}, G_{n,1}, \dots, G_{n,L}$ *graphs*, they are in fact *multi-graphs* at the end of phase 5. To ensure that G_n is simple, we perform a series of *rewirings* in G_n . A rewiring takes two edges as input, splits them into four half-edges, and creates two new edges distinct from the input. We first rewire each community graph $G_{n,j}$ ($j \in [L]$), and the background graph $G_{n,0}$, independently as follows.

1. For each edge $e \in E(G_{n,j})$ that is a loop, we add e to a *recycle* list that is assigned to $G_{n,j}$. Similarly, if $e \in E(G_{n,j})$ contributes to a multi-edge, we put all but one copies of this edge to the *recycle* list.
2. We shuffle the *recycle* list and, for each edge e in the list, we choose another edge e' uniformly from $E(G_{n,j}) \setminus \{e\}$ (not necessarily in the list) and attempt to rewire these two edges. We save the result only if the rewiring does not lead to any further self-loops or multi-edges, otherwise we give up. In either case, we then move to the next edge in the *recycle* list.
3. After we attempt to rewire every edge in the *recycle* list, we check to see if the new *recycle* list is smaller. If yes, we repeat step 2 with the new list. If no, we give up and move all of the “bad” edges from the community graph to a collective *global recycle* list.

As a result, after ignoring edges in the *global recycle* list, all community graphs are simple and the background graph is simple. However, as is the case in the original **ABCD** model, an edge in the background graph can form a multi-edge with an edge in a community graph. Another problem that might occur, specific to **ABCD+o²** model, is that an edge from one community can form a multi-edge with an edge from a different but overlapping community. All of these problematic edges are added to the *global recycle* list. We merge all community graphs with the background graph. Finally, the *global recycle* list is transformed into a list of half-edges and new edges are created from it. We follow the same procedure as for the community graphs. However, we do not “give up” recycling and follow the process until all required edges are created. As the background graph is sparse, this final rewiring is very fast in practice.

3 Properties of the ABCD+o² Model

In this section, we present experiments highlighting the properties of the **ABCD+o²** model. We begin by investigating the communities generated by the reference layer as described in Section 3.1. We then continue with additional features of the model in Section 3.2. We are mainly interested in comparing the model’s properties with properties of real-world networks. Thus, for the model parameters, we use values derived from several networks that have known overlapping communities. The networks were presented in [37], and are available on the Stanford Network Analysis Project (SNAP) website⁴.

⁴<https://snap.stanford.edu/data/index.html#communities>

Parameter	Description	DBLP	Amazon	YouTube
n	Number of nodes	317,080	334,863	52,675
s_0	Number of outliers	56,082	17,669	0
η	Average number of communities a non-outlier node is part of	2.76	7.16	2.45
ρ	Pearson correlation between the degree of nodes and the number of communities they belong to	0.76	0.22	0.37
γ	Power-law degree distribution with exponent γ	2.30	3.04	1.87
δ	Min degree at least δ	5	5	5
Δ	Max degree at most Δ	343	549	1,928
β	Power-law community size distribution with exponent γ	1.88	2.03	2.13
s	Min community size at least s	10	10	10
S	Max community size at most S	7,556	53,551	3,001
ξ	Level of noise	0.11	0.11	0.59

Table 2: The parameters used for generating **ABCD+o**² graphs based on the DBLP, Amazon, and YouTube datasets.

DBLP: a network where nodes are authors and an edge exists between two authors if they published a paper together. The ground-truth communities are defined by the publication venues of the journals in which the authors published.

Amazon: a network where nodes are products and an edge exists between two products if they are frequently co-purchased. The ground-truth communities are defined by product categories.

YouTube: a network where nodes are YouTube channels and an edge exists between two channels if they are friends on the platform. The ground-truth communities are user-generated social groups. In the full network over 95% of nodes are outliers and the level of noise ξ is 0.96. We consider the subgraph induced by the non-outlier nodes to avoid such extreme values.

For each network, all parameters other than δ and s are measured directly from the data. We set $\delta = 5$ to avoid cases where nodes have no neighbours in their own community, and insist on $s \geq 10$. The power-law exponents γ and β are fit using maximum-likelihood estimation described in [12] and implemented in the Python package **powerlaw**⁵[1]. In this fitting process, we set the minimum degree as $\delta = 5$ and the minimum community size as $s = 10$ for consistency. The level of noise ξ is computed directly from the graphs as the fraction $|E_b|/|E|$ where E_b is the set of edges where the incident nodes share no communities. The parameters for each network are shown in Table 2.

Additionally, some of the experiments are affected by the dimension d of the hidden reference layer. In such cases, we compare dimensions 2, 8, and 64.

⁵<https://github.com/jeffalstott/powerlaw>

3.1 Effects of the Reference Layer

The new and novel feature of the **ABCD+o²** model is the hidden reference layer used to generate overlapping communities. Here, we analyze the effects of this reference layer on the community overlap sizes. For comparison, we consider the **CKB** model [11] following a description from [31]. The **CKB** model is an application of the random bipartite community affiliation graph proposed for overlapping **LFR** [26] that uses a power-law distribution for the number of communities per node. Let us briefly summarize this graph construction process. The community affiliation graph is a bipartite graph where one part consists of nodes and the other consists of communities. As input, the **CKB** model requires the number of nodes n , a truncated power-law distribution $\mathcal{P}(\Omega, x_{\min}, x_{\max})$ for the number of communities per node, and another truncated power-law distribution $\mathcal{P}(\beta, s, S)$ for the community sizes. The number of communities is set to

$$\left\lfloor \frac{n \cdot \mathbb{E}[\mathcal{P}(\Omega, x_{\min}, x_{\max})]}{\mathbb{E}[\mathcal{P}(\beta, s, S)]} \right\rfloor.$$

First, each node in the bipartite graph is assigned a number of half-edges sampled from $\mathcal{P}(\Omega, x_{\min}, x_{\max})$, and likewise each community is assigned a number of half-edges sampled from $\mathcal{P}(\beta, s, S)$. A small adjustment is made to ensure that the number of half-edges in both parts is equal. Lastly, the nodes are matched with the communities based on the bipartite configuration model. Thus, each community is populated with nodes, and each node is assigned to some number of communities.

Community sizes.

In the first experiment (see Figure 2), we compare the distributions of community sizes. Of course, since both **ABCD+o²** and **CKB** sample community sizes directly from the specified distribution, it is unsurprising that the empirical distributions are near-perfect fits. Thus, this first experiment merely acts as a sanity check ensuring that the **ABCD+o²** and **CKB** models generate community sizes correctly. Moreover, this experiment verifies that the distributions coming from the real networks are indeed power-law.

