

# Active learning for node classification in assortative and disassortative networks

Cristopher Moore\*  
Computer Science Dept.  
University of New Mexico  
Albuquerque NM 87131 USA  
and Santa Fe Institute  
moore@cs.unm.edu

Xiaoran Yan  
Computer Science Dept.  
University of New Mexico  
Albuquerque NM 87131 USA  
everyxt@gmail.com

Yaojia Zhu  
Computer Science Dept.  
University of New Mexico  
Albuquerque NM 87131 USA  
yaojia.zhu@gmail.com

Jean-Baptiste Rouquier  
Complex Systems Institute  
Rhône-Alpes  
ENS Lyon, France  
jrouquie@gmail.com

Terran Lane  
Computer Science Dept.  
University of New Mexico  
Albuquerque NM 87131 USA  
terran@cs.unm.edu

## ABSTRACT

In many real-world networks, nodes have class labels or variables that affect the network’s topology. If the topology of the network is known but the labels of the nodes are hidden, we would like to select a small subset of nodes such that, if we knew their labels, we could accurately predict the labels of all the other nodes. We develop an active learning algorithm for this problem which uses information-theoretic techniques to choose which nodes to explore. We test our algorithm on networks from three different domains: a social network, a network of English words that appear adjacently in a novel, and a marine food web. Our algorithm makes no initial assumptions about how the groups connect, and performs well even when faced with quite general types of network structure. In particular, we do not assume that nodes of the same class are more likely to be connected to each other—only that they connect to the rest of the network in similar ways.

## Keywords

complex networks, structure and function, community detection, information theory, active learning, collective classification, transductive graph labeling

## 1. INTRODUCTION

In many networks, whether social, biological, or technological, nodes have underlying variables that are correlated with the network’s topology. In online social networks, blogs tend to link to other blogs with similar political views [1]. In vertebrate food webs, predators tend to eat prey whose mass is

\*This work was supported by the McDonnell Foundation.

smaller, but not too much smaller, than their own [11]. Networks of word adjacencies are correlated with those words’ parts of speech [30]. In the Internet, different types of service providers form different kinds of links based on their capacities and business relationships [3, 13]—and so on.

There has been a great deal of work on efficient algorithms for community detection in networks (see [12, 32] for reviews). However, most of this work defines a “community” as a group of nodes with high density of connections within the group and a low density of connections to the rest of the network. While this type of *assortative* community structure is common in social networks, we are interested in a more general definition of *functional* community—a group of nodes that connect to the rest of the network in similar ways. A set of predators might form a functional group in a food web, not because they eat each other, but because they eat similar prey. In English, nouns often follow adjectives, but seldom follow other nouns. Even some social networks have *disassortative* structure where pairs of nodes are more likely to be connected if they are from different classes. For example, some human societies are divided into moieties, and only allow marriages between different moieties [21].

We consider a setting where the topology of the network is known, but the class labels of the nodes are not. This could be the case, for instance, if we have a network of blogs and hyperlinks between them (like citations, trackbacks, blogrolls, etc.) and we are trying to classify the blogs according to their political leaning. Another possible application is in online social networks, where friendships are known and we are trying to infer hidden demographic variables. This problem is sometimes referred to as collective classification [35]. However, in that work the focus is on classification of individual nodes. In contrast, our focus is on the discovery of functional communities in the network, and our underlying generative model is designed around the assumption of that these communities exist.

We make no initial assumptions about the structure of the network—for instance, whether its groups are assortative, disassortative, or some mixture of the two. We assume that

we can learn the label of any given node, but at a cost, say in terms of work in the field or laboratory. Our goal is to identify a small subset of nodes such that, once we explore them and learn their labels, we can accurately predict the labels of all the others.

We present a general approach to this problem. Our algorithm uses information-theoretic measures to decide which node to explore next—which one will give us the most information about the rest of the network. We start with a probabilistic generative model of the network, called a *stochastic block model* [20, 38], in which groups connect to each other according to a matrix of probabilities. This model allows an arbitrary mixture of assortative and disassortative structure, as well as directed links from one group to another, and has been used to model networks in many fields (e.g. [4, 19, 33]).

We stress, however, that our approach could be applied equally well to many other probabilistic models—for instance, those where nodes belong to a mixture of classes [2], a hierarchy of classes and subclasses [10], locations in a latent geographical or social space [17], or niches in a food web [39]. It could also be applied to degree-corrected block models such as those in [23, 26, 31], which treat the nodes’ degrees as parameters rather than data to be predicted.

At each stage of the algorithm, some of the nodes’ labels are already known and we need to decide which node to explore next. We do this by estimating, for each node, the mutual information between its label and the joint distribution of all the others’ labels, conditioned on the labels of the nodes that are known so far. We obtain this estimate by Gibbs sampling, giving each classification of nodes a probability integrated over the parameters of the block model. We then explore the node for which this mutual information is largest.

A key fact about the mutual information, which we argue is essential to our algorithm’s performance, is that it is not just a measure of uncertainty: it is a combination of our uncertainty about a node’s label and the extent to which it is correlated with the labels of other nodes. Thus the algorithm explores nodes which maximize the expected amount of information we will gain about the entire network. It skips nodes whose labels seem obvious to it, or which are uncertain but have little effect on other nodes. In an assortative network, for instance, it starts by exploring nodes which are central to their communities, and then explores nodes along the boundaries between them, without being told in advance to pursue this strategy.

We also present an alternate approach which maximizes a quantity we call the *average agreement*. For each node  $v$ , this is the average number of nodes at which two independent samples of the Gibbs distribution agree, conditioned on the event that they agree at  $v$ . Like mutual information, average agreement is high for nodes that are highly correlated with the rest of the network. A similar idea (but not applied to networks) is present in [34].

