

RESEARCH

Open Access



Understanding the limitations of network online learning

Timothy LaRock^{1*} , Timothy Sakharov¹, Sahely Bhadra³ and Tina Eliassi-Rad^{1,2}

*Correspondence:

larock.t@husky.neu.edu

¹Network Science Institute,
Northeastern University, 360
Huntington Ave, Boston MA,
02115-5000 USA
Full list of author information is
available at the end of the article

Abstract

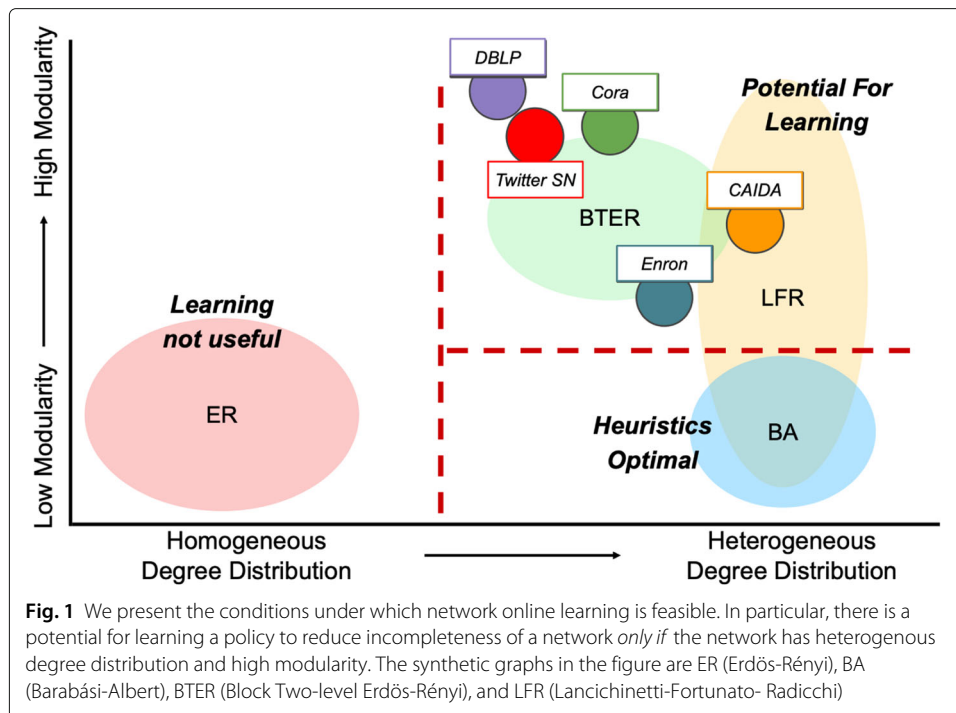
Studies of networked phenomena, such as interactions in online social media, often rely on incomplete data, either because these phenomena are partially observed, or because the data is too large or expensive to acquire all at once. Analysis of incomplete data leads to skewed or misleading results. In this paper, we investigate limitations of learning to complete partially observed networks via node querying. Concretely, we study the following problem: given (i) a partially observed network, (ii) the ability to query nodes for their connections (e.g., by accessing an API), and (iii) a budget on the number of such queries, sequentially learn which nodes to query in order to maximally increase observability. We call this querying process **Network Online Learning** and present a family of algorithms called **NOL***. These algorithms learn to choose which partially observed node to query next based on a parameterized model that is trained online through a process of exploration and exploitation. Extensive experiments on both synthetic and real world networks show that (i) it is possible to sequentially learn to choose which nodes are best to query in a network and (ii) some macroscopic properties of networks, such as the degree distribution and modular structure, impact the potential for learning and the optimal amount of random exploration.

Keywords: Partially observed networks, Online learning, Heavy-tailed target distributions

Introduction

Incomplete datasets are common in the analysis of networks because the phenomena under study are often partially observed. It has been shown that analysis of incomplete networks may lead to biased results (Sanz et al. 2012; Ghosh et al. 2013; González-Bailón et al. 2014; Sampson et al. 2015; Alves et al. 2020). Our work seeks to address the following problem: *Given a partially observed network with no information about how it was observed and a budget to query the partially observed nodes, can we learn to sequentially ask optimal queries relative to some objective?* In addition to introducing new methodology, we study when learning is feasible in this problem (see Fig. 1).