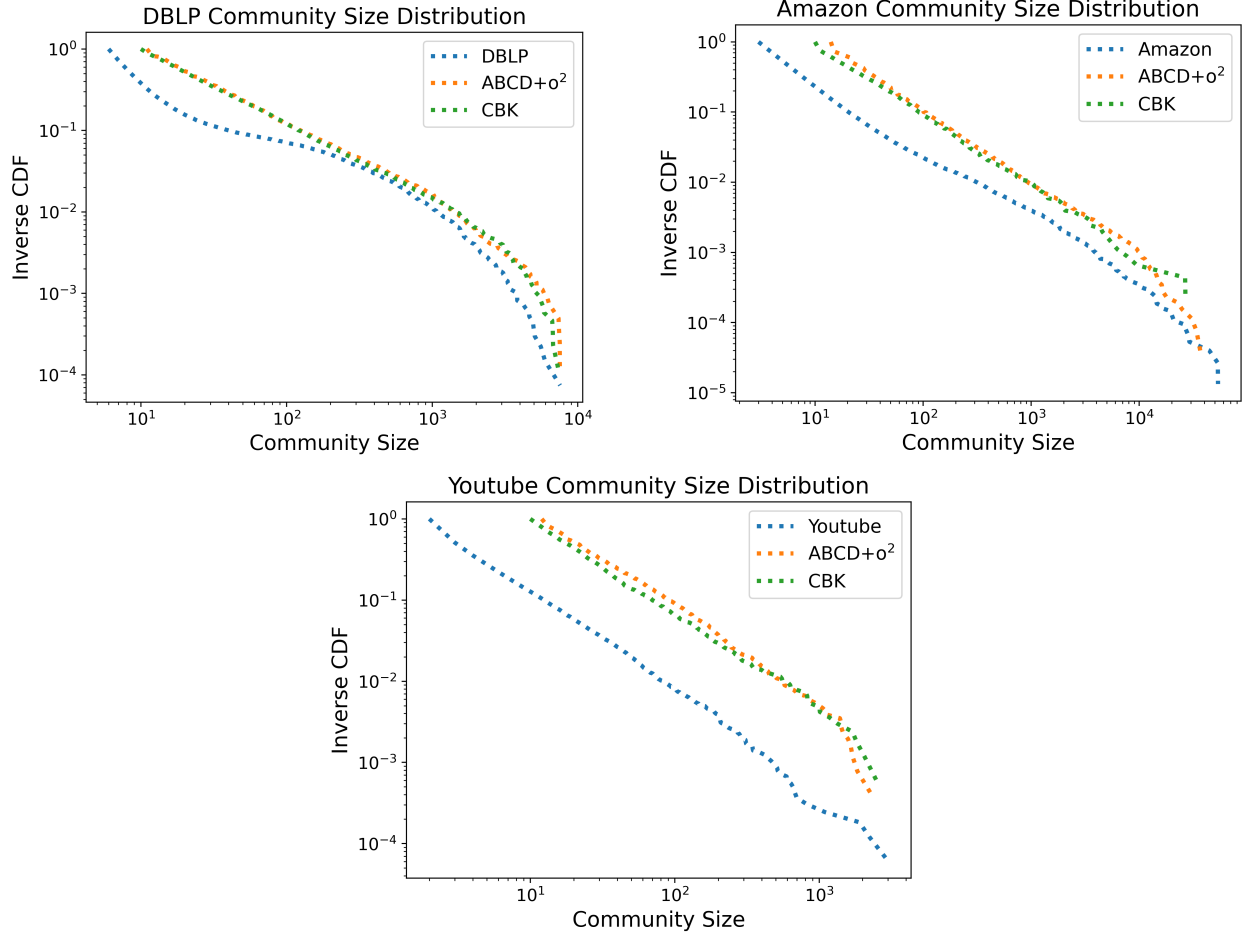


Figure 2: Distribution of community sizes for the three real-world networks and their synthetic counterparts.

Number of communities per node.

Results are presented in Figure 3. In contrast with the previous experiment, there is less agreement between the two models and the real distributions. In the case of the YouTube and DBLP networks, the power-law distribution generated by the **CBK** model is a good fit for the empirical distribution, whereas the Amazon distribution does not appear to be power-law. In contrast, the $\text{ABCD}+\mathbf{o}^2$ model is able to produce a distribution similar to each of the three networks with a properly tuned dimension parameter d ; for the Amazon graph a low dimensional reference layer captures the behaviour well, and for the DBLP and YouTube graphs a high dimensional reference layer captures the behaviour well. The flexibility and realistic distributions are even more encouraging when we recall that the distribution is not given to the $\text{ABCD}+\mathbf{o}^2$ model, and is instead a natural feature of the hidden reference layer.

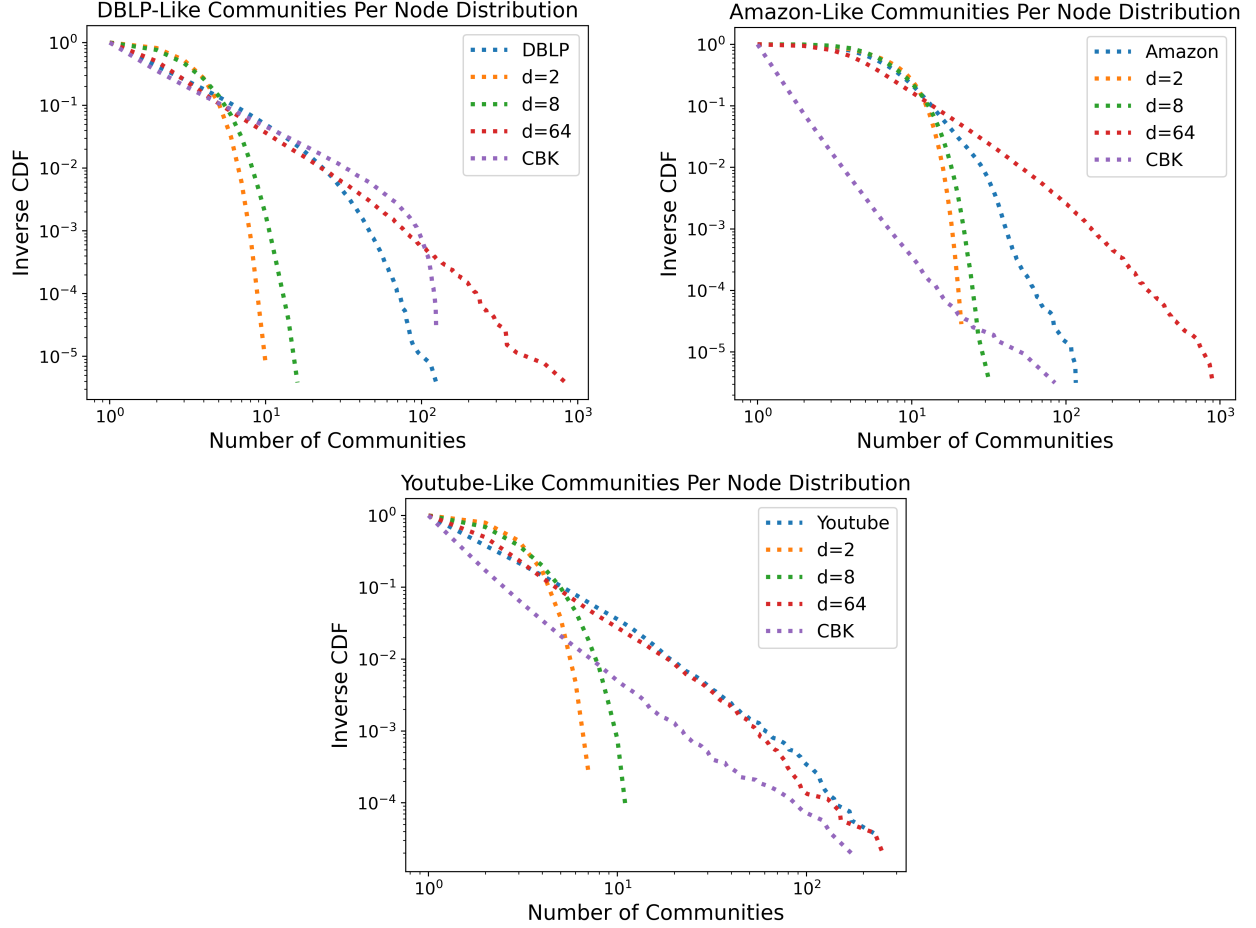


Figure 3: Distribution of the number of communities per node.

Community intersection sizes.

Results are presented in Figure 4. We examine the size of the overlaps produced by each model. Neither the **ABCD**+ \mathbf{o}^2 model nor the **CKB** model specifies a particular distribution, but we can see the **ABCD**+ \mathbf{o}^2 model, with a well-tuned dimension parameter, is able to create a similar distribution to the real graph. In contrast, the **CKB** model does not fit the Amazon and YouTube graphs well. Specifically, the **CKB** model produces too many small intersections and, moreover, the largest intersection is several orders of magnitude too small. We consider this experiment strong evidence for using a reference layer as the default option for generating overlapping communities.