We test our algorithm on three real-world networks: the social network of a karate club, a network of common adjacent words in a Charles Dickens novel, and a marine food web of

species in the Antarctic. Each of these networks is curated in the sense that we possess the correct node labels, such as the faction of the social network each individual belongs to, the part of speech of each word, or the part of the habitat each species lives in. We judge our algorithm according to how accurately it predicts the labels of the unexplored nodes, as a function of the number of nodes it has explored so far. We also compare our algorithm with several simple heuristics, such as exploring nodes based on their degree or betweenness centrality, and find that it significantly outperforms them.

## 2. RELATED WORK

The idea of designing experiments by maximizing the mutual information between the variable we learn next and the joint distribution of the other variables, or equivalently the expected amount of information we gain about the joint distribution, has a long history in statistics, artificial intelligence, and machine learning, e.g. Mackay [25] and Guo and Greiner [16]. Indeed, it goes back to the work of Lindley [24] in the 1950s. However, to our knowledge this is the first time it has been coupled with a generative model to discover hidden variables in networks.

In recent work, Zhu, Lafferty, and Ghahramani [41] study active learning of node labels using Gaussian fields and harmonic functions defined using the graph Laplacian. However, this technique only applies to networks where neighboring nodes are likely to be in the same class—that is, networks with assortative community structure. In contrast, our techniques are capable of learning about much more general types of network structure, including disassortative and directed relationships between functional communities.

Another approach to active learning of node labels is found in the work of Bilgic and Getoor [6] and Bilgic, Mihalkova, and Getoor [7], who use collective vector-based classifiers. By properly defining the collective relationships between nodes, both assortative or disassortative communities can be learned in this framework. However, our technique differs from theirs by using mutual information as the active learning criterion, which takes into account not only uncertainty, but correlation as well.

Additional works by Goldberg, Zhu, and Wright [14] and Tong and Jin [36] also perform semi-supervised learning on graphs, and handle the disassortative case. But they work in a setting where they know, for each link, if the ends should have the same or different labels, such as if one writer quotes another using negative words. In contrast, we work in a setting where we have no such information: only the topology is available to us, and the edges don’t have any labels telling us whether we should propagate similar or dissimilar labels.

## 3. MODEL AND METHODS

We represent our network as a directed graph  $G = (V, E)$  with  $n$  nodes. We assume that there are  $k$  classes of nodes, so that each node  $v$  has a class label  $t(v) \in \{1, \dots, k\}$ . We are given the graph  $G$ , and our goal is to learn the labels  $t(v)$ . To do this, we will assume that  $G$  is generated by some probabilistic model, in which its topology is correlated with these classes.

The simplest such model, although by no means the only one to which our methods could be applied, is a stochastic block model [20, 38]. It assumes that each pair of nodes  $u, v$  have an edge between them with a probability  $p_{t(u),t(v)}$  that depends only on their labels, and that these events are independent. Given a classification of the nodes  $t : V \rightarrow \{1, \dots, k\}$ , the probability of generating a given graph  $G$  in this model is

$$\begin{aligned} P(G|t, p) &= \left( \prod_{(u,v) \in E} p_{t(u),t(v)} \right) \left( \prod_{(u,v) \notin E} (1 - p_{t(u),t(v)}) \right) \\ &= \prod_{i,j=1}^k p_{ij}^{e_{ij}} (1 - p_{ij})^{n_i n_j - e_{ij}}. \end{aligned} \quad (1)$$

Here  $n_i = |\{v \in V : t(v) = i\}|$  is the number of nodes of class  $i$ , and  $e_{ij} = |\{(u, v) \in E : t(u) = i, t(v) = j\}|$  is the number of edges from nodes of class  $i$  to nodes of class  $j$ . If we wish to focus on undirected graphs, we can modify this expression by restricting the product over pairs of classes with  $i \leq j$ . We can also forbid self-loops, if we wish, by replacing  $n_i^2$  in the term  $i = j$  with  $n_i(n_i - 1)$  or  $\binom{n_i}{2}$  in the directed or undirected case respectively.

This kind of stochastic block model is well-known in the machine learning, statistics, and network communities (e.g. [5, 37, 15, 18, 19, 33]) and has also been used in ecology to identify groups of species in food webs [4]. Unlike e.g. [37, 18, 19], we do not assume that  $p_{ij}$  takes one value when  $i = j$  and a smaller value when  $i \neq j$ . In other words, we do not assume an assortative community structure, where nodes are more likely to be connected to other nodes of the same class. Nor do we require in general that  $p_{ij} = p_{ji}$ , since the directed nature of the edges may be important—for instance, in a food web or word adjacency network.

If all classifications  $t$  are equally likely *a priori*, then Bayes' rule implies that the Gibbs distribution on the classifications, i.e., the probability of  $t$  given  $G$ , is proportional to the probability of  $G$  given  $t$ :

$$P(t|G) \propto P(G|t). \quad (2)$$

In order to define  $P(G|t)$ , we need to integrate  $P(G|t, p)$  over some prior probability distribution on  $p$ . If we assume that the  $p_{ij}$  are independent, then this integral factors over the product (1). In particular, if each  $p_{ij}$  follows a beta prior, we have the Bayesian estimate of edge probabilities

$$\begin{aligned} P(G|t) &= \iint \int_0^1 d\{p_{ij}\} P(G|t, p) \\ &= \prod_{i,j=1}^k \int_0^1 dp_{ij} \text{Beta}(p_{ij}|\alpha, \beta) p_{ij}^{e_{ij}} (1 - p_{ij})^{n_i n_j - e_{ij}} \\ &= \prod_{i,j=1}^k \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} \\ &\int_0^1 dp_{ij} p_{ij}^{e_{ij} + \alpha - 1} (1 - p_{ij})^{n_i n_j - e_{ij} + \beta - 1} \\ &= \prod_{i,j=1}^k \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} \frac{\Gamma(e_{ij} + \alpha) \Gamma(n_i n_j - e_{ij} + \beta)}{\Gamma(n_i n_j + \alpha + \beta)}. \end{aligned} \quad (3)$$

For reasonable choices of the hyperparameters  $\alpha$  and  $\beta$ , the prior dominates only in small data cases, such as very small networks or sparsely populated classes. For such small data cases, the beta prior allows the user to input some domain knowledge about, say, the (dis)assortativity of the target network's community structure. In the limit of large data, the prior will wash out and the data-driven community structure will dominate.

