

# Testing Alignment of Node Attributes with Network Structure Through Label Propagation

Natalie Stanley

The University of North Carolina at  
Chapel Hill  
NatalieStanley1318@gmail.com

Marc Niethammer

The University of North Carolina at  
Chapel Hill  
mn@cs.unc.edu

Peter J. Mucha

The University of North Carolina at  
Chapel Hill  
much@unc.edu

## ABSTRACT

Attributed network data is becoming increasingly common across fields, as we are often equipped with information about nodes in addition to their pairwise connectivity patterns. This extra information can manifest as a classification, or as a multidimensional vector of features. Recently developed methods that seek to extend community detection approaches to attributed networks have explored how to most effectively combine connectivity and attribute information to identify quality communities. These methods often rely on some assumption of the dependency relationships between attributes and connectivity. In this work, we seek to develop a statistical test to assess whether node attributes align with network connectivity. The objective is to quantitatively evaluate whether nodes with similar connectivity patterns also have similar attributes. To address this problem, we use a node sampling and label propagation approach. We apply our method to several synthetic examples that explore how network structure and attribute characteristics affect the empirical  $p$ -value computed by our method. Finally, we apply the test to a network generated from a single cell mass cytometry (CyTOF) dataset and show that our test can identify markers associated with distinct sub populations of single cells.

## KEYWORDS

Community detection, attributed networks, label propagation

### ACM Reference format:

Natalie Stanley, Marc Niethammer, and Peter J. Mucha. 2018. Testing Alignment of Node Attributes with Network Structure Through Label Propagation. In *Proceedings of workshop on mining and learning with graphs, London, United Kingdom, August 2018 (MLG '18)*, 8 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 INTRODUCTION

Community detection in networks is a common pursuit that seeks to partition the network's nodes into sets of structurally coherent groups, where members of a *community* have strong similarity in connectivity patterns [7, 11, 16]. While the identification of communities based solely on the network's adjacency matrix is straightforward, the implications of having node attributes and how to integrate that into the community detection problem is not as well

understood. We refer to a *structural community* as a community identified according to only the adjacency matrix, while we define an *attribute community* as a community that has been identified using the attribute information. Recently, there have been numerous approaches extending common community detection techniques to attributed networks [5, 8, 13–15, 22]. While each of these methods provide extensions to a variety of community detection approaches, they also differ in their assumption about the dependence relationships between the attributes and connectivity. On one hand, it seems reasonable to assume that members of a structural community should be highly similar in attribute space. However, work by Newman *et al.* [13] and Peel *et al.* [14] have provided phenomenal examples and insight into when this assumption could be invalid.

In this work, we seek to develop a test that returns a statistic reflecting how closely node attributes correlate with connectivity patterns. Our test is based on label propagation and ultimately returns an empirical  $p$ -value that can be interpreted as the significance of the relationship between network connectivity and node attributes. We demonstrate that the computed empirical  $p$ -value is meaningful with several synthetic examples and on a network representation of a single-cell mass cytometry CyTOF dataset.

This paper is organized as follows: First, we describe the latest advances in attributed community detection. Next, we define our method and show several synthetic examples to evaluate the performance and meaningfulness of our computed empirical  $p$ -value. Finally, we apply our test to a single cell mass cytometry CyTOF dataset, identifying identify particular attributes that can distinguish populations of cells.

### 1.1 Community detection in attributed networks

There are a variety of approaches for identifying structural communities based only on connectivity patterns such as probabilistic methods [17, 21], quality function maximization [12], and higher-order motif-based clustering [2]. Most of these classes of methods have inspired extensions in attributed networks. Here, we discuss extensions to probabilistic and quality function maximization approaches. Note that throughout this text we use the phrases *attributes* and *metadata* interchangeably to be consistent with the terminology in the discussed references.

**1.1.1 Probabilistic approaches.** Probabilistic community detection methods edge existence in a network based on latent community structure. After learning model parameters through likelihood optimization, samples generated from the model should align closely with the true underlying network. Two common approaches

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).  
MLG '18, August 2018, London, United Kingdom  
© 2018 Copyright held by the owner/author(s).  
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM.  
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

are the stochastic block model (SBM) [17] and the affiliation model [21].

The assumption of the stochastic block model is that nodes within a community are connected to nodes within and between communities in a characteristic way. Moreover, the objective in the model fitting and parameter inference of a stochastic block model in a network with  $K$  communities is to learn the node-to-community assignments and the within and between community connection probabilities that maximize the model likelihood. The stochastic block model has been extensively studied in the literature and has been extended to attributed networks in at least four ways [8, 13, 14, 18]. Newman *et al.* provided an extension to the stochastic block model capable of incorporating discrete or continuous metadata (attribute) information [13]. Peel *et al.* proposed the neoSBM [14], which explores the effects of ‘fixing’ and freeing nodes from their metadata label on the SBM inference. Along with this work, the authors developed BESTest, a statistical test to measure the relationship between a node’s metadata label and community structure. Hric *et al.* constructed a joint stochastic block model for both the attributes and metadata through a nonparametric, Bayesian framework [8]. They assessed the alignment of the attributes with the connectivity based on performance in link prediction tasks. Stanley *et al.* introduced a version of the stochastic block model for networks with nodes having multiple continuous attributes [18].

Another useful probabilistic model for community structure is the affiliation model. This model assumes that nodes can be affiliated to multiple communities to varying extents [21]. Moreover, the edge between a pair of nodes is based on their similarity in community affiliations. A useful method for integrating multidimensional vectors of binary attributes was introduced by Yang *et al.* in a method called CESNA [22], which modifies the affiliation model likelihood to incorporate this information. This is achieved by allowing the attributes and connectivity information to be modeled as conditionally independent, giving the node-to-community affiliations and feature importance weights for the attributes.