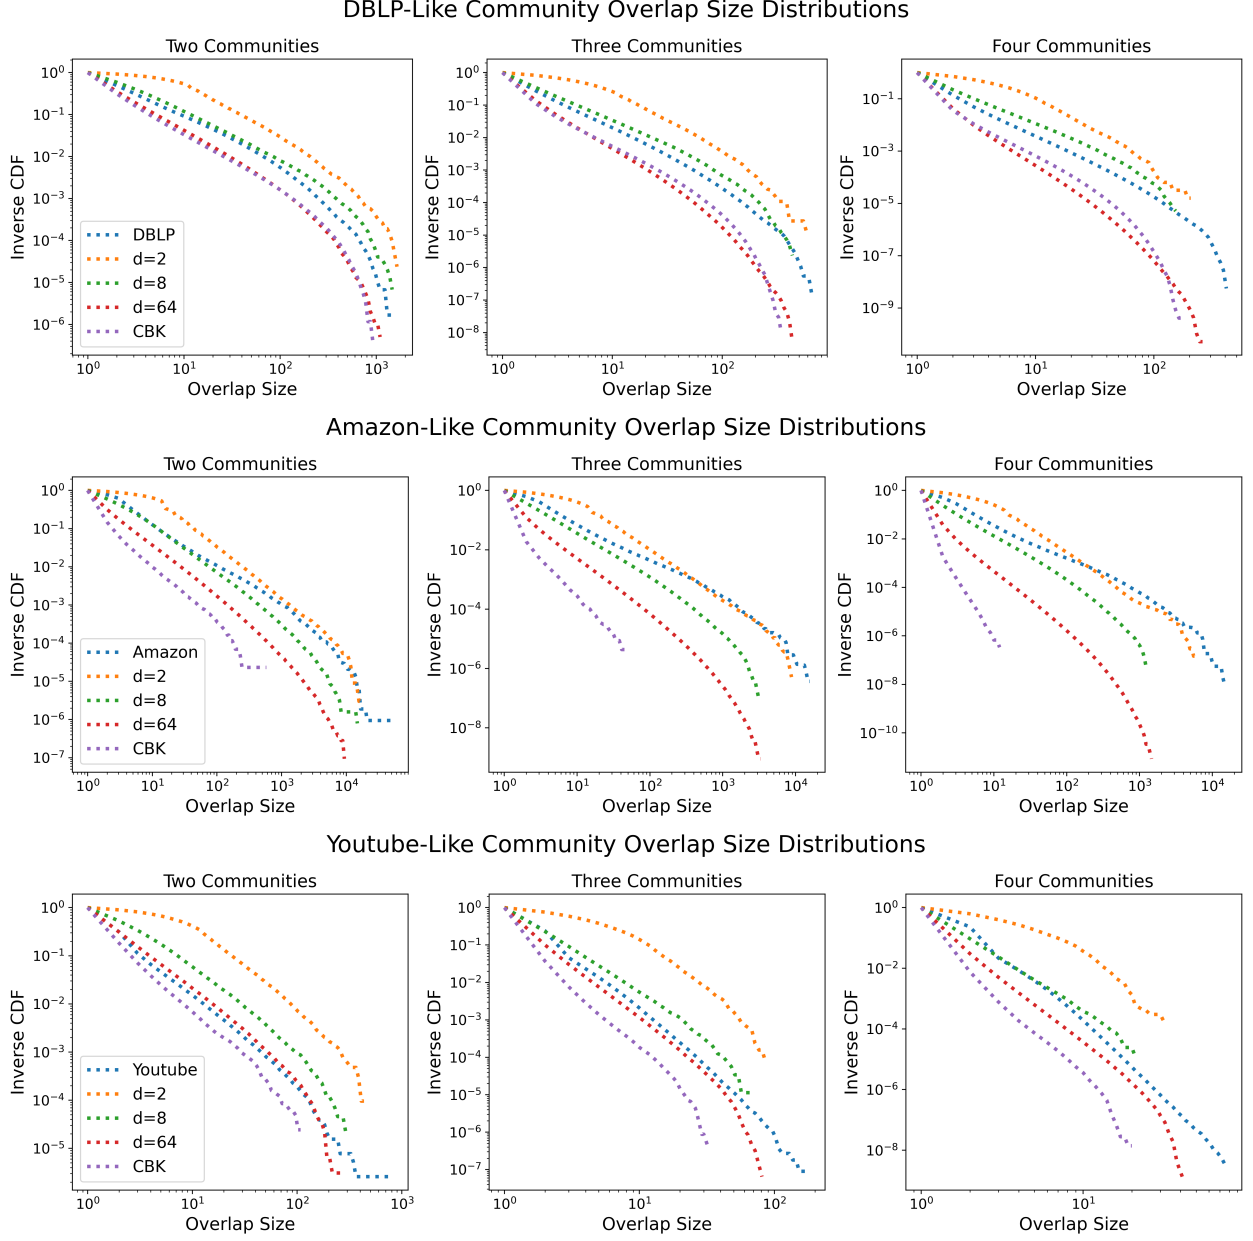


Figure 4: Distribution of the size of overlaps.

3.2 Additional properties of the model.

We now turn to properties of the **ABCD+o²** model beyond those governed by the hidden reference layer.

Graph	Empirical	$d = 2$	$d = 8$	$d = 64$
DBLP	0.76	0.43	0.56	0.68
Amazon	0.22	0.20	0.19	0.20
YouTube	0.37	0.37	0.37	0.38

Table 3: Measured ρ in both the real networks and the corresponding **ABCD**+ \mathbf{o}^2 graphs.

Degree vs. number of communities.

The **ABCD**+ \mathbf{o}^2 model attempts to produce a correlation ρ between the degree of a node and the number of communities it belongs to. In Table 3 we report the desired correlation ρ for each of the networks, as well as the achieved correlation for the **ABCD**+ \mathbf{o}^2 graphs. Note that the extreme positive correlation found in the DBLP network cannot be matched by the model, although it can still achieve a strong positive correlation. The maximum achievable correlation increases with the dimension as there are more nodes belonging to a large number of communities (see Figure 3).

Density of intersections

One consequence of forcing nodes of large degree into community intersections is that those intersections will be denser than the individual communities. We show this phenomenon in Figure 5 by comparing the community and intersection densities for various ρ . Note that, given communities C_j and C_k with $s_j \gg s_k$, the density of C_k is nearly always larger than that of C_j since both community graphs are sparse (at most a $1 - \xi$ fraction of each node’s degree is delegated to one of its communities). Figures for the other graphs and dimensions can be found in Appendix A. For computational reasons, we only consider overlaps of size at least 25, and to ensure we are not simply measuring the density of the smaller community, the size of the overlap must be at most half the size of the smaller community.




Figure 5: Community overlap densities compared to individual community densities for various values of ρ . The line corresponds to the median, and the shaded region is the 25th to 75th quantile. We only consider pairs of communities where the overlap is at least 50 nodes and the size of the overlap is less than half the size of the smaller community.

Having community intersections that are denser than the respective communities is a desirable property and one of our main motivators for introducing the parameter ρ to the **ABCD+o²** model. In [38] it was found that, in real networks, community intersections can be significantly denser than the individual communities. In Figure 6 we confirm that the DBLP, Amazon, and YouTube networks all exhibit this phenomenon. Moreover, we compare these densities with those coming from the corresponding **ABCD+o²** graphs. Aside from confirming that the model mimics the behaviour of real networks, we also find that the dimension of the reference layer affects the density of both the communities and the intersections. Moreover, we see that a low-dimensional reference layer yields a better approximation of the real data in the case of Amazon, whereas a high-dimensional reference layer yields a better approximation in the cases of DBLP and YouTube. Note that this high-dimensional vs. low-dimensional discrepancy is consistent with the experiment presented in Figure 3.