We present a family of online learning algorithms, called **NOL***, which are based on online regression and related to reinforcement learning techniques like approximate Q-learning. **NOL*** algorithms do not assume any a priori knowledge or estimate of the true



underlying network structure, overall size, or the sampling method used to collect the partially observed network.

We describe two algorithms from the NOL* family, both of which learn an *interpretable* policy for growing an incomplete network through successive node queries. Interpretability of these algorithms is important because we want to be able to reason about when and why a policy is or is not learnable. The first algorithm, referred to as NOL¹, uses online linear regression as a model for predicting the value of querying available nodes, then uses those predictions to choose the best node to query next. The second algorithm, NOL-HTR, uses a heavy-tail regression method to account for heavy-tailed (equivalently heterogeneous) reward functions. This is necessary because in the case of a heavy-tailed reward distribution such as node degree, NOL will under-predict the objective for hubs (which are the “big and rare” instances).

We conduct experiments using NOL* algorithms on graphs generated by synthetic models, as well as real-world network data.² We focus on the following objective or reward function: discover the maximum number of initially unobserved nodes in the network through successive node querying.

Formal problem definition Given an *incomplete* network $\hat{G}_0 = \{\hat{V}_0, \hat{E}_0\}$, which is a partial observation of an underlying network $G = \{V, E\}$, sequentially learn a policy that maximizes the number of nodes $u \in \hat{V}_b, u \notin \hat{V}_0$ after b queries of the incomplete network. Querying the network involves selecting a node and asking an oracle or an API for all the neighbors of the selected node.

¹An earlier version of NOL was presented at the 14th Annual Workshop on Mining and Learning with Graphs, a non-archival venue, in 2018 (LaRock et al. 2018).

²We use the terms graph and network interchangeably in this paper.

This problem is a sequential decision-learning task that can be formulated as a Markov Decision Process (MDP), where the *state* of the process at any time step is the partially observed network, the *action space* is the set of partially observed nodes available for probing, and the *reward* is a user-defined function (e.g., the increase in the number of observed nodes). The goal of an MDP learning algorithm is to learn a mapping from states to actions such that an agent using this mapping will maximize its expected reward. In our case, the state-action space of the problem can be arbitrarily large given that we make no assumptions about the underlying network; thus any solution will need to learn to generalize over states and actions from experience (i.e., answers to successive queries).

Our framework NOL* can be viewed as solving an MDP problem. Its algorithms learn models to predict the expected reward gained by probing³ a partially observed node. The current model is then used at each time step to decide the best action (i.e., which node to query to observe as many new nodes as possible). In this way, NOL* algorithms choose the action that leads to maximizing the total reward over time within an arbitrary budget constraint.

NOL* algorithms query nodes in the partially observed network *sequentially*. That is, in each iteration the algorithm queries one node, adds all of its neighbors to the observed network, and updates the pool of partially observed nodes available to be queried in the next iteration. This makes NOL* algorithms *adaptive*, since the parameters of the model change based on recent experience. Initially, all nodes in the network are assumed to be partially observed, thus NOL* is agnostic to the underlying observation or sampling method. At any iteration, there are three “classes” of nodes: fully observed (probed), partially observed (unprobed but visible) and unobserved (unprobed and invisible).

In the present work, we assume networks that are undirected, unweighted, and static, meaning that the data being queried does not change over time (no insertion or deletion of nodes or edges). However, we note that our approach is flexible enough to be modified to incorporate features and objective functions that take edge directionality and weight into account. For example, one could incorporate in and out degree as features in the case of a directed network, and weighted degree in weighted networks. Our methodology can also be applied to discover previously unobserved nodes in a bipartite network, as long as there is a suitable method for querying the data. Finally, our method can be extended to allow for repeated probing of nodes under any API access model, which is useful when the network is not static.