**1.1.2 Quality function maximization.** Quality function maximization methods have also been extended to attributed networks. When community detection is formulated with a quality function, the objective is to specify a null model for a network with no community structure and find the partition of nodes to communities that maximizes the difference from this null model. A standard quality function for communities is known as modularity [12]. The state-of-the-art optimization heuristic for maximizing modularity is the Louvain algorithm [3]. Work by Combe *et al.* adapted the modularity to take into account multidimensional attribute vectors and optimized this quantity in a Louvain-style manner with I-Louvain [5].

Recent work by Perozzi *et al.* defines a modularity-inspired quantity known as *community normality* [15]. This measure prioritizes partitions where members of a community are very similar to each other in attribute space (and obviously in connectivity patterns). Further, members of a community are also expected to be different from nodes on the community boundary or in a different community.


## 1.2 Novelty and paper objectives

The methods described above indicate great progress in the integration of attributes in community detection. However, there has not been much work focusing on the interplay between attributes and connectivity information and the extent to which these data should be combined. In this paper, we seek to develop a statistical test to evaluate the relationship between connectivity and attribute information. Our approach is meant to be generalizable to all networks and agnostic to the particular community detection algorithm applied to the network. Our method can also accommodate multiple attributes that can be either discrete, continuous, or a combination of both. After defining the method for implementing this test, we seek to systematically study the properties of our test and its empirical  $p$ -value output across various types of networks and attributes. Finally, we apply our tool to a single cell mass cytometry (CyTOF) dataset, where our results suggest that our method can successfully identify attributes with discriminative ability for distinguishing between communities.

## 2 METHODS

Our method is built on label propagation (LP), where given a partially labeled network of  $N$  nodes (i.e. only a subset of nodes have community assignments), the objective is to use this information to predict the community assignments of the unlabeled nodes [20]. In this work, we first *label* the nodes according to their attribute information. We then take several sub samples of  $l$  labeled nodes and use the prediction of the remaining  $N - l$  unlabeled nodes as a proxy for how closely the attributes align with the network connectivity. In particular, we use a label propagation approach that returns a probability distribution for each node over each of the attribute-defined node-to-community assignments. We then quantify the uncertainty of the prediction with a simple entropy measure. In doing this, we assume that if the attributes are aligned with the network connectivity patterns, the entropy should be low. Alternatively, if attributes and connectivity are disparate, then predicting the unlabeled nodes will be difficult and entropy should be higher.

As an overview of this process, we first label the nodes according to their attribute information. This can be achieved by classifying the nodes according to a single discrete value, or through simple clustering of the nodes based on their attributes. After obtaining a labeling of the nodes, we begin our label propagation and null label propagation tasks. For a large number of  $T^*$  trials, we take a sample of  $l$  nodes and their attribute-based labels and denote these nodes as *labeled*. We then try to predict the labels of the remaining  $N - l$  nodes, comprising the *unlabeled* set, by propagating the labels outward. Since the label propagation method returns a probability distribution over possible community assignments, we can compute an entropy measure. Along with this true label propagation task, in each trial we also permute the labels of the nodes in our sample set to generate a null distribution of entropy values for the unlabeled nodes. Finally, the overlap between the null and empirical entropy distributions are used to compute a  $p$ -value. This process is outlined in Figure 1. We will now provide a detailed description of each step in this process.



**Figure 1: Our test first labels the nodes according to the attribute-based node-to-community assignment,  $\tilde{z}$ . Next, in each of  $T^*$  trials, a sample of  $l$  labeled nodes is used as input to the label propagation task to predict the probability distribution over possible communities for the unlabeled  $N - l$  nodes. The entropy of the node-to-community assignment probabilities is used to estimate how well the attributes align with connectivity. In each trial,  $\tilde{z}$  is permuted and subjected to the label propagation task to compute a ‘null’ entropy value. This process is repeated  $T^*$  times and the empirical  $p$ -value is calculated based on the overlap between the null entropy distribution and the empirical entropy distribution.**

## 2.1 Notation

We first define some notation that assists in setting up this problem. For a network with  $N$  nodes, we let  $\mathbf{z}$  be the  $N$ -length vector of node-to-community assignments, based on only the network connectivity information given in the adjacency matrix,  $\mathbf{A} = \{A_{ij}\}$ . This implies that the  $i$ -th entry,  $z_i$ , gives the community assignment for node  $i$ . Alternatively, when nodes are labeled according to the attribute information, we denote their community assignments with  $\tilde{\mathbf{z}}$ . Finally, our permutation test involves taking a subset of nodes and their labels in  $\tilde{\mathbf{z}}$  as the *labeled* set to propagate labels outward to a set of *unlabeled* nodes. We denote this distinction between the attribute-based partition of the labeled and unlabeled subsets of nodes by  $\tilde{\mathbf{z}}^L$  and  $\tilde{\mathbf{z}}^U$ , respectively. Finally, we assume that each node has  $p$  associated attributes, which are stored in the  $N \times p$  matrix,  $\mathbf{X}$ . That is, the  $i$ th row of  $\mathbf{X}$ ,  $X_i$ , gives the values of the  $p$  attributes for node  $i$ .

## 2.2 Classifying Nodes

The first step is to classify nodes according to attributes, denoted by  $\tilde{\mathbf{z}}$ . We assume some prior knowledge for the  $K$ , specifying how many communities are in the data. Hence, each  $\tilde{z}_i$  takes on 1 of  $K$  values. In the case where nodes are classified discretely, according to a single source of information, this labeling occurs without any effort. In the case where each node has multiple attributes, we have found that a simple clustering method, such as  $k$ -means works well. Because this first step of labeled the nodes is achieved through a clustering procedure, our approach can accommodate discrete and continuous attributes.