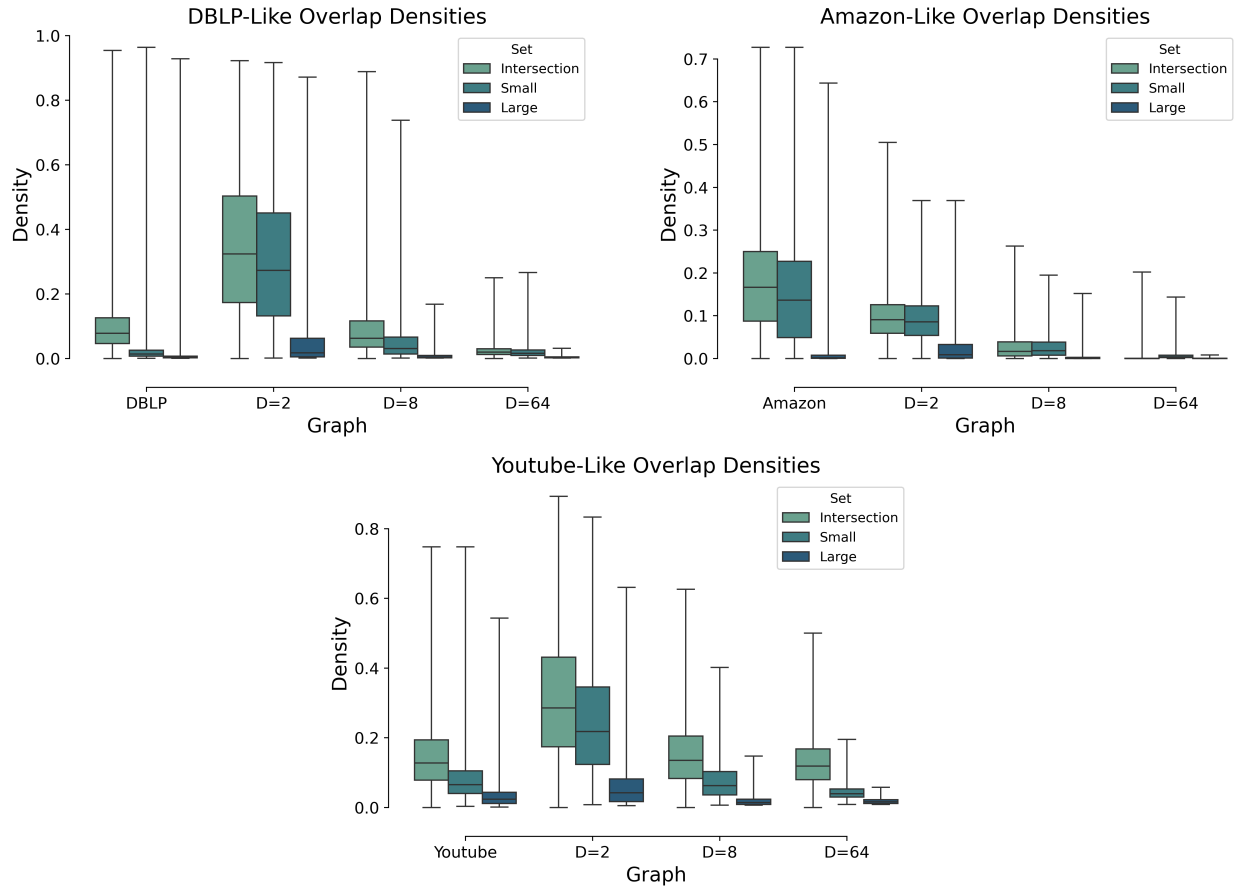


Figure 6: Density of the Overlap compared to the density of the communities. We restrict to overlaps of size at least 10 due to the discrete and unpredictable nature of small communities.

Internal edge fraction

Finally, we analyze how strongly nodes are connected to their communities. There are many measures for community association strength, some of which we recently analyzed in [4]. For the

sake of simplicity, the measure we will use here is the internal edge fraction $\text{IEF}(v, C)$, defined as

$$\text{IEF}(v, C) = \frac{|\{\{u, v\} \in E : u \in C\}|}{\deg(v)}.$$

For a graph G , a collection of communities $\{C_i, i \in [k]\}$, and a node v , let $C_1^{(v)}, \dots, C_k^{(v)}$ be an ordering of the communities such that

$$\text{IEF}(v, C_1^{(v)}) \geq \text{IEF}(v, C_2^{(v)}) \geq \dots \geq \text{IEF}(v, C_k^{(v)}).$$

Then, if v is a member of ℓ communities, we would expect a jump in value between $\text{IEF}(v, C_\ell^{(v)})$ and $\text{IEF}(v, C_{\ell+1}^{(v)})$. In Figure 7 we show the range of the top five internal edge fractions, sorted in descending order, on DBLP (left) and a generated **ABCD+o²** graph (right), binned by the number of communities the node v is in. As predicted, a node in ℓ communities tends to have ℓ relatively large internal edge fractions. Furthermore, despite the **ABCD+o²** model splitting the community degree of a node v evenly among its ℓ communities, it is possible for v to have a higher internal edge fraction into community C if either some of v 's edges lies in an intersection containing C or some of v 's background edges land in C . Thus, we see a pattern in **ABCD+o²** graphs where there are ℓ large, but not equal, internal edge fraction values. This feature is also present in the DBLP graph, the other real graphs, and the other **ABCD+o²** graphs (see Appendix B for a complete collection of figures).

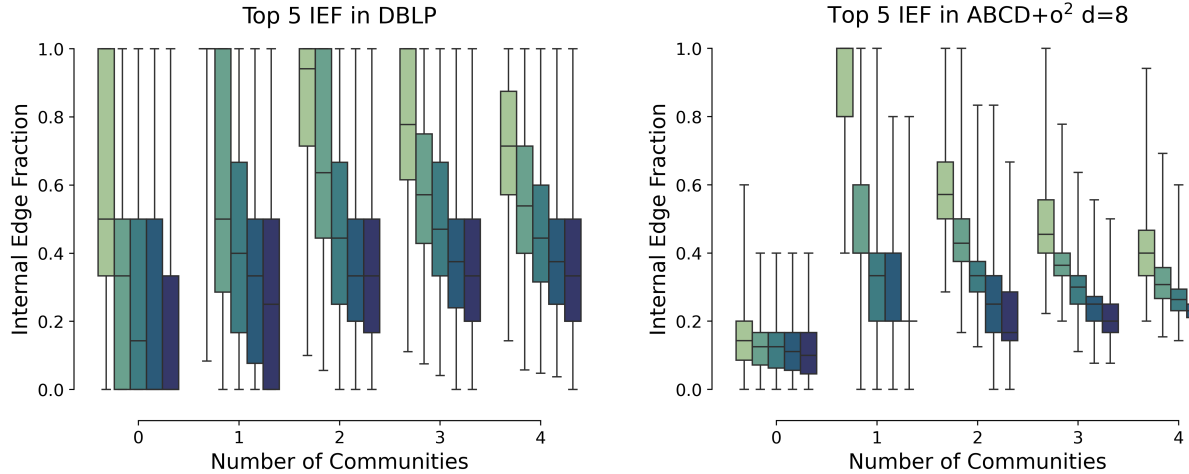


Figure 7: Five largest internal edge fractions, sorted in descending order, grouped by the number of communities.

4 Benchmarking Community Detection Algorithms

The main purpose of having synthetic models with ground-truth community structure is to test, tune, and benchmark community detection algorithms. To showcase **ABCD+o²** in this light, we use the model to evaluate the performance of four community detection algorithms. The algorithms are as follows.

Leiden [34]: a greedy algorithm that attempts to optimize the modularity function. Note that this algorithm returns a partition, and we use it merely as a baseline to compare with algorithms that attempt to find overlapping communities.

Edge Clustering [25]: an edge-partitioning algorithm that translates to overlapping clusters of nodes. Pairs of edges are measured based on similarity of neighbourhoods, and these similarity measures dictate the order in which edge-communities merge. As edge-communities merge, the modularity is tracked on the line-graph, and the maximum modularity attained yields the edge-communities, which in turn yields overlapping node-communities.

Ego-Split [14]: a method which finds overlapping clusters in a graph G by applying a partitioning algorithm such as Leiden to an auxiliary graph G' and mapping the resulting partition onto G . The auxiliary graph G' is constructed from G by creating multiple copies, or “egos”, of each node based on its neighbourhood.