If the user wishes to remain agnostic, however, he or she can specify a uniform prior ( $\alpha = \beta = 1$ ) and allow the learning algorithm to estimate the degree of (dis)assortativity (or other structure) entirely from the data. We take the latter approach in this paper, in which case

$$P(G|t) = \prod_{i,j=1}^k \frac{1}{(n_i n_j + 1) \binom{n_i n_j}{e_{ij}}}. \quad (4)$$

An even simpler approach is to assume that the  $p_{ij}$  take their maximum likelihood values

$$\hat{p}_{ij} = \underset{p}{\text{argmax}} P(G|t, p) = e_{ij}/n_i n_j, \quad (5)$$

and set  $P(G|t) = P(G|t, \hat{p})$ . This approach was used, for instance, for a hierarchical block model in [10]. When  $k$  is fixed and the  $n_i$  are large, this will give results similar to (4), since the integral over  $p$  is tightly peaked around  $\hat{p}$ . However, for any particular finite graph it makes more sense, at least to a Bayesian, to integrate over the  $p_{ij}$ , since they obey a posterior distribution rather than taking a fixed value. Moreover, averaging over the parameters as in (4) discourages overfitting, since the average likelihood goes down when we increase  $k$  and hence the volume of the parameter space. This gives us a principled way to determine  $k$  automatically, although in this paper we set  $k$  by hand. Other methods to determine  $k$  include minimum description length (MDL) techniques [33] and the Akaike information criterion [4].

## 4. ACTIVE LEARNING

In the active learning setting, the algorithm can learn the class label of any given node, but at a cost—say, by devoting resources in the laboratory or the field. Since these resources are limited, it has to decide which node to explore. Its goal is to explore a small set of nodes and use their labels to guess the labels of the remaining nodes.

One natural approach is to explore the node  $v$  with the largest mutual information (MI) between its label  $t(v)$  and the labels  $t(G \setminus v)$  of the other nodes according to the Gibbs distribution (2). We can write this as the difference between the entropy of  $t(G \setminus v)$  and its conditional entropy given  $t(v)$ ,

$$\text{MI}(v) = I(v; G \setminus v) = H(G \setminus v) - H(G \setminus v | v). \quad (6)$$

Here  $H(G \setminus v | v)$  is the entropy, averaged over  $t(v)$  according to the marginal of  $t(v)$  in the Gibbs distribution, of the joint distribution of  $t(G \setminus v)$  conditioned on  $t(v)$ . In other words,  $\text{MI}(v)$  is the expected amount of information we will gain about  $t(G \setminus v)$ , or equivalently the expected decrease in the entropy, that will result from learning  $t(v)$ . Since the mutual information is symmetric, we also have

$$\text{MI}(v) = I(v; G \setminus v) = H(v) - H(v | G \setminus v), \quad (7)$$

where  $H(v)$  is the entropy of the marginal distribution of  $t(v)$ , and  $H(v|G \setminus v)$  is the entropy, on average, of the distribution of  $t(v)$  conditioned on the labels of the other nodes. Thus  $MI(v)$  is large if (1) we are uncertain about  $v$ , so that  $H(v)$  is large, and (2)  $v$  is strongly correlated with the other nodes, so that  $H(v|G \setminus v)$  is small.

We estimate these entropies by sampling from the space of classifications according to the Gibbs distribution. Specifically, we use a single-site heat-bath Markov chain. At each step, it chooses a node  $v$  uniformly from among the unexplored nodes, and chooses its label  $t(v)$  according to the conditional distribution proportional to  $P(G|t)$ , assuming that the labels of all other nodes stay fixed. In addition to exploring the space, this allows us to collect a sample of the conditional distribution of the chosen node  $v$  and its entropy. Since  $H(v|G \setminus v)$  is the average of the conditional entropy, and since  $H(v)$  is the entropy of the average conditional distribution, we can write

$$I(v; G \setminus v) = - \sum_{i=1}^k \langle P_i \rangle \ln \langle P_i \rangle + \left\langle \sum_{i=1}^k P_i \ln P_i \right\rangle, \quad (8)$$

where  $P_i$  is the probability that  $t(v) = i$  and  $\langle \cdot \rangle$  denotes the average, according to the Gibbs distribution, over the labels of the other nodes.

We offer no theoretical guarantees about the mixing time of this Markov chain, and it is easy to see that there are families of graphs and values of  $k$  for which it takes exponential time. However, for the real-world networks we have tried so far, it appears to converge to equilibrium in a reasonable amount of time. We test for equilibrium by measuring whether the marginals change noticeably when the number of updates is increased by a factor of 2. We improve our estimates by averaging over many runs, each one starting from an independently random initial state.

We say that the algorithm is in *stage*  $j$  if it has already explored  $j$  nodes. In that stage, it estimates  $MI(v)$  for each unexplored node  $v$ , using the Markov chain to sample from the Gibbs distribution conditioned on the labels of the nodes explored so far. It then explores the node  $v$  with the largest  $MI$ . We provide it with the correct value of  $t(v)$ , and it moves on to the next stage.