## 2.3 Sampling Nodes and Creating Entropy Distributions

In the sampling step, for a large number of  $T^*$  trials, we randomly select  $l$  nodes,  $\{L\}$ , and their corresponding labels,  $\tilde{\mathbf{z}}^L$ . From here,

we seek to use  $\tilde{\mathbf{z}}^L$  and the network structure to predict the labels for the the remaining  $N - l$  nodes in the unlabeled set,  $\{U\}$ .

After splitting all  $N$  nodes into their labeled and unlabeled sets, we use the label propagation approach described by Zhu *et al.* [23] to generate a probability distribution over the communities from  $\tilde{\mathbf{z}}^L$  for each of the nodes in  $\{U\}$ . Ultimately under this LP approach, we seek to define the  $N \times K$  matrix,  $\mathbf{Y}$ , where  $Y_{ic}$  is the probability that node  $i$  belongs to class  $c$ . We can split this matrix into two matrices,  $\mathbf{Y}^L$  and  $\mathbf{Y}^U$ , containing the subset of rows corresponding to nodes in  $\{L\}$  and  $\{U\}$ , respectively. Therefore, the label propagation task is to effectively estimate  $\mathbf{Y}^U$ .

To compute  $\mathbf{Y}^U$  following the approach by Zhu *et al.* [23], we first use the adjacency matrix,  $\mathbf{A}$ , to define an  $N \times N$  transition probability matrix,  $\mathbf{T} = \{T_{ij}\}$ . Here, entry  $T_{ij}$  gives the probability of jumping from node  $j$  to node  $i$ . This is given by,

$$T_{ij} = \frac{A_{ij}}{\sum_{f=1}^N A_{fj}}. \quad (1)$$

We then compute  $\tilde{\mathbf{T}}$ , which is simply the row normalized version of  $\mathbf{T}$ . From here,  $\tilde{\mathbf{T}}$  is rearranged so that the first  $L$  rows and columns correspond to the labeled nodes, and the remaining  $N - L$  columns correspond to the unlabeled nodes. This rearrangement allows for  $\tilde{\mathbf{T}}$  to be written as four submatrices, obtained by splitting  $\tilde{\mathbf{T}}$  after the  $l$ th row and  $l$ th column as,

$$\tilde{\mathbf{T}} = \begin{bmatrix} \tilde{\mathbf{T}}_{ll} & \tilde{\mathbf{T}}_{lu} \\ \tilde{\mathbf{T}}_{ul} & \tilde{\mathbf{T}}_{uu} \end{bmatrix}. \quad (2)$$

Based on the fact that Zhu *et al.*, define their label propagation algorithm as  $\mathbf{Y} \leftarrow \mathbf{T}\mathbf{Y}$ ,  $\mathbf{Y}^U$  can be defined as,

$$\mathbf{Y}^U = (\mathbf{I} - \tilde{\mathbf{T}}_{uu})^{-1} \tilde{\mathbf{T}}_{ul} \mathbf{Y}^L. \quad (3)$$

More details about this label propagation approach are described in Ref. [23]. Computing  $\mathbf{Y}^U$  for one pair of labeled and unlabeled node sets,  $\{L\}$  and  $\{U\}$ , comprises the true label propagation task of one trial. To perform a null LP task, we first permute the entries of  $\tilde{\mathbf{z}}^L$ , and denote this permuted version as  $\tilde{\mathbf{z}}^{L,\text{perm}}$ . Just as we showed in the true label propagation task, we use  $\tilde{\mathbf{z}}^{L,\text{perm}}$  to define a corresponding permuted version of  $\mathbf{Y}^U$ , with  $Y_{ic}^{U,\text{perm}}$  set to be 1 if node  $i$  belongs to community  $c$ , under the permuted labels, given by  $\tilde{\mathbf{z}}^{L,\text{perm}}$ . The analogous update relationship shown in equation 3 gives  $\mathbf{Y}^{U,\text{perm}}$ .

After computing  $\mathbf{Y}^U$  and  $\mathbf{Y}^{U,\text{perm}}$ , the next step is to compute their corresponding entropies,  $E$  and  $E^{\text{perm}}$ . We chose to use a cross entropy measure to account for both the attribute-based community assignment ( $\tilde{\mathbf{z}}^U$ ) and the prediction under our label propagation task,  $\mathbf{Y}^U$ . From our attribute-based classification vector  $\tilde{\mathbf{z}}$ , we create the  $(N - l) \times K$  indicator matrix,  $\tilde{\mathbf{Z}}^U$ , where  $\tilde{Z}_{ic}^U = 1$  if node  $i$  is assigned to community  $c$ . With this notation, we compute cross entropy  $H(\tilde{\mathbf{Z}}^U, \mathbf{Y}^U)$  as,

$$H(\tilde{\mathbf{Z}}^U, \mathbf{Y}^U) = - \sum_{ic} \tilde{Z}_{ic}^U \log(Y_{ic}^U). \quad (4)$$

Moreover,  $H(\tilde{\mathbf{Z}}^U, \mathbf{Y}^U)$  and  $H(\tilde{\mathbf{Z}}^{U,\text{perm}}, \mathbf{Y}^{U,\text{perm}})$  give  $E$  and  $E^{\text{perm}}$ , respectively.

We let  $\mathcal{E} = \{E_1, E_2, \dots, E_T\}$  and  $\mathcal{E}_{\text{perm}} = \{E_1^{\text{perm}}, E_2^{\text{perm}}, \dots, E_T^{\text{perm}}\}$  be the collection of entropies over the  $T^*$  trials.