Ego-Split+IEF [4]: the same algorithm as Ego-Split, but with a post-processing step that re-assigns nodes to communities based on the IEF measure.

This is by no means an exhaustive list of community detection algorithms. Moreover, we use our own straightforward Python implementations of the above algorithms which may not be optimal. We wish only to showcase the usefulness of **ABCD+o²** in comparing various algorithms related to the community structure of graphs.

In the first experiment, we compare the accuracy of each algorithm with respect to the ground-truth communities. The measure we use to determine accuracy is the overlapping Normalized Mutual Information (oNMI) measure [28]. For this experiment we fix the number of nodes $n = 5,000$, the minimum degree $\delta = 5$, the maximum degree $\Delta = 100$, the minimum community size $s = 50$, and the maximum community size $S = 500$. The parameters of **ABCD+o²** with the most influence on the quality of detection algorithms are ξ (the level of noise) and η (the average number of communities a non-outlier is part of), and we consider every combination of $\xi \in \{0.1, 0.2, \dots, 0.6\}$ and $\eta \in \{1, 1.25, \dots, 2.5\}$. We fix the remaining parameters according to the YouTube graph, namely, $\gamma = 1.87$, $\beta = 2.13$, and $\rho = 0.37$.

Figure 8 summarizes the results of the experiment. We see that the **Leiden** algorithm performs the best when $\eta = 1$, whereas for $\eta > 1$ the **Ego-Split+IEF** algorithm performs better. Given that the **Leiden** algorithm finds a partition of the nodes, it is unsurprising that it quickly depreciates in quality as η increases. We also see a general trend of all algorithms performing worse as the graph gets noisier, either by increasing ξ or η . This trend is highlighted separately in Figure 9 where we consider only results with $\eta = 1.25$ (left) or $\xi = 0.1$ (right). From numerous and varying tests, we have found that, in general, increasing η is far more damaging to community detection algorithms than increasing ξ .




Figure 8: Comparing oNMI of four clustering algorithms on $\mathbf{ABCD}+\mathbf{o}^2$ graphs with 5,000 nodes.

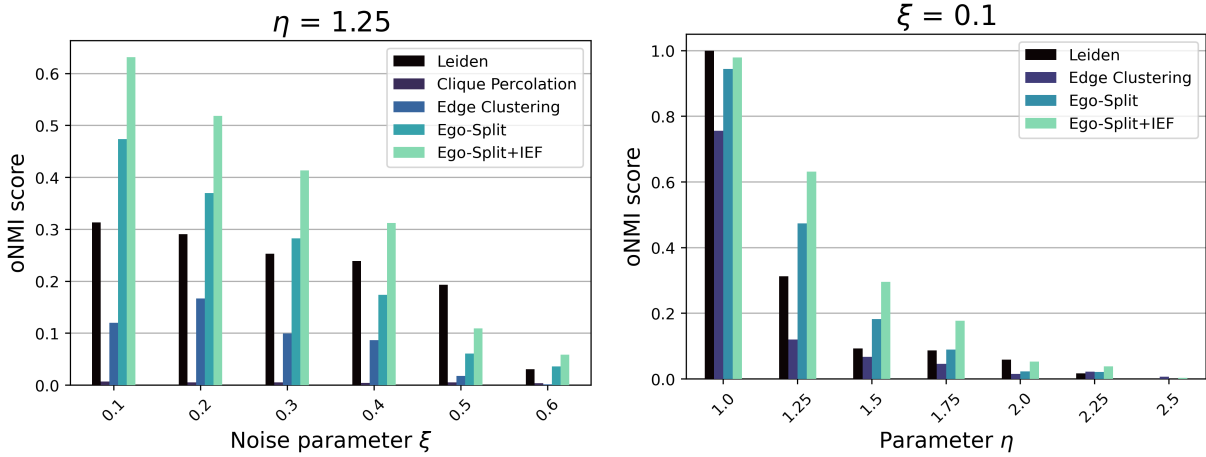


Figure 9: Comparing oNMI of four clustering algorithms on $\mathbf{ABCD}+\mathbf{o}^2$ graphs with 5,000 nodes. In the left plot, we fix $\eta = 1.25$ and vary ξ . In the right plot, we fix $\xi = 0.1$ and vary η .

Another important aspect of clustering algorithms is their performance when running on large graphs. In our second experiment, we fix $\eta = 1.25$ and $\xi = 0.1$ and we vary the number of nodes from $n = 1,000$ to $n = 15,000$; the parameters other than n, η and ξ are the same as in the previous experiment. The results are presented in Figure 10. We see that all algorithms trying to recover overlapping communities are significantly slower than the simple **Leiden** partitioning algorithm. We also see that adding the IEF-based post-processing stage to the Ego-Splitting algorithm is negligible in terms of performance. Together with the experiment presented in Figure 8, this negligible slowdown suggests that including a post-processing stage to an overlapping community detection algorithm yields a worthwhile improvement.

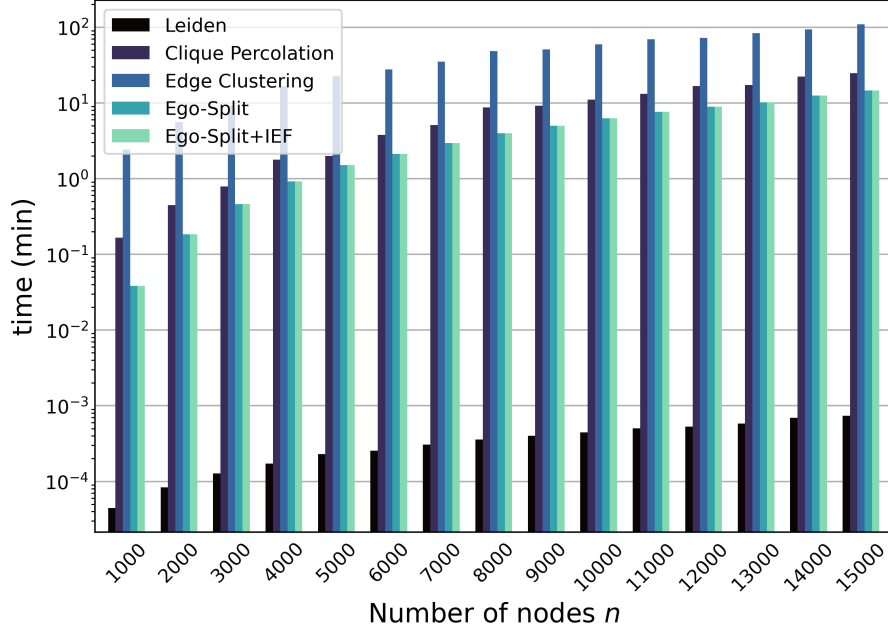


Figure 10: The performance of four clustering algorithms on **ABCD+o²** graphs with 1,000 to 15,000 nodes.

5 Conclusion

We presented **ABCD+o²**: a generalization of the **ABCD+o** model that allows for overlapping communities. We proposed a novel mechanism for producing overlapping communities, a hidden reference layer, that requires two new parameters: the average number of communities per node η , and the dimension for the reference layer d . Furthermore, we introduced a parameter ρ that controls the correlation between degree and number of communities for non-outlier nodes. In a series of experiments, we compared the properties of the **ABCD+o²** model to those of real graphs using parameters measured from the real graph. We found that the hidden reference layer can accurately create a structure of overlapping communities similar to those in the real graph. We also found that the model can successfully create an empirical ρ close to the requested value (except for extreme cases). Finally, we showcased the model’s ability to benchmark community detection algorithms and compare their quality.

This paper acts as a first step in our study of the **ABCD+o²** model and gives way to several open problems for future research. First, a theoretical description of the distribution of the number of communities per node and the effect of the dimension of the reference layer are clear goals. A similar direction is understanding the effects of modifying the geometry of the reference layer. For example, one could estimate the reference space of a network from auxiliary data and use it for a corresponding parameter-fitted **ABCD+o²** model. Using this fitting procedure, can we achieve additional structural similarities to real networks?