Contributions Our contributions are as follows:

- We propose a family of interpretable algorithms, NOL*, for learning to grow incomplete networks.
- We present two algorithms from the NOL* family: an online regression algorithm, simply called NOL; and an algorithm that can effectively learn for heavy-tailed objective functions, called NOL-HTR.
- Our extensive experiments on both synthetic and real network data showcase the limitations of online learning to improve incomplete networks.

³We use the terms querying and probing interchangeably in this paper.

In the next section we summarize the literature related to the network discovery problem. Then, we describe and evaluate NOL* before closing the paper with a discussion of future directions.

Related work

Incomplete data affects numerous areas of research. It has been shown to be a problem in social networks, including in public health (Gile 2011; Wejnert and Heckathorn 2008) and economics (Breza et al. 2017), as well as the mining of large systems such as the World Wide Web (WWW) (Cho et al. 1998; Avrachenkov et al. 2014) or Internet infrastructure (Vázquez et al. 2002). It is also a problem for understanding gene regulatory networks (Sanz et al. 2012). It has also been shown that over-simplified models, for example those that do not incorporate directionality or weight of edges, can lead to anomalies in the analysis of centrality (Alves et al. 2020).

The problem of network online learning is different from network sampling. In traditional network sampling, the goal is to gather a representative sample of the underlying network from which statistical characteristics are then estimated. In our setting, the data collection is guided by an initial sample graph and a user-defined objective function, which may or may not be directly related to a notion of statistical representativeness of the data. For an excellent survey on network sampling, we refer the reader to (Ahmed et al. 2013).

One approach to growing incomplete networks is to assume that the graph is being generated by an underlying network model, then use that model to infer the missing nodes and links. Examples of this approach are (Hanneke and Xing 2019) and (Peixoto 2018), both of which model and address sampling errors within a Stochastic Block Model (Wang and Wong 1987) framework. Another example is (Kim and Leskovec 2011), which assumes Kronecker graph models (Leskovec et al. 2010). Finally, we refer the reader to (Chen et al. 2019) for a recent paper on model selection for mechanistic network models.

Avrachenkov et al. (2014) explore methods for maximally covering a graph by adaptive crawling and querying. Their work introduces the Maximum Expected Uncovered Degree (MEUD) method and shows that for a certain class of random power law networks, MEUD reduces to choosing the node with maximum observed degree for each query (which is equivalent to our high degree heuristic; see [Experiments](#) section).

MAXOUTPROBE (Soundarajan et al. 2015) estimates the degree of each observed node and the average clustering coefficient of the graph; it does not assume knowledge of how the incomplete network was collected. MAXREACH (Soundarajan et al. 2016) estimates the degree of each observed node and the per-degree average clustering coefficient. It assumes that the incomplete network was collected via random node or edge sampling and that one knows the number of nodes and edges in the fully observed graph.

Multi-armed bandit approaches are well-suited for the problem of growing incomplete networks because they are designed explicitly to facilitate the exploration vs. exploitation tradeoff. Soundarajan et al. (2017) present a multi-armed bandit approach to network completion that trades off densification vs. expansion. Their approach also focuses on the problem of probing edges rather than nodes. Murai et al. (2018) propose a multi-armed bandit algorithm for reducing network incompleteness that probabilistically chooses from an ensemble of classifiers that are trained simultaneously. Madhawa and Murata (2019) describe a nonparametric multi-armed bandit based on a k-nearest

neighbor upper-confidence bound algorithm, which will be described in more detail in [Experiments](#) section.

Active Exploration (Pfeiffer III et al. 2014) and Selective Harvesting (Murai et al. 2018) address variants of the problem definition where the general task is to iteratively *search* a partially observed network through node querying. The goal is to maximize the number of nodes in the expanded network with a particular binary target attribute. In this paper we address a different problem, where the goal is to maximally grow the network regardless of particular node attributes.

In the following section, we introduce NOL*, which provides a unified framework for interpretable and scalable network online learning. It incorporates exploration vs. exploitation and does not assume how the incomplete network was originally collected or a specific model generating the underlying network.