## 2.4 Computing the empirical $p$ -value

After having repeated this LP task over  $T^*$  trials, we compute the empirical  $p$ -value for the test. This  $p$ -value is interpreted as the overlap between  $\mathcal{E}$  and  $\mathcal{E}_{\text{perm}}$ . In the case where attributes ( $X$ ) and connectivity ( $A$ ) are well-aligned with connectivity,  $\mathcal{E}$  and  $\mathcal{E}_{\text{perm}}$  should not overlap because the entropy for the label propagation task should be very low. Alternatively, as  $X$  and  $A$  become less aligned, the entropy of the prediction from the label propagation task should be higher and hence  $\mathcal{E}$  and  $\mathcal{E}_{\text{perm}}$  will overlap. Then the empirical  $p$ -value,  $p$ , is calculated as,

$$p = P(\mathcal{E}_{\text{perm}} < \max(\mathcal{E})). \quad (5)$$

Note that this  $p$ -value is strictly empirical and intended to quantify the overlap between  $\mathcal{E}$  and  $\mathcal{E}_{\text{perm}}$ .

## 3 RESULTS

We present results on synthetic networks and on a network representation of a single cell mass cytometry CyTOF dataset. In this section, we seek to confirm that the empirical  $p$ -value leads to an accurate and interpretable conclusion. The results on synthetic data are useful because we have an understanding of when the  $p$ -value should be significant, due to our knowledge of how the data were generated. Similarly, in the single cell mass cytometry CyTOF dataset, we use particular marker features and their discriminative ability to validate our computed empirical  $p$ -values.

### 3.1 Synthetic Examples


We sought to verify that our empirical  $p$ -value was capturing desirable behavior. First, we expected the  $p$ -value to decrease in significance as the LP entropy distribution increased in overlap with the empirical null LP entropy distribution. Second, we sought to have a  $p$ -value that became less significant (i.e. higher) as the correlation between attributes and network structure decreased. In Figure 2, we considered a network generated from a stochastic block model with  $N = 200$  nodes,  $K = 4$  communities, within-community edge probability ( $p_{in}$ ),  $p_{in} = 0.6$ , and between-community edge probability, ( $p_{out}$ ),  $p_{out} = 0.02$ . That is for a pair of nodes,  $i$  and  $j$ , the probability of an edge existing between them is modeled as  $P(A_{ij} = 1) = p_{in}$  if  $z_i = z_j$  and  $P(A_{ij} = 1) = p_{out}$  if  $z_i \neq z_j$ .

Associated with each node is a 3-dimensional Gaussian attribute vector, drawn from 1 of  $K$  multivariate Gaussian distributions. Under this formulation, each community has its own associated multivariate Gaussian distribution. The attribute vector for a node in community  $k$  is parameterized by mean  $\mu_k = [\mu_1, \mu_2, \mu_3]$  and covariance matrix  $\Sigma_k$ .

To generate each  $\mu_k = [\mu_1, \mu_2, \mu_3]$ , we draw each of the  $\mu_1, \mu_2$  and  $\mu_3$  from a standard Normal distribution with mean 0 and unit variance. For a community  $k$ ,  $\Sigma_k$  is also the identity covariance matrix.

When performing our label propagation task, in each of the  $T^* = 1000$  trials, we used a sample of  $l = 100$  labeled nodes as the input to the LP task.

First, we performed a tSNE [10] 2-dimensional embedding of the nodes based only on the attribute information,  $X$ . This is shown in Figure 2A. Each point represents a node and is colored by its community assignment,  $z$ . We can see that there are clearly clusters



**Figure 2: To understand the properties of our empirical  $p$ -value, we generated a synthetic network, A, from an SBM with  $N = 200$  nodes, and  $K = 4$  communities. The vector of continuous attributes for a node  $i$ , ( $X_i$ ) was drawn from a multivariate Gaussian distribution parameterized by its community assignment ( $z_i$ ) or  $\{\mu_{z_i}, \Sigma_{z_i}\}$ . In these experiments, we permuted varying proportions of  $\tilde{z}$  and observed the effects on entropy and empirical  $p$ -value. A. We used tSNE to visualize the two dimensional projection of the 200 nodes based on attribute information. B. We plotted the empirical  $p$ -value of our test as a function of the proportion of attribute-based community assignments,  $\tilde{z}$ , permuted. We varied the proportion of entries of  $\tilde{z}$  permuted between 0.01 and 1 (horizontal axis). We observed decreased statistical significance (increased empirical  $p$ -value) with an increasing proportion of permuted labels. C. We plotted the empirical  $p$ -value as a function of the mean entropy ( $\mathcal{E}$ ) across  $T^* = 1000$  trials used to generate the entropy distributions for each experiment described in (B.). Increased entropy corresponding to a larger proportion of  $\tilde{z}$  permuted leads to a decreased  $p$ -value.**

of nodes from the same community, but there is also some mixing. This implies that attributes and connectivity are not perfectly correlated.

In this first experiment, we sought to explore how the empirical  $p$ -value behaved as the LP entropy distribution converged to the null LP entropy distribution. To study this, we consciously made the label propagation task incrementally more difficult by perturbing various proportions of the initial attribute-based node-to-community assignments,  $\tilde{z}$ . This test was implemented to verify that with a higher proportion of permuted (i.e. incorrect) entries in  $\tilde{z}$ , the empirical entropy distributions,  $\mathcal{E}$  and  $\mathcal{E}_{\text{perm}}$  would have more extensive overlap. As expected, in Figure 2B. we observed that by permuting a larger proportion of the labels,  $\tilde{z}$ , there was an associated increase in the empirical  $p$ -value (decreased significance). Here, the proportion of labels permuted in  $\tilde{z}$  was varied between 0.01 and 1 (horizontal axis). Next, we examined the relationship between the entropy distribution,  $\mathcal{E}$  and our empirical  $p$ -value,  $p$  in the experiments described in Figure 2B. In Figure 2C., we plot the empirical  $p$ -value against the mean entropy of  $\mathcal{E}$  over the  $T^* = 1000$  trials. As expected, these quantities are highly related, with a higher entropy leading to a higher (less significant) empirical  $p$ -value.