We are also interested in results for the **ABCD+o²** model that generalize results of the **ABCD** and **ABCD+o** models. In [20] the modularity was studied and it was found that the maximum modularity came from the ground truth communities until a certain level of noise, after which a higher modularity could be attained. A similar behaviour should be seen with **ABCD+o²** and its overlap parameter η , although we would need to select an appropriate generalization of modularity for overlapping sets, such as the generalization proposed in [13]. Additionally, in [6] it was shown that **ABCD** graphs exhibit self-similar behaviour, namely, the degree distributions of communities are asymptotically the same as the degree distribution of the whole graph (up to an appropriate normalization). We suspect that this self-similar property persists in **ABCD+o²**.

References

- [1] Jeff Alstott, Ed Bullmore, and Dietmar Plenz. powerlaw: A python package for analysis of heavy-tailed distributions. *PLoS ONE*, 9(1), 2014. doi:10.1371/journal.pone.0085777.
- [2] Samin Aref, Mahdi Mostajabdaveh, and Hriday Chheda. Bayan algorithm: Detecting communities in networks through exact and approximate optimization of modularity. *Phys. Rev. E*, 110:044315, 2024. doi:10.1103/PhysRevE.110.044315.
- [3] Albert-László Barabási. *Network science*. Cambridge University Press, 2016. URL: <http://barabasi.com/networksciencebook/>.
- [4] Jordan Barrett, Ryan DeWolfe, Bogumił Kamiński, Paweł Prałat, Aaron Smith, and François Théberge. Improving community detection via community association strength scores. In Mindaugas Bloznelis, Paulius Drungilas, Bogumił Kamiński, Paweł Prałat, Matas Šileikis, François Théberge, and Rimantas Vaicekauskas, editors, *Modelling and Mining Networks*, pages 81–95, Cham, 2025. Springer Nature Switzerland. doi:10.1007/978-3-031-92898-7_6.
- [5] Jordan Barrett, Bogumił Kamiński, Paweł Prałat, and François Théberge. Self-similarity of communities of the ABCD model. *Theoretical Computer Science*, 1026:115012, 2025. doi:10.1016/j.tcs.2024.115012.
- [6] Edward A Bender and E Rodney Canfield. The asymptotic number of labeled graphs with given degree sequences. *Journal of Combinatorial Theory, Series A*, 24(3):296–307, 1978. doi:10.1016/0097-3165(78)90059-6.
- [7] Jon Louis Bentley. Multidimensional binary search trees used for associative searching. *Commun. ACM*, 18(9):509–517, September 1975. doi:10.1145/361002.361007.

- [8] Marian Boguna, Ivan Bonamassa, Manlio De Domenico, Shlomo Havlin, Dmitri Krioukov, and M Ángeles Serrano. Network geometry. *Nature Reviews Physics*, 3(2):114–135, 2021. doi:10.1038/s42254-020-00264-4.
- [9] Béla Bollobás. A probabilistic proof of an asymptotic formula for the number of labelled regular graphs. *European Journal of Combinatorics*, 1(4):311–316, 1980. doi:10.1016/S0195-6698(80)80030-8.
- [10] Kyrylo Chykhraze, Anton Korshunov, Nazar Buzun, Roman Pastukhov, Nikolay Kuzyurin, Denis Turdakov, and Hangkyu Kim. Distributed generation of billion-node social graphs with overlapping community structure. In Pierluigi Contucci, Ronaldo Menezes, Andrea Omicini, and Julia Poncela-Casasnovas, editors, *Complex Networks V*, pages 199–208, Cham, 2014. Springer International Publishing. doi:10.1007/978-3-319-05401-8_19.
- [11] Aaron Clauset, Cosma Rohilla Shalizi, and M. E. J. Newman. Power-law distributions in empirical data. *SIAM Review*, 51(4):661–703, 2009. doi:10.1137/070710111.
- [12] J. Chitra Devi and E. Poovammal. An analysis of overlapping community detection algorithms in social networks. *Procedia Computer Science*, 89:349–358, 2016. Twelfth International Conference on Communication Networks, ICCN 2016, August 19– 21, 2016, Bangalore, India Twelfth International Conference on Data Mining and Warehousing, ICDMW 2016, August 19-21, 2016, Bangalore, India Twelfth International Conference on Image and Signal Processing, ICISP 2016, August 19-21, 2016, Bangalore, India. doi:10.1016/j.procs.2016.06.082.
- [13] Alessandro Epasto, Silvio Lattanzi, and Renato Paes Leme. Ego-splitting framework: From non-overlapping to overlapping clusters. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 145–154, 2017. doi:10.1145/3097983.3098054.
- [14] Santo Fortunato. Community detection in graphs. *Physics reports*, 486(3-5):75–174, 2010. doi:10.1016/j.physrep.2009.11.002.
- [15] Santo Fortunato and Marc Barthelemy. Resolution limit in community detection. *Proceedings of the national academy of sciences*, 104(1):36–41, 2007. doi:10.1073/pnas.0605965104.
- [16] Michelle Girvan and Mark EJ Newman. Community structure in social and biological networks. *Proceedings of the national academy of sciences*, 99(12):7821–7826, 2002. doi:10.1073/pnas.122653799.
- [17] Adam Douglas Henry, Paweł Prałat, and Cun-Quan Zhang. Emergence of segregation in evolving social networks. *Proceedings of the National Academy of Sciences*, 108(21):8605–8610, 2011. doi:10.1073/pnas.1014486108.
- [18] Bogumił Kamiński, Tomasz Olczak, Bartosz Pankratz, Paweł Prałat, and François Théberge. Properties and performance of the ABCDe random graph model with community structure. *Big Data Research*, 30:100348, 2022. doi:10.1016/j.bdr.2022.100348.
- [19] Bogumił Kamiński, Bartosz Pankratz, Paweł Prałat, and François Théberge. Modularity of the ABCD random graph model with community structure. *Journal of Complex Networks*, 10(6):cnac050, 2022. doi:10.1093/comnet/cnac050.

- [20] Bogumił Kamiński, Paweł Prałat, and François Thériberge. Artificial benchmark for community detection (ABCD)—fast random graph model with community structure. *Network Science*, pages 1–26, 2021. doi:10.1017/nws.2020.45.
- [21] Bogumił Kamiński, Paweł Prałat, and François Thériberge. *Mining Complex Networks*. Chapman and Hall/CRC, 2021. doi:10.1201/9781003218869.
- [22] Bogumił Kamiński, Paweł Prałat, and François Thériberge. Artificial benchmark for community detection with outliers (ABCD+o). *Applied Network Science*, 8(1):25, 2023. doi:10.1007/s41109-023-00552-9.
- [23] Bogumił Kamiński, Paweł Prałat, and François Thériberge. Hypergraph artificial benchmark for community detection (h-ABCD). *Journal of Complex Networks*, 11(4):cnad028, 2023. doi:10.1093/comnet/cnad028.
- [24] Paul Kim and Sangwook Kim. Detecting overlapping and hierarchical communities in complex network using interaction-based edge clustering. *Physica A: Statistical Mechanics and its Applications*, 417:46–56, 2015. doi:10.1016/j.physa.2014.09.035.
- [25] Andrea Lancichinetti and Santo Fortunato. Benchmarks for testing community detection algorithms on directed and weighted graphs with overlapping communities. *Physical Review E*, 80(1):016118, 2009. doi:10.1103/PhysRevE.80.016118.
- [26] Andrea Lancichinetti, Santo Fortunato, and Filippo Radicchi. Benchmark graphs for testing community detection algorithms. *Physical review E*, 78(4):046110, 2008. doi:10.1103/PhysRevE.78.046110.
- [27] Aaron F McDaid, Derek Greene, and Neil Hurley. Normalized mutual information to evaluate overlapping community finding algorithms. *arXiv preprint arXiv:1110.2515*, 2011. doi:10.48550/arXiv.1110.2515.
- [28] Günce Keziban Orman and Vincent Labatut. A comparison of community detection algorithms on artificial networks. In *Discovery Science: 12th International Conference, DS 2009, Porto, Portugal, October 3-5, 2009 12*, pages 242–256. Springer, 2009. doi:10.1007/978-3-642-04747-3_20.
- [29] Gergely Palla, Imre Derényi, Illés Farkas, and Tamás Vicsek. Uncovering the overlapping community structure of complex networks in nature and society. *Nature*, 435(7043):814–818, 2005. doi:10.1038/nature03607.
- [30] Neha Sengupta, Michael Hamann, and Dorothea Wagner. Benchmark generator for dynamic overlapping communities in networks. In *2017 IEEE International Conference on Data Mining (ICDM)*, pages 415–424, 2017. doi:10.1109/ICDM.2017.51.
- [31] M Ángeles Serrano and Marián Boguñá. *The Shortest Path to Network Geometry: A Practical Guide to Basic Models and Applications*. Cambridge University Press, 2021. doi:10.1017/9781108865791.
- [32] M Ángeles Serrano, Dmitri Krioukov, and Marián Boguñá. Self-similarity of complex networks and hidden metric spaces. *Physical review letters*, 100(7):078701, 2008. doi:10.1103/PhysRevLett.100.078701.