The mutual information is not the only quantity we might use to identify which node to explore. Another is the *average agreement*, which we define as follows. Given two classifications  $t_1, t_2$ , define their *agreement* as the number of nodes on whose labels they agree,

$$|t_1 \cap t_2| = |\{v : t_1(v) = t_2(v)\}|. \quad (9)$$

Since our goal is to label as many nodes correctly as possible, what we would really like to maximize is the agreement between an assignment  $t_1$ , drawn from the Gibbs distribution, and the correct assignment  $t_2$ . But since we don't know  $t_2$ , we assume that it is drawn from the Gibbs distribution as well. Exploring  $v$  projects onto the part of the joint distribution of  $(t_1, t_2)$  where  $t_1(v) = t_2(v)$ . So, we define  $AA(v)$  as the expected agreement between two classifications  $t_1, t_2$  drawn independently from the Gibbs distribution, conditioned on the event that they agree at  $v$ . This




Figure 1: Zachary's Karate Club.

gives us the following quantity,

$$AA(v) = \frac{\sum_{t_1, t_2: t_1(v)=t_2(v)} P(t_1)P(t_2) |t_1 \cap t_2|}{\sum_{t_1, t_2: t_1(v)=t_2(v)} P(t_1)P(t_2)}. \quad (10)$$

We estimate the numerator and denominator of  $AA(v)$  using the same heat-bath Gibbs sampler as for  $MI(v)$ , except that we sample independent pairs of classifications  $(t_1, t_2)$  by starting the Markov chain at two independently random initial states.

## 5. RESULTS AND DISCUSSION

We tested our algorithms on three different networks from three different fields. The first is Zachary's Karate Club [40].

As shown in Fig. 1, this is a social network consisting of 34 members of a karate club, where undirected edges represent friendships. The club split into two factions, indicated by diamonds and circles respectively. One of them centered around the instructor (node 1) and the other around the club president (node 34), each of which formed their own club. Shaded nodes are more peripheral, and have weaker ties to their communities. This network is highly assortative, with a high density of edges within each faction and a low density of edges between them.

We judge the performance of each algorithm by asking, at each stage and for each node, with what probability the Gibbs distribution assigns it the correct label. In each stage we sample the Gibbs distribution using 100 independently chosen initial conditions, doing  $2 \times 10^4$  steps of the heat-bath Markov chain for each one, and computing averages using the last  $10^4$  steps. Increasing the number of Markov chain steps to  $10^5$  per stage produced only marginal improvements in performance. Fig. 2 shows what fraction of the unexplored nodes are assigned the correct label with probability at least  $q$ , for various thresholds  $q = 0.1, 0.3, 0.5, 0.7, 0.9$ , as a function of the stage  $j$ . After exploring just four or five nodes, our algorithms succeed in correctly predicting the labels of most of the remaining nodes—i.e., to which faction they belong—with high accuracy. The  $AA$  algorithm performs slightly better than  $MI$ , achieving an accuracy close to 100% after exploring nine nodes. Of course, this network is small, and there are many community-finding algorithms that classify the two factions with perfect or near-perfect accuracy [32, 12].




Figure 2: Results of the active learning algorithms on Zachary's Karate Club network.

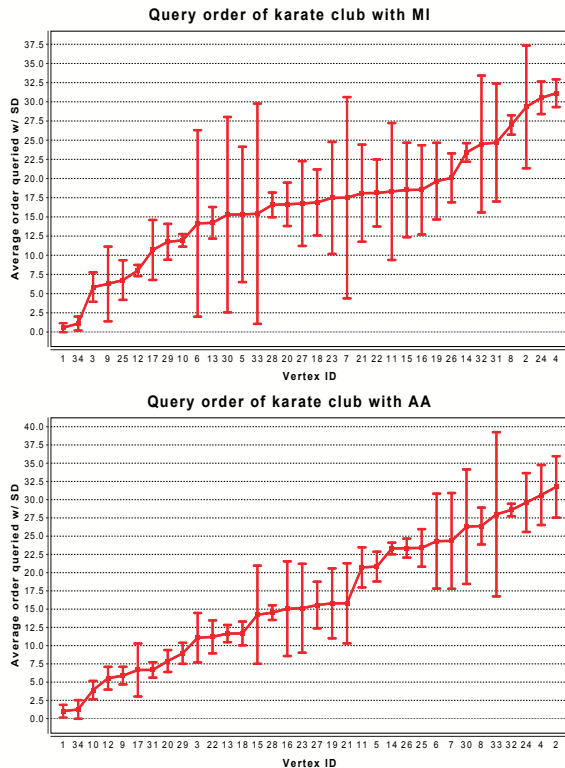


Figure 3: The order in which the active learning algorithms explore nodes in Zachary's Karate Club.

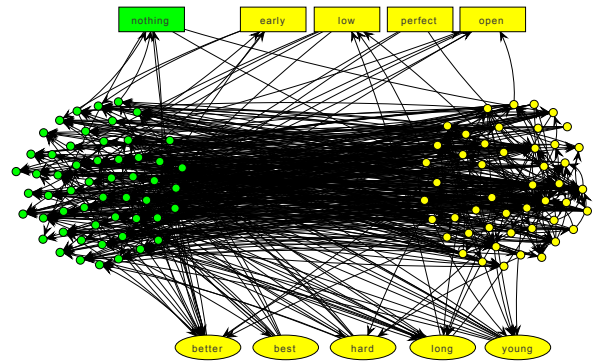



Figure 4: The order in which the active learning algorithm MI explores nodes in word adjacency network from the novel *David Copperfield*.