**3.1.1 Comparison to BESTest.** We used the synthetic data from the experiment described in Figure 2 to compare our results to those obtained using BESTest [14]. Recall that BESTest is the method developed to measure the relationship between communities and a piece of node metadata in the context of a stochastic block


model. This comparison is especially appropriate since the synthetic data were generated from a stochastic block model. BESTest works first by labeling the nodes according to  $\tilde{z}$ , based on the attribute information. Under this partition of the nodes, the SBM parameters are optimized, where the maximum likelihood estimate for the connection probability between a pair of communities  $r$  and  $s$  is given by  $\hat{\omega}_{rs}$ . This maximum likelihood estimate  $\hat{\omega}_{rs}$  is computed as  $\hat{\omega}_{rs} = m_{rs}/n_r n_s$ . Here,  $m_{rs}$  is the number of edges between communities  $r$  and  $s$ , while  $n_r$  and  $n_s$  are the number of nodes in communities  $r$  and  $s$ , respectively. The entropy,  $\mathcal{H}$  of this partition across the communities is computed as,

$$\mathcal{H}(\tilde{z}) = -\frac{1}{2} \left[ \sum_{rs} m_{rs} \log \hat{\omega}_{rs} + (n_r n_s - m_{rs}) \log(1 - \hat{\omega}_{rs}) \right] + O(N^{-1}). \quad (6)$$

The empirical  $p$ -value is computed with BESTest through a permutation test which computes  $\mathcal{H}(\tilde{z}^{\text{perm}})$  many times and reports the fraction of  $\mathcal{H}(\tilde{z}^{\text{perm}}) < \mathcal{H}(\tilde{z})$ . The BESTest entropy measure was developed in the context of a stochastic block model. While we show in subsequent experiments that our results are highly correlated with the BESTest results, our approach is developed outside of the context of the stochastic block model. We analyzed the similarity between BESTest and our label propagation approach by studying the relationship between the BESTest entropy, and our entropy and empirical  $p$ -value.

As described in section 3.1, our experiment involved permuting varying proportions of  $\tilde{z}$ , which resulted in a range of entropies and hence computed empirical  $p$ -values. Each experiment corresponding to a particular proportion of permuted entries of  $\tilde{z}$  served as the attribute-based node classifications in our calculations of label propagation entropy and BESTest entropy. Note that to compute the BESTest entropies, the whole network is used. In Figure 3A, we plot our computed  $p$ -value against the BESTest entropy and observe that these quantities are highly related. Even more related are the BESTest and label propagation entropies, plotted against each other in Figure 3B. These analyses suggest that these tests are highly related for this particular experiment. Since we did use stochastic block models to generate our synthetic data, interesting future work could examine the relationship between these tests for more diverse types of network structures.

**3.1.2 Strength of community structure.** Given that aspects of network structure can influence label propagation results, we sought to experimentally probe how our test performed for community structure of varying strengths. We refer to a strong community structure as one that has many within-community connections and few between-community connections. In this case, communities are easy to identify, based on the prominent organizational structure. To approximate this, we considered the  $p_{in}$  to  $p_{out}$  ratio for a stochastic block model. As previously described,  $p_{in}$  is defined as the probability of observing an edge between a pair of nodes in the same community, while  $p_{out}$  is the probability of observing an edge between a pair of nodes in different communities. We expected that the entropy and empirical  $p$ -value would decrease with an increasing  $p_{in}/p_{out}$  ratio. That is, as the community structure becomes less prominent with an increased number of connections between communities, the label propagation task should become



**Figure 3: To understand the relationship between our test and BESTest, we performed our analysis on the synthetic experiments described in section 3.1. In this experiment, we permute varying proportions of node-based attribute classifications,  $\tilde{z}$ , to produce a spectrum of entropy values. From this range of entropy values that we observed with our test, we wanted to better understand the with similarity with the BESTest results. A. There is a strong relationship between the BESTest entropy and the LP empirical  $p$ -value returned by our test. B. We plotted the BESTest entropy against our label propagation entropy and observed a very strong positive correlation. This analysis demonstrates that the entropies computed by our LP test and BESTest are highly similar.**


more difficult. To study this with synthetic data, we varied the  $p_{in}/p_{out}$  ratio, by considering a four community stochastic block model with values of  $p_{in}$  between 0.05 and 0.45 and choosing a corresponding  $p_{out}$ , such that the mean degree was equal 30. For each pair of  $p_{in}$  and  $p_{out}$ , we generated 10 realizations of the corresponding stochastic block model. Accompanying each synthetic network was a fixed 3-dimensional attribute matrix,  $X$ , where the attribute vectors for the members of community  $k$  were drawn from a 3-dimensional multivariate Gaussian, parameterized by  $\{\mu_k, \Sigma_k\}$ . For each synthetic network, we computed the entropy under our label propagation method and the corresponding  $p$ -value.

In Figure 4A, we plot the mean LP entropy over the  $T^* = 100$  samples used to construct the empirical entropy distribution,  $\mathcal{E}$ , across the 10 network realizations for each set of  $p_{in}$  and  $p_{out}$ . The shaded region denotes the standard deviation of the LP entropy. As the ratio between  $p_{in}$  and  $p_{out}$  increases, the empirical LP entropy decreases. We see a similar effect in Figure 4B, where we plot the empirical  $p$ -value as a function of the  $p_{in}/p_{out}$  ratio. In this plot, the shaded region denotes the standard deviation of the empirical  $p$ -value. Here, a significant  $p$ -value (at  $\alpha = 0.05$ ) was sometimes reached (implying attributes and connectivity are aligned) when  $p_{in}/p_{out} > 3$ . Finally in Figure 4C, we examined the relationship between the mean empirical entropy ( $\mathcal{E}$ ) and the associated mean empirical  $p$ -value across the 10 network realizations generated under each parameter pair. We observe that when the LP entropy (horizontal axis) reaches 130, the mean empirical  $p$ -value increases (i.e. decreases in significance) very quickly.