- [33] V. Traag, L. Waltman, and Nees Jan van Eck. From Louvain to Leiden: guaranteeing well-connected communities. *Scientific Reports*, 9:5233, 03 2019. doi:10.1038/s41598-019-41695-z.
- [34] N. C. Wormald. *Models of Random Regular Graphs*, page 239–298. London Mathematical Society Lecture Note Series. Cambridge University Press, 1999. doi:10.1017/CB09780511721335.010.
- [35] Nicholas C Wormald. Generating random regular graphs. *Journal of algorithms*, 5(2):247–280, 1984. doi:10.1016/0196-6774(84)90030-0.
- [36] Jaewon Yang and Jure Leskovec. Defining and evaluating network communities based on ground-truth. In *2012 IEEE 12th International Conference on Data Mining*, pages 745–754, 2012. doi:10.1109/ICDM.2012.138.
- [37] Jaewon Yang and Jure Leskovec. Overlapping communities explain core–periphery organization of networks. *Proceedings of the IEEE*, 102(12):1892–1902, 2014. doi:10.1109/JPROC.2014.2364018.

A Density of Intersection Figures




Figure 11: The density of overlaps compared to the empirical ρ on YouTube-like **ABCD+o²** graphs.




Figure 12: The density of overlaps compared to the empirical ρ on DBLP-like **ABCD+o²** graphs.




Figure 13: The density of overlaps compared to the empirical ρ on Amazon-like **ABCD+o²** graphs.

B Internal Edge Fraction Figures

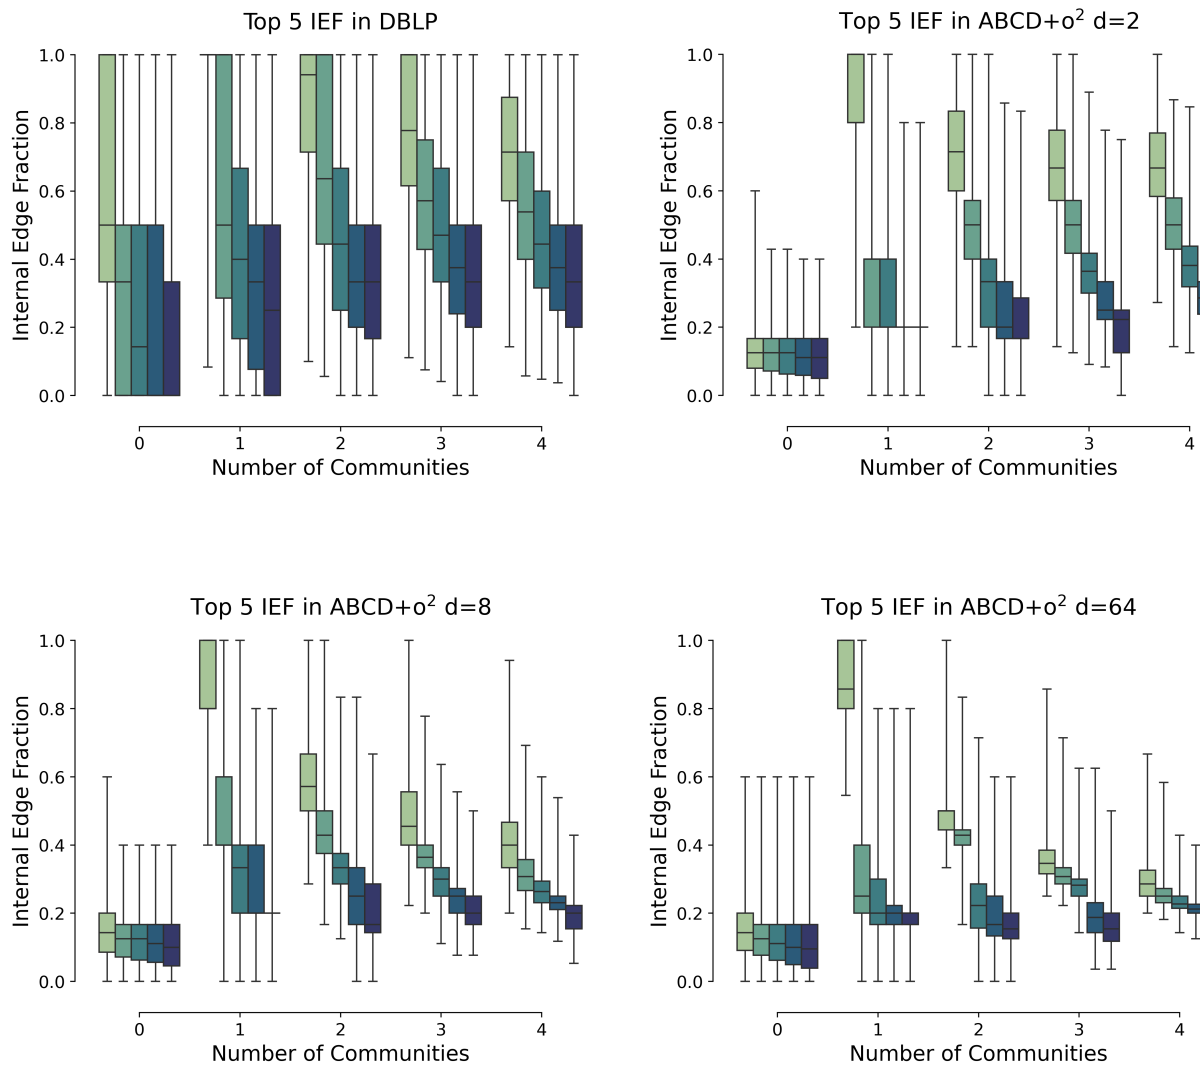


Figure 14: The top 5 IEF values in decreasing order, grouped by the number of communities the node belongs to, for DBLP and DBLP-like **ABCD+ σ^2** graphs.