Perhaps more interesting is the *order* in which our algorithms choose to explore the nodes. In Fig. 3, we sort the nodes in order of the median stage at which they are explored. Error bars show 90% confidence intervals over 100 independent runs of each algorithm. Some nodes show a large variance in the stage in which they are explored, while others are reliably explored at the beginning or end of the process. Both algorithms start by exploring nodes 1 and 34, which are central to their respective communities. Note that these nodes are chosen, as we argued above, not just because their labels are uncertain, but because they are highly correlated with the labels of other nodes.

After learning that nodes 1 and 34 are in class 1 and 2 respectively, the algorithms “know” that the network consists of two assortative communities. They they explore nodes such as 3, 9, and 10 which lie at the boundary between these communities. Once the boundary is clear, they can easily predict the labels of the remaining nodes. The last nodes to be explored are those such as 2, 4, and 24, which lie so deep inside their communities that their labels are not in doubt.

The second network consists of the 60 most commonly occurring nouns and the 60 most commonly occurring adjectives in Charles Dickens' novel *David Copperfield*. A directed edge connects any pair of words that appear adjacently in the text, pointing from the preceding word to the following one. Excluding eight words which are disconnected from the rest leaves a network with 112 nodes [29]. Unlike Zachary's Karate Club, this network is both directed and highly disassortative. Of the 1494 edges, 1123 of them point from adjectives to nouns. This lets us classify most nodes early on, simply by labeling a node as an adjective or noun if its out-degree or in-degree is large.

Accordingly, our algorithms focus their attention on words about which they are uncertain, like “early,” “low,” and “nothing,” whose out-degrees and in-degrees in the text are roughly equal, and words like “perfect” that precede words of both classes (see Fig. 4, where green and yellow nodes represent nouns and adjectives respectively; rectangular nodes are ex-



**Figure 5: Results of the active learning algorithms on word adjacency network in the novel *David Copperfield* by Charles Dickens.**

plored first, and elliptical ones last). Once these nodes are resolved, both algorithms achieve high accuracy—80% accuracy after exploring 20 nodes and close to 100% after exploring 65 nodes (see Fig. 5).

In each stage we sample the Gibbs distribution using 100 independently chosen initial conditions, doing  $5 \times 10^4$  steps of the heat-bath Markov chain for each one, and computing averages using the last  $2.5 \times 10^4$  steps. Increasing the number of Markov chain steps to  $10^5$  per stage produced only marginal improvements in performance. As in Fig. 2, the  $y$ -axis shows the fraction of unexplored nodes which are labeled correctly by the conditional Gibbs distribution with probability at least  $q$ , for  $q = 0.1, 0.3, 0.5, 0.7, 0.9$ . The performance of the two algorithms is similar in the later stages, but unlike the Karate Club, here MI performs noticeably better than AA in the early stages.

The third network is a food web of 488 species in the Weddell Sea in the Antarctic [11, 9, 22], with edges pointing to each predator from its prey. This data set is very rich, but we focus on two particular variables—the feeding type and the habitat in which the species lives. The feeding type takes  $k = 6$  values, namely primary producer, omnivorous, herbivorous/detrivorous, carnivorous, detrivorous, and carnivorous/necrovorous. The habitat variable takes  $k = 5$  values, namely pelagic, benthic, benthopelagic, demersal, and land-based.

We show results of our algorithms for both variables in Fig. 6. The results are averaged over 100 runs of each algorithm. In each stage we sample the Gibbs distribution using 100 independently chosen initial conditions, doing  $5 \times 10^4$  steps of the heat-bath Markov chain for each one, and computing averages using the last  $2.5 \times 10^4$  steps. For the feeding type, after exploring half the nodes, both algorithms correctly label about 75% of the remaining nodes. For the habitat variable, both algorithms are less accurate, although AA performs somewhat better than MI. Note that the accuracy only includes the unexplored nodes, not the nodes we have already explored. Thus it can decrease if we explore easily-classified nodes early on, so that hard-to-classify nodes form a larger fraction of the remaining ones.

Fig. 6 shows that both algorithms get to a state where they are confident, but wrong, about many of the unexplored nodes. For the feeding type variable, for instance, after the AA algorithm has explored 300 species, it labels 75% of the remaining nodes correctly with probability 90%, but it labels the other 25% correctly with probability less than 10%. In other words, it has a high degree of confidence about all the nodes, but is wrong about many of them. Its accuracy improves as it explores more nodes, but it doesn’t achieve high accuracy on all the unexplored nodes until there are only about 60 of them left.

Why is this? We argue that the fault lies, not with our learning algorithms and the order in which they explore the nodes, but with the stochastic block model and its ability to model the data. For example, for the habitat variable, these algorithms perform well on pelagic, demersal, and land-based species. But the benthic habitat, which is the largest and most diverse, includes species with many feeding types and trophic levels.

These additional variables have a large effect on the topology, but they are not taken into account by the block model. As a result, more than half the benthic species are mislabeled by the block model in the following sense: even if we condition on the correct habitats of *all* the other species, the species’ most likely habitat is pelagic, benthopelagic, demersal, or land-based. Specifically, 219 of the 488 species are mislabeled by the most likely block model, 94% of them with confidence over 0.9. Of course, we can also regard our algorithms’ mistakes as evidence that these habitat classifications are not cut and dried. Indeed, ecologists recognize that there are “connector species” that connect one habitat to another, and belong to some extent to both.