## 4 MASS CYTOMETRY NETWORK EXAMPLE


We applied our test to a single cell mass cytometry CyTOF dataset. Mass cytometry with CyTOF [1] is an immunological profiling technique that gives simultaneous quantification of various immune





**Figure 4:** To understand the effect of the strength of a network’s community structure on our test, we generated synthetic networks from stochastic block models with various  $p_{in}$  (within-community) and  $p_{out}$  (between-community) parameters. Networks were generated with  $p_{in}$  varying between 0.05 and 0.45 and we chose a corresponding  $p_{out}$  such that the mean degree was 30. We used  $p_{in}/p_{out}$  as a proxy for the strength of community, with a higher value of this ratio indicating a stronger community structure with more within-community edges and fewer between community edges. For each  $p_{in}, p_{out}$  combination, we generated 10 synthetic network realizations. A. We plotted the relationship between our LP entropy and  $p_{in}/p_{out}$ . The shaded area denotes standard deviation of the mean entropy over the 10 networks for each  $p_{in}, p_{out}$  combination. B. Similar to (A.), we plotted the mean empirical  $p$ -value over the  $T^* = 100$  trials used to generate the entropy distributions,  $\mathcal{E}$  and  $\mathcal{E}_{perm}$ . For large  $p_{in}/p_{out}$ , the empirical  $p$ -value became more significant. The shaded area denotes standard deviation of empirical  $p$ -value over the 10 networks for each  $p_{in}, p_{out}$  combination. C. We plotted the relationship between the mean entropy ( $\mathcal{E}$ ) over the  $T^*=100$  trials and 10 network realizations for each SBM parameter pair and the empirical  $p$ -value.

features. The output of this technology is approximately 50 immune features profiled for each cell in a large collection of single cells. We used a freely available mass cytometry dataset, originally described in Ref. [19], but pre-processed in an R tool called Cytokit [4]. The dataset profiles 51 immune features across single cells on human T helper cells from peripheral blood and tonsils, which have shown to be heterogeneous within a sample. Note that consistent with the immunology literature, we also refer to these immune features as markers. To untangle the heterogeneity and infer cellular phenotypes, dimension reduction and clustering are typically applied to single cell data. In this pursuit, the objective is to cluster the single cells into predicted phenotypes, based on the measured features. A powerful way to segment the single cells into their respective phenotypes is by constructing a similarity network between the cells and clustering with community detection. This method for studying single cell data is called PhenoGraph and is described in Ref. [9]. We studied the data in an analogous way by constructing a 5-nearest neighbor network between the cells and applying community detection to cluster them. In this representation, each node in the  $k$  nearest neighbor network is a single cell and is connected to its 5 nearest neighbors, based on the pairwise Euclidean distance for the 51 measured immune features. In this example we considered a subset of 1000 single cells. After constructing the network, we predicted phenotypes by identifying communities ( $z$ )



**Figure 5:** We considered each of the 51 features in the CyTOF data and their potential to be used as discriminative markers for particular communities. We identified 11 communities (or inferred phenotypes) under the Louvain algorithm. We then created a partition,  $\tilde{z}$ , based on each attribute in isolation. For each attribute and its induced partition of the nodes,  $\tilde{z}$ , normalized mutual information (NMI) was used to measure the discriminative power of the marker in distinguishing network communities, or  $\text{NMI}(\tilde{z}, z)$ . We expected that our  $p$ -value should align with this NMI measure in that markers leading to high NMI between the induced  $\tilde{z}$  and  $z$  should have more significant  $p$ -values. A. We used a histogram to visualize the distribution of NMI values across the 51 possible markers, with many of them leading to low NMI (between 0 and 0.1). B. Similar to (A.), we visualized the empirical  $p$ -value for the 51 possible markers. C. We compared the relationship between the empirical  $p$ -value (vertical axis) and  $\text{NMI}(\tilde{z}, z)$  (horizontal axis) across the 51 possible markers. As expected, we observed these quantities to be anti-correlated in that more significant (lower) empirical  $p$ -values were obtained for higher values of  $\text{NMI}(\tilde{z}, z)$ .

with the Louvain algorithm [3]. Applying the Louvain algorithm to the network’s adjacency matrix,  $A$ , identified 11 communities. As shown in Ref. [4], one further analysis after clustering the single cells is to identify features with discriminative power between inferred phenotypes. We find the application of our LP task to CyTOF data to be an appealing validation task for our algorithm because there should be a set of features (i.e. the profiled markers) that have strong discriminative ability in separating communities in the network. Note that in this context, communities in the network are also closely linked to inferred cell phenotype.

The first test we performed on CyTOF  $k$  nearest neighbor network was to examine how each marker feature related to the community partition,  $z$ , identified with the Louvain algorithm. To understand the interplay between the  $k$  nearest neighbor network structure and the measured immune features, we performed 51 separate analyses. Each analysis considered the correlation of each immune feature individually with community structure. To produce a partition of the network,  $\tilde{z}$ , corresponding to a single marker,  $M$ , we simply clustered the 1000 nodes into 1 of 11 clusters, based on the value of marker  $M$ .