NOL* family of algorithms

Algorithm 1 presents the general NOL* framework. The goal is to sequentially learn to predict the reward value of partially observed nodes in an incomplete network under a resource constraint or budget, b , on the number of queries we can ask an oracle or API, leveraging previous queries as sample data. One node is queried at every time step $t = 0, 1, 2, \dots b$. The partially observed network $\hat{G}_t = \{\hat{V}_t, \hat{E}_t\} \subset G$, with nodes V_t , edges E_t and the list of nodes which have been already probed P_t , are updated by the algorithm after every query. The incomplete network \hat{G}_t grows after every probe by incorporating the neighbors of the probed node j , selected from $(\hat{V}_t - P_t)$. The number of new nodes added to the observed network by probing j in timestep t is the *reward* earned in that timestep, denoted by $r_t(j)$.

The goal is to make a prediction about the reward to be earned by querying any partially observed node j that maximizes reward at time t . In general, given an initial state s_0 , we wish to maximize the *cumulative reward* earned after b queries, given that the starting point was s_0 :

$$c_r(b) = \sum_{t=0}^b r_t | s_0$$

In each time step t , we want to maximize the reward r_t earned by taking action a from state s_t . Therefore we learn to choose an action, which corresponds to choosing a node u to query, such that

$$u = \operatorname{argmax}_{j \in (\hat{V}_t - P_t)} r_t(j)$$

where $r_t(j)$ is the reward earned by querying node j from state s_t .

Although a more general Temporal Difference Learning (Sutton and Barto 2018) solution to this problem can be formulated, we did not find an advantage in using discounting or credit assignment in experiments. We conjecture that this is due to a combination of the finite resource constraint and the fact that we do not visit the same states multiple times in our setting, making standard planning tools less effective.

To predict the reward value of every node in \hat{G}_t , a feature vector $\phi_t(j) \in \mathbb{R}^d$ is maintained that represents the knowledge available to the model at time t for node $j \in \hat{G}_t$. Given features $\phi_t(i)$ for all nodes in \hat{G}_t , NOL* algorithms learn a function $\mathcal{V}_\theta : \mathbb{R}^d \rightarrow \mathbb{R}$

with parameter θ to predict the expected reward to be earned by probing the node. Any available information about a node can be included as a feature, but since learning is to happen online, features should be feasible for online computation. Further, it is desirable for features to be interpretable so that the resulting regression parameters can be interpreted to help understand the performance of the algorithm in a given dataset (see [Appendix A.2](#)).

Given parameter θ_t at time t , the predicted number of unobserved nodes attached to node $j \in \hat{V}_t$ is estimated as $\mathcal{V}_{\theta_t}(j) = \theta_t \phi_t(j)$ such that $\mathcal{V}_{\theta_t}(j)$ should be close to the observed reward value of probing the node j at time t . At each step, the expected loss $E_t[\mathcal{V}_{\theta_t}(j) - r_t(j)]$ is minimized, where $r_t(j)$ is the true value of the reward function for node j at time t .

In NOL, θ is updated after each probe through online stochastic gradient descent based on Strehl and Littman (2007) (see lines 10-13 of Algorithm 1). If the variable of interest is heavy-tailed (e.g., the degree distributions of real networks exhibit this property), NOL-HTR adopts the generalized median of means approach to regression with heavy tails found in Hsu and Sabato (2016). These example functions illustrate the flexibility of

Algorithm 1 NOL* Framework

Require: \hat{G}_0 (initial incomplete network), b (probing budget), ϵ_0 (initial jump rate), k and λ (for NOL-HTR), α (for NOL).

Ensure: \hat{G}_b (network after b probes), θ_b (learned parameters).