To test our hypothesis that it is the block model’s inability to model the data that causes some nodes to be misclassified, we artificially modified the data set to make it consistent with the block model. Starting with the original variables, we updated the habitat of each species to its most likely value according to the block model, given the habitats of all the other species. After iterating this process six times, we reached a fixed point where each species’ habitat is consistent with the block model’s predictions. On this synthetic



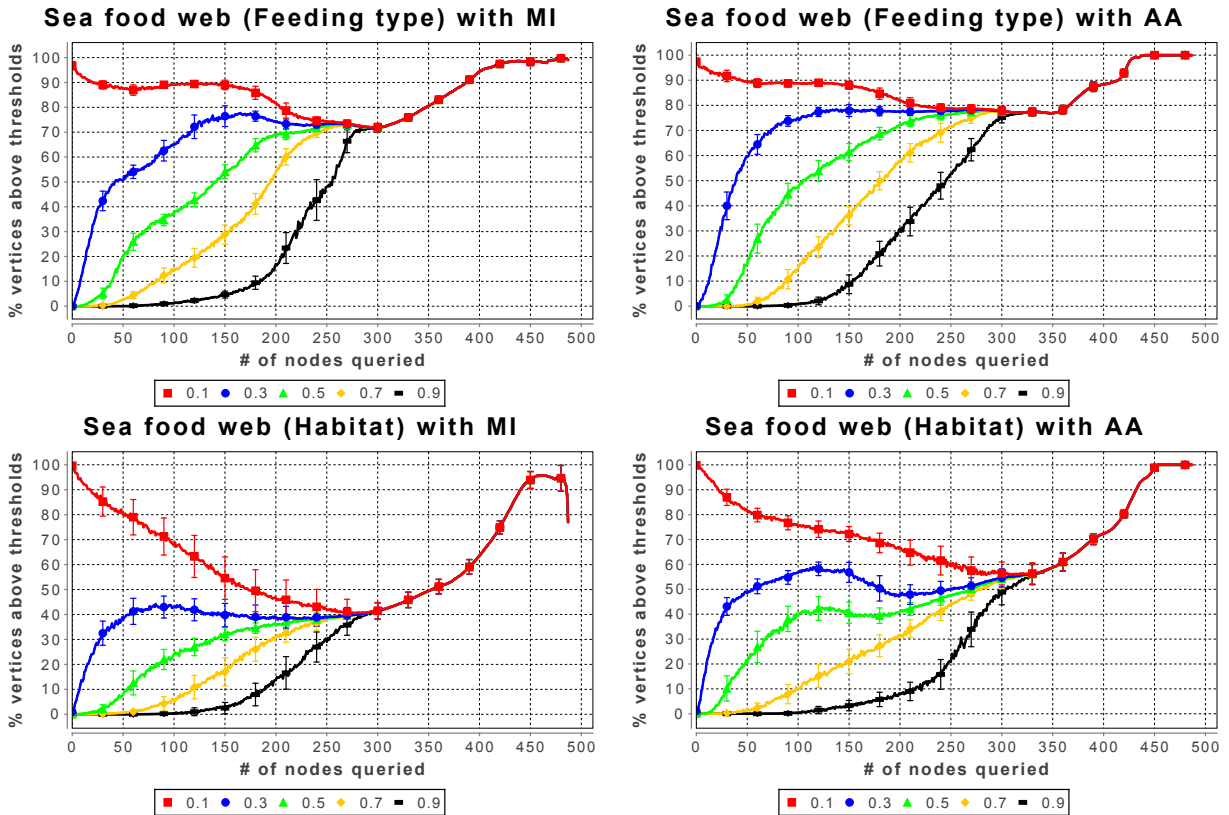


Figure 6: Results for the Weddell Sea food web.

data set both of our learning algorithms perform perfectly, predicting the habitat of every species with close to 100% accuracy after exploring just 18% of them.

More generally, it’s important to remember that the topology of the network is only imperfectly correlated with the nodes’ variables. Zachary [40] notes that one of members of the Karate Club joined the instructor’s faction even though the network’s topology suggests that he was more strongly connected to the president. The reason is that he was only three weeks away from a test for his black belt when the split occurred. He had already invested four years learning the instructor’s style of karate, and if he had joined the president’s club he would have had to start over with a white belt. In any real-world network, there is information of this kind that is not reflected in the topology and which is hidden from our algorithm. If a node is of a given class for idiosyncratic reasons like these, we cannot expect any algorithm based solely on topology and the other nodes’ class labels—no matter how sophisticated a probabilistic model we use—to correctly classify it.

## 6. COMPARISON WITH SIMPLE HEURISTICS

We compared our active learning algorithms with some simple heuristics. These include exploring the node with highest degree in the subgraph of unexplored nodes, exploring the node with highest betweenness centrality (the fraction of shortest paths that go through it, see [8, 27, 28]) in the

subgraph of unexplored nodes, and exploring a node chosen uniformly at random from the unexplored ones. We judge the performance of these heuristics using the same Gibbs sampling process as for MI and AA.

In Fig. 7, we show the results of these heuristics at the 0.9 accuracy threshold on all three networks, including both the habitat and feeding type variables in the food web. On Zachary’s Karate Club (left) our algorithms outperform these heuristics consistently. In the *David Copperfield* network (right), the high-degree and high-betweenness heuristics enjoy an early lead, but quickly hit a ceiling and are surpassed by MI and AA.

For the Weddell Sea food web (bottom), the high-degree and high-betweenness heuristics perform poorly throughout the process. One reason for this is that many nodes with high degree or high betweenness are easy to classify from the labels of their neighbors. By exploring these nodes first, these heuristics leave themselves mainly with hard-to-classify nodes. The random heuristic performs surprisingly well early on, but all three heuristics are worse than MI or AA once they have explored half the nodes.