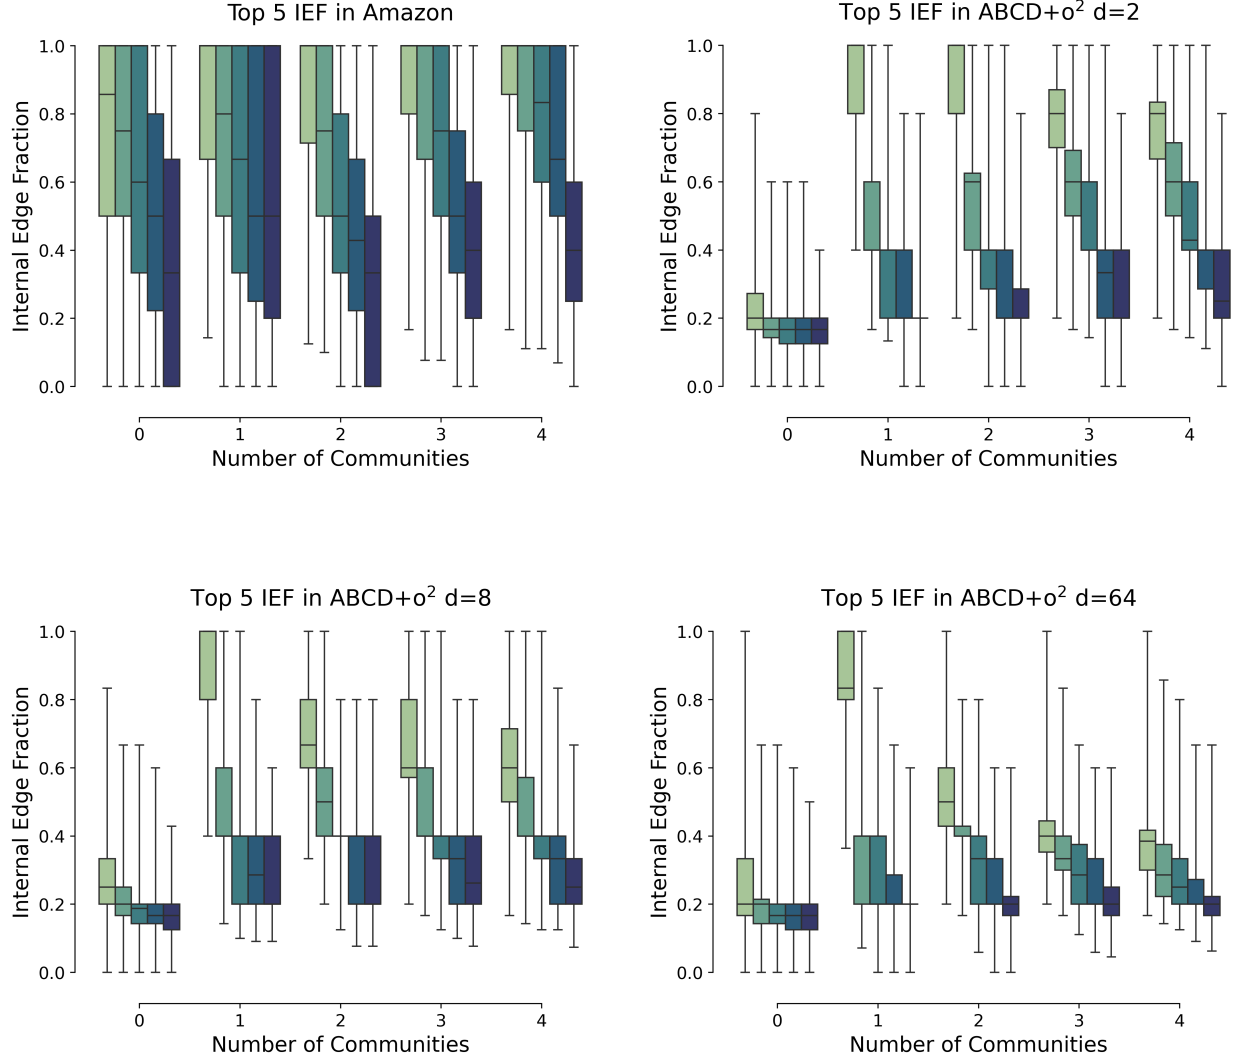


Figure 15: The top 5 IEF values in decreasing order, grouped by the number of communities the node belongs to, for Amazon and Amazon-like $\mathbf{ABCD+o^2}$ graphs.

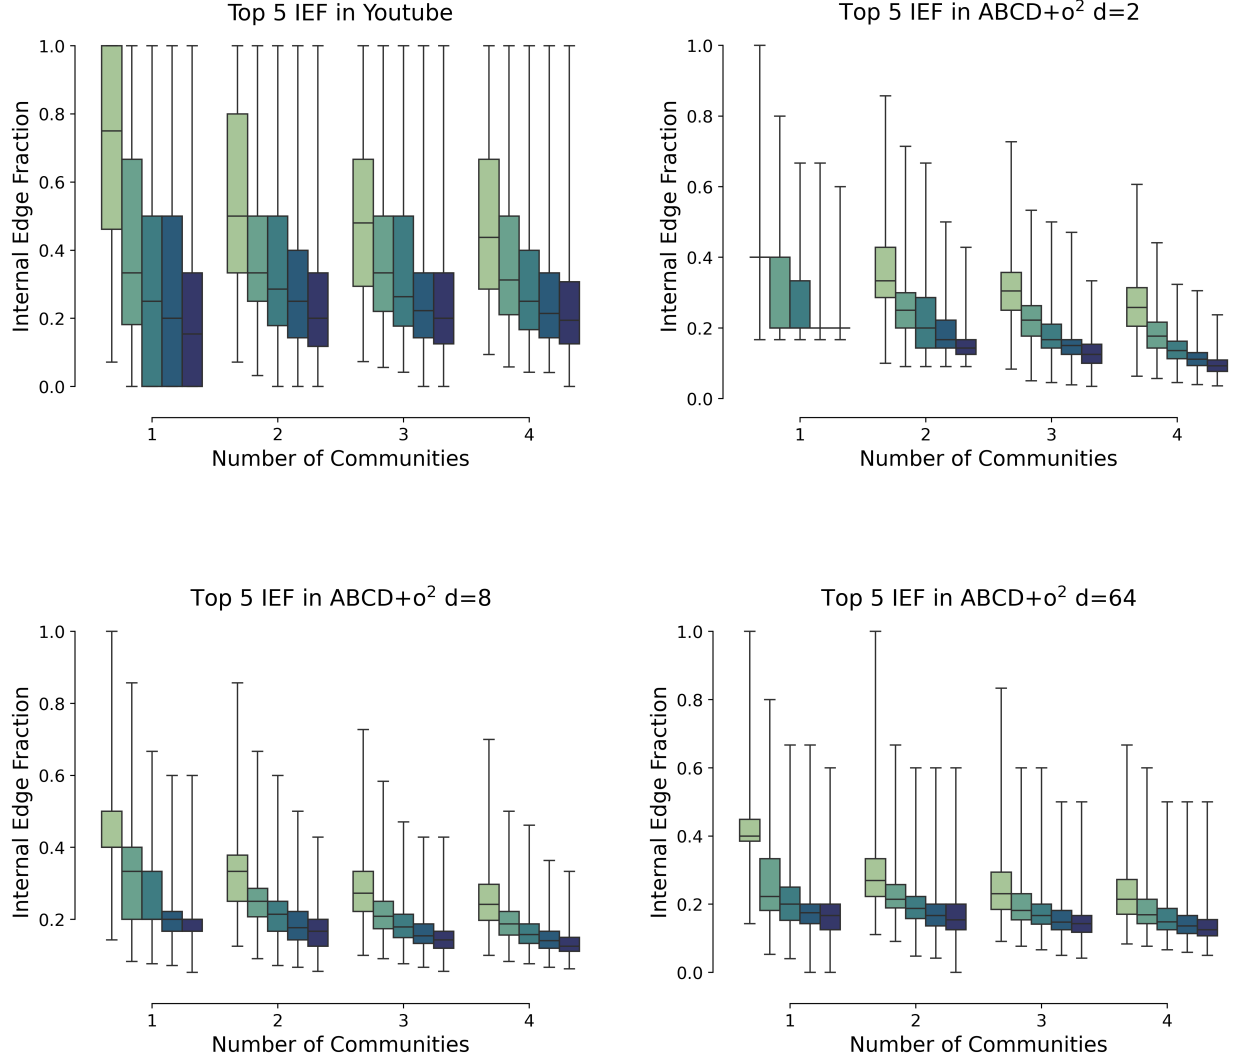


Figure 16: The top 5 IEF values in decreasing order, grouped by the number of communities the node belongs to, for YouTube and YouTube-like **ABCD+o²** graphs.