Before applying our LP test to this network, we used normalized mutual information (NMI) [6] to quantify the similarity between  $z$  and  $\tilde{z}$ . A high NMI (i.e. close to 1), indicates that the single attribute,  $M$ , used to create  $\tilde{z}$  creates a similar partition to the partition  $z$ , obtained from the Louvain algorithm on just the network structure (i.e. connectivity information only). Conversely, an NMI near 0




Figure 6: Under our label propagation task, we chose two markers with significant  $p$ -values (CD8 and TNFa) and two markers with insignificant  $p$ -values (IL4 and CD14). We colored the nodes in the  $k$ -nearest neighbor network of single cells by the expression of each of these four markers. The significant markers (shown in the top row) have expression that varies across communities. Conversely, the expression of the insignificant markers does not vary across communities. This analysis suggests that our test is effective in the identification of attributes that are well aligned with community structure.

indicates that when nodes (cells) are clustered based on the particular feature, their partition is very different to that obtained using connectivity information (A).

In Figure 5A, we show the distribution of NMIs computed between  $z$  and  $\tilde{z}$  for each of the 51 markers. We observe a fairly broad range of marker qualities represented. Similarly, we applied our LP task for  $T^* = 30$  trials and a sample of 500 labeled nodes in each trial. Figure 5B, shows the distribution of empirical  $p$ -values from our LP method. In this experiment we did not correct for multiple testing since significantly scoring features had very low empirical  $p$ -values (i.e. 0 or close to 0, implying no overlap between  $\mathcal{E}$  and  $\mathcal{E}_{perm}$ ). We noticed that there are approximately 30 markers with a low  $p$ -value (between 0 and 0.2), according to our LP test. Finally, in Figure 5C, we examined the relationship between the NMI between  $z$  and  $\tilde{z}$  and the empirical  $p$ -value, across each of the 51 markers.




Figure 7: We computed the empirical  $p$ -values induced by the partition  $\tilde{z}$  for each of the 51 markers in isolation. Using our test, we looked closely at 5 of the most and least significant markers under the computed empirical  $p$ -value. Since a quality marker in this case is said to be one that induces a labeling of the nodes,  $\tilde{z}$ , similar to the result obtained under  $z$ , we expect the expression of such a marker to vary across communities. In this plot, we show the expression of each marker as a function of the community index. The family of orange-colored lines correspond to the top 5 significant markers (according to empirical  $p$ -value). From all of these lines, the expression varies across communities. Conversely, looking at the markers with low significance under our test, expression is relatively constant across all communities.

As expected, these quantities are highly related, with high values of NMI corresponding to lower, more significant  $p$ -values.

To visualize how particular markers correlated with communities in the network, through their induced partition,  $\tilde{z}$ , we closely examined four different markers. Again, we used each attribute in isolation to perform an attribute-based partition of the nodes,  $\tilde{z}$  into 1 of 11 clusters. Two of these markers (CD8 and TNFa), had a  $p$ -value of 0 with our label propagation task. The other two markers (IL4 and CD14) had non-significant  $p$ -values of 0.47 and 0.97, respectively. In Figure 6, we show the  $k$  nearest neighbor networks of single cells with nodes colored by the expression for each of these four markers. Note that lighter colors indicate high expression and darker colors indicate lower expression. In the top row, we show the networks corresponding to CD8 and TNFa. We observe that expression patterns are indeed localized in the network. That is, dense clusters in the network tend to have similar marker expression. Further, this marker expression varies across communities in the network. Alternatively in the bottom row where we visualization IL4 and CD14, we observe that there is not much variability in marker expression across the network. In other words, using these markers to classify nodes would not be effective. This analysis further suggests that our empirical  $p$ -value can successfully identify markers with discriminative power because they are well-correlated with particular communities.

As an final experiment, we sought to see if the markers with significant empirical  $p$ -values (implying that they are effective in distinguishing cellular phenotypes) did indeed vary across communities in the network, through their induced partition,  $\mathcal{Z}$ . To do this, we selected 10 markers from the 51 measured features of the CyTOF data. In particular, we looked at the 5 of the most and least significant markers, in terms of the computed empirical  $p$ -value. For each of these 10 markers, we computed the mean marker expression across each of the 11 communities identified by applying the Louvain algorithm applied to the network's adjacency matrix,  $A$ , in Figure 7, we then plotted the mean marker expression across communities for the 5 most and least significant markers. The least significant markers are shown in the family of blue lines and are relatively static across each of the 10 communities. In contrast, the orange family of lines corresponds to the markers for the more significant features and do vary across communities. Since a marker with a significant low empirical  $p$ -value should correlate well with communities, this is the pattern we expected. The 5 poorly ranked markers clearly do not correlate with communities because their expression is constant across all communities.

## 5 CONCLUSION

In this paper, we introduced a label propagation based approach to determine how closely attributes align with network connectivity. Over  $T^*$  trials, the label propagation task uses a subset of nodes labeled according to attribute information to predict the labels for a set of unlabeled nodes. The label propagation task we adopt returns a probability distribution for each of the unlabeled nodes over the possible communities. The empirical  $p$ -value of our test is computed by comparing the empirical entropy distributions across the  $T^*$  trials from our label propagation task, and a permuted label propagation task, denoted by  $\mathcal{E}$  and  $\mathcal{E}_{\text{perm}}$ , respectively. The intuition is that if attributes are well aligned with network connectivity patterns, then the label propagation task should produce results that are more certain, and hence have lower entropy. Our results indicate that the computed entropy and empirical  $p$ -value are behaving as expected on synthetic examples, where we designed the experiments in a way that we knew how well the attributes and connectivity correlated. We also show that our test is useful in the identification of important marker features for distinguishing communities in the single cell mass cytometry  $k$  nearest neighbor network. Here, features (markers) with low empirical  $p$ -value are features that vary across communities and hence give insight into immune features that distinguish communities or cell subsets.