## 7. CONCLUSION

Active learning, using mutual information or average agreement coupled with a generative model, offers a new approach to analyzing networks where our knowledge of class labels is incomplete and costly to obtain. We have shown for three networks, one social, one lexical, and one biological, that

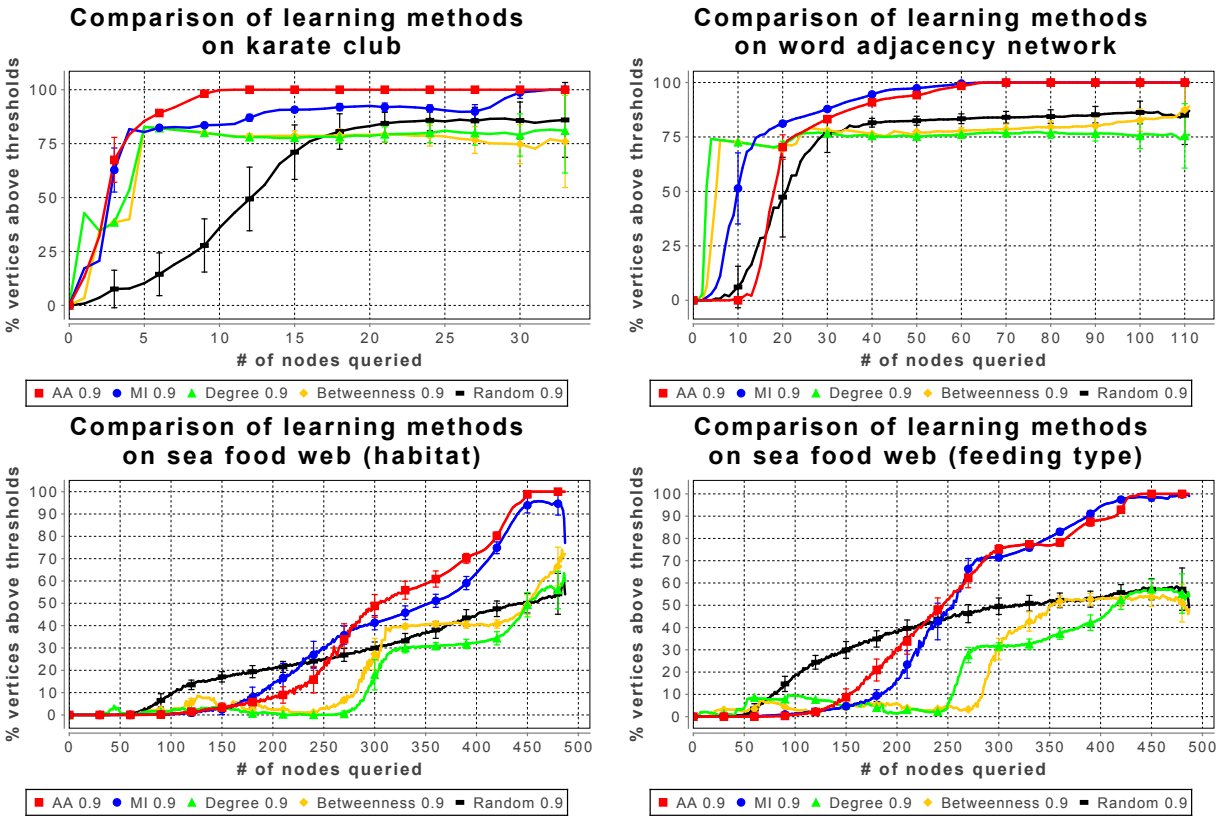


Figure 7: A comparison of the MI and AA learning algorithms with three simple heuristics.

our algorithms do a good job of predicting the labels of unexplored nodes after exploring a relatively small fraction of the network, correctly recognizing both assortative and disassortative functional communities. Certainly not all networks are well-described by the simple block model we use here, but our approach can be generalized to probabilistic network models which take information on the nodes' locations or degrees into account.

**Acknowledgments.** We are grateful to Joel Bader, Aaron Clauset, Jennifer Dunne, Nathan Eagle, Brian Karrer, Jon Kleinberg, Mark Newman, Cosma Shalizi, and Jerry Zhu for helpful conversations, and to Ute Jacob for the Weddell Sea food web data. J.-B. R. is also grateful to the Santa Fe Institute for their hospitality. This work was supported by the McDonnell Foundation.

## 8. REFERENCES

- [1] L. Adamic and N. Glance. The political blogosphere and the 2004 US election: Divided they blog. In *Proc 3rd Intl Workshop on Link Discovery.*, 2005.
- [2] E. M. Airoldi, D. M. Blei, S. E. Fienberg, and E. P. Xing. Mixed membership stochastic blockmodels. *J. Machine Learning Research*, 9:1981–2014, 2008.
- [3] D. Alderson, L. Li, W. Willinger, and J. C. Doyle. Understanding internet topology: principles, models, and validation. *IEEE/ACM Trans. Networks*, 13(6):1205–1218, 2005.
- [4] S. Allesina and M. Pascual. Food web models: a plea for groups. *Ecology Letters*, 12:652–662, 2009.
- [5] P. J. Bickel and A. Chen. A nonparametric view of network models and newman-girvan and other modularities. *Proc. Natl. Acad. Sci.*, 106:21068–21073, 2009.
- [6] M. Bilgic and L. Getoor. Link-based active learning. In *NIPS Workshop on Analyzing Networks and Learning with Graphs*, 2009.
- [7] M. Bilgic, L. Mihalkova, and L. Getoor. Active learning for networked data. In *Proc. Intl. Conf. on Machine Learning*, 2010.
- [8] U. Brandes. A faster algorithm for betweenness centrality. *Journal of Mathematical Sociology*, 25(2):163–177, 2001.
- [9] U. Brose, L. Cushing, E. L. Berlow, T. Jonsson, C. Banasek-Richter, L. F. Bersier, J. L. Blanchard, T. Brey, S. R. Carpenter, M. F. Blandenier, et al. Body sizes of consumers and their resources. *Ecology*, 86(9):2545–2545, 2005.
- [10] A. Clauset, C. Moore, and M. E. J. Newman. Hierarchical structure and the prediction of missing links in networks. *Nature*, 453(7191):98–101, 2008.
- [11] U. B. et al. Consumer-resource body-size relationships in natural food webs. *Ecology*, 87(10):2411–2417, 2006.
- [12] S. Fortunato. Community detection in graphs. *Physics Reports*, 2009.
- [13] L. Gao and J. Rexford. Stable internet routing



- without global coordination. *IEEE/ACM Trans. Networks*, 9(6):681–692, 2001.
- [14] A. B. Goldberg, X. Zhu, and S. Wright. Dissimilarity in graph-based semi-supervised classification. *J. Machine Learning Research W&P*, 2:155–162, 2007.
- [15] R. Guimera and M. Sales-Pardo. Missing and spurious interactions and the reconstruction of complex networks. *Proc. Natl. Acad. Sci.*, 106:22073–22078, 2009.
- [16] Y. Guo and R. Greiner. Optimistic active learning using mutual information. In *Proc. Intl. Joint Conf. on Artificial Intelligence*, 2007.
- [17] M. S. Handcock, A. E. Raftery, and J. M. Tantrum. Model-based clustering for social networks. *J. Royal Statist. Soc. A*, 170(2):1–22, 2007.
- [18] M. B. Hastings. Community detection as an inference problem. *Physical Review E*, 74(3):035102, 2006.
- [19] J. M. Hofman and C. H. Wiggins. Bayesian approach to network modularity. *Physical Review Letters*, 100(25):258701, 2008.
- [20] P. W. Holland, K. B. Laskey, and S. Leinhardt. Stochastic blockmodels: first steps. *Social networks*, 5:109–137, 1983.
- [21] M. Houseman and D. R. White. *Taking Sides: Marriage Networks and Dravidian Kinship in Lowland South America*, pages 214–243. Transformations of Kinship. Smithsonian Institution Press, 1998.
- [22] U. Jacob. *Trophic Dynamics of Antarctic Shelf Ecosystems, Food Webs and Energy Flow Budgets*. PhD thesis, University of Bremen, 2005.
- [23] B. Karrer and M. E. J. Newman. Stochastic blockmodels and community structure in networks, 2010. Preprint, arXiv:1008.3926v1.
- [24] D. V. Lindley. On a measure of the information provided by an experiment. *Ann. Math. Statist.*, 27(4):986–1005, 1956.
- [25] D. J. MacKay. Information-based objective functions for active data selection. *Neural Computation*, 4(4):590–604, 1992.
- [26] M. Mørup and L. K. Hansen. Learning latent structure in complex networks. In *NIPS Workshop on Analyzing Networks and Learning with Graphs*, 2009.
- [27] M. Newman. Scientific collaboration networks. II. shortest paths, weighted networks, and centrality. *Physical Review E*, 64(1), 2001.
- [28] M. E. J. Newman. A measure of betweenness centrality based on random walks. *Social Networks*, 27(1):39–54, 2005.
- [29] M. E. J. Newman. Finding community structure in networks using the eigenvectors of matrices. *Physical Review E*, 74(3):36104, 2006.
- [30] M. E. J. Newman and E. A. Leicht. Mixture models and exploratory analysis in networks. *Proc. Natl. Acad. Sci.*, 104:9564–9569, 2006.
- [31] A. S. Patterson, Y. Park, and J. S. Bader. Degree-corrected block models. Manuscript.
- [32] M. A. Porter, J. P. Onnela, and P. J. Mucha. Communities in networks. *Notices of the American Mathematical Society*, 56(9):1082–1097, 2009.
- [33] M. Rosvall and C. T. Bergstrom. An information-theoretic framework for resolving community structure in complex networks. *Proc. Natl. Acad. Sci.*, 104(18):7327, 2007.
- [34] N. Roy and A. McCallum. Toward optimal active learning through sampling estimation of error reduction. In *Proc. 18th Intl. Conf. on Machine Learning*, pages 441–448, 2001.
- [35] P. Sen, G. M. Namata, M. Bilgic, L. Getoor, B. Gallagher, and T. Eliassi-Rad. Collective classification in network data. *AI Magazine*, 29(3):93–106, 2008.
- [36] W. Tong and R. Jin. Semi-supervised learning by mixed label propagation. In *Proc. 22nd Intl. Conf. on Artificial intelligence*, volume 1, pages 651–656, 2007.
- [37] J. Čopič, M. O. Jackson, and A. Kirman. Identifying community structures from network data. *B.E. Press Journal of Theoretical Economics*, 9(1):Article 30, 2009.
- [38] Y. J. Wang and G. Y. Wong. Stochastic blockmodels for directed graphs. *J. American Statistical Assn.*, 82(397):8–19, 1987.
- [39] R. J. Williams, A. Anandanadesan, and D. Purves. The probabilistic niche model reveals the niche structure and role of body size in a complex food web. *PLoS One*, 5(8):e1209, 2010.
- [40] W. W. Zachary. An information flow model for conflict and fission in small groups. *J. Anthropological Research*, 33(4):452–473, 1977.
- [41] X. Zhu, J. Lafferty, and Z. Ghahramani. Combining active learning and semi-supervised learning using gaussian fields and harmonic functions. In *Proc. ICML-2003 Workshop on the Continuum from Labeled to Unlabeled Data*, 2003.