As future work, one might examine how the entropy and empirical  $p$ -value relate to different communities identified using different algorithms or approaches. Future work could also benefit from the analysis of the number of trials ( $T^*$ ) and the optimal number of nodes in each of the labeled and unlabeled node sets that lead to the most statistically stable and meaningful results. Finally, similar to how we detected particular marker features that were aligned with the identified communities, perhaps we can use our tool as a feature selection method that can be used to create meaningful network representations of data.

## ACKNOWLEDGMENTS

This work was supported by the National Science Foundation under award #1610762.

## REFERENCES

- [1] Sean C Bendall, Garry P Nolan, Mario Roederer, and Pratip K Chattopadhyay. 2012. A deep profiler's guide to cytometry. *Trends in immunology* 33, 7 (2012), 323–332.
- [2] Austin R Benson, David F Gleich, and Jure Leskovec. 2016. Higher-order organization of complex networks. *Science* 353, 6295 (2016), 163–166.
- [3] Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. 2008. Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment* 2008, 10 (2008), P10008.
- [4] Hao Chen, Mai Chan Lau, Michael Thomas Wong, Evan W Newell, Michael Poidinger, and Jimmiao Chen. 2016. Cytokit: a bioconductor package for an integrated mass cytometry data analysis pipeline. *PLoS computational biology* 12, 9 (2016), e1005112.
- [5] David Combe, Christine Largeron, Mathias Géry, and Előd Egyed-Zsigmond. 2015. I-Louvain: An Attributed Graph Clustering Method. In *Advances in Intelligent Data Analysis XIV*. Springer, 181–192.
- [6] Leon Danon, Albert Diaz-Guilera, Jordi Duch, and Alex Arenas. 2005. Comparing community structure identification. *Journal of Statistical Mechanics: Theory and Experiment* 2005, 09 (2005), P09008.
- [7] Santo Fortunato and Darko Hric. 2016. Community detection in networks: A user guide. *Physics Reports* 659 (2016), 1–44.
- [8] Darko Hric, Tiago P Peixoto, and Santo Fortunato. 2016. Network structure, metadata, and the prediction of missing nodes and annotations. *Physical Review X* 6, 3 (2016), 031038.
- [9] Jacob H Levine, Erin F Simonds, Sean C Bendall, Kara L Davis, D Amir El-ad, Michelle D Tadmor, Oren Litvin, Harris G Fienberg, Astraea Jager, Eli R Zunder, et al. 2015. Data-driven phenotypic dissection of AML reveals progenitor-like cells that correlate with prognosis. *Cell* 162, 1 (2015), 184–197.
- [10] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, Nov (2008), 2579–2605.
- [11] Peter J Mucha, Thomas Richardson, Kevin Macon, Mason A Porter, and Jukka-Pekka Onnela. 2010. Community structure in time-dependent, multiscale, and multiplex networks. *science* 328, 5980 (2010), 876–878.
- [12] Mark EJ Newman. 2006. Modularity and community structure in networks. *Proceedings of the National Academy of Sciences* 103, 23 (2006), 8577–8582.
- [13] Mark EJ Newman and Aaron Clauset. 2016. Structure and inference in annotated networks. *Nature Communications* 7 (2016), 11863.
- [14] Leto Peel, Daniel B Larremore, and Aaron Clauset. 2017. The ground truth about metadata and community detection in networks. *Science Advances* 3, 5 (2017), e1602548.
- [15] Bryan Perozzi and Leman Akoglu. 2018. Discovering Communities and Anomalies in Attributed Graphs: Interactive Visual Exploration and Summarization. *ACM Trans. Knowl. Discov. Data* 12, 2, Article 24 (Jan. 2018), 40 pages. <https://doi.org/10.1145/3139241>
- [16] Saray Shai, Natalie Stanley, Clara Granell, Dane Taylor, and Peter J Mucha. 2017. Case studies in network community detection. *arXiv preprint arXiv:1705.02305* (2017).
- [17] Tom AB Snijders and Krzysztof Nowicki. 1997. Estimation and prediction for stochastic blockmodels for graphs with latent block structure. *Journal of classification* 14, 1 (1997), 75–100.
- [18] Natalie Stanley, Thomas Bonacci, Roland Kwitt, Marc Niethammer, and Peter J Mucha. 2018. Stochastic Block Models with Multiple Continuous Attributes. *arXiv preprint arXiv:1803.02726* (2018).
- [19] Michael T Wong, Jimmiao Chen, Sriram Narayanan, Wenyu Lin, Rosslyn Anicete, Henry Tan Kun Kiaang, Maria Alicia Curotto De Lafaille, Michael Poidinger, and Evan W Newell. 2015. Mapping the diversity of follicular helper T cells in human blood and tonsils using high-dimensional mass cytometry analysis. *Cell reports* 11, 11 (2015), 1822–1833.
- [20] Jierui Xie and Boleslaw K Szymanski. 2011. Community detection using a neighborhood strength driven label propagation algorithm. In *Network Science Workshop (NSW), 2011 IEEE*. IEEE, 188–195.
- [21] Jaewon Yang and Jure Leskovec. 2013. Overlapping community detection at scale: a nonnegative matrix factorization approach. In *Proceedings of the sixth ACM international conference on Web search and data mining*. ACM, 587–596.
- [22] Jaewon Yang, Julian McAuley, and Jure Leskovec. 2013. Community detection in networks with node attributes. In *Data mining (ICDM), 2013 IEEE 13th international conference on*. IEEE, 1151–1156.
- [23] Xiaojin Zhu and Zoubin Ghahramani. 2002. Learning from labeled and unlabeled data with label propagation. (2002).