- 1: Initialize: θ_0 (randomly or heuristically); $P_0, X, Y = \emptyset$; $t = 0$; $\epsilon_t = \epsilon_0$
 - 2: **repeat**
 - 3: Calculate feature vectors: $\phi_t(i), \forall$ nodes $i \in \hat{V}_t - P_t$
 - 4: Calculate estimated rewards: $\mathcal{V}_{\theta_t}(i) = \theta_t \phi_t(i)$
 - 5: Explore: With probability ϵ_t , choose u_t uniformly at random from $\hat{V}_0 - P_t$.
 - 6: Exploit: With probability $1 - \epsilon_t$,
 $u_t = \operatorname{argmax}_{i \in \hat{V}_t - P_t} \mathcal{V}_{\theta_t}(i)$
 - 7: Probe node u_t , add to $P_{t+1} = P_t \cup u_t$
 - 8: Update the observed graph: $\hat{G}_{t+1} = \{\hat{G}_t \cup \text{neighbors of } u_t\}$
 - 9: Collect reward: $r_t(u_t) = |\hat{V}_{t+1}| - |\hat{V}_t|$
 - 10: **if** NOL **then**
 - 11: $loss_t = (r_t - \mathcal{V}_{\theta_t}(u_t))^2$
 - 12: $\nabla_{\theta_t} loss_t = -2(r_t - \mathcal{V}_{\theta_t}(u_t)) \phi_t(u_t)$
 - 13: $\theta_{t+1} = \theta_t - \alpha \nabla_{\theta_t} loss_t$
 - 14: **else if** NOL-HTR **then**
 - 15: Append $\phi_t(u_t)$ to X and r_t to Y
 - 16: $\theta_{t+1} = \text{HeavyTailRegression}(X, Y, k, \lambda)$
 - 17: **end if**
 - 18: $t = t + 1$
 - 19: **if** decay **then**
 - 20: $\epsilon_t = \epsilon_0 e^{-\frac{t}{b}}$
 - 21: **end if**
 - 22: **until** $t=b$
 - 23: **return** θ_b and \hat{G}_b
-

NOL*: the choice of reward function and learning algorithm should correspond to individual goals and circumstances. We describe the details of how we learn the parameters for NOL-HTR in the next section and Algorithm 2.

Algorithm 2 *HeavyTailRegression*

Require: \mathbf{X} (matrix of probed node features), Y (associated reward values), k , and λ

Ensure: θ (updated parameters)

- 1: Partition data into k uniformly random subsamples S_1, \dots, S_k of size $\frac{n}{k}$, where each $S_i = \mathbf{X}_i \subset \mathbf{X}, Y_i \subset Y$
 - 2: **for** each subsample S_i **do**
 - 3: Compute covariance Σ_i of \mathbf{X}_i
 - 4: Compute regression parameters $\omega_i \in \operatorname{argmin}_{\omega} L^{\lambda}(\mathbf{X}_i, Y_i)$
 - 5: **end for**
 - 6: **for** each subsample S_i **do**
 - 7: Compute $m_i = \mathbf{median}_{j \neq i}((\omega_i - \omega_j), (\Sigma_{S_j} + \lambda Id)(\omega_i + \omega_j))$
 - 8: **end for**
 - 9: Assign θ to be the ω_i associated with the minimum m_i
 - 10: **return** θ
-

Parameter estimation for NOL-HTR

Our goal at every time t is to choose a node that maximizes the earned reward. However, since the reward distribution is based on node degree, and the degree distribution in many real-world networks is heavy-tailed (e.g. hubs are present), we expect the reward distribution of an incomplete version of the network to be heavy-tailed as well. In order to deal with the heavy tailed nature of the target variable, we adapt the methods from (Hsu and Sabato 2016) for regression in the presence of heavy-tailed distributions. This process is a generalization of the median of means approach to parameter estimation. Intuitively, the algorithm splits the previously observed feature vectors and associated rewards into k subsamples, then computes parameters ω for each subsample. Then, the algorithm chooses the set of parameters that has the minimum median distance from all of the other parameters. This procedure provides guarantees and confidence bounds on the distance of the learned parameters from the true parameters (Hsu and Sabato 2016), discussed further in [Scalability & guarantees](#) section.

Algorithm 2 presents our adopted process. At time t , $t - 1$ nodes with feature vectors $x_i \in X$ have been probed, and their corresponding reward values $r_i \in Y$ observed. We select an integer $k \leq t$ and randomly sample the data into k subsamples S_i of size $\frac{t}{k}$. Then, the covariance matrix Σ_i and maximum likelihood regression estimate ω_i are computed for each S_i . For each i , the Mahalanobis distance between ω_i and every other ω_j is computed, and the median distance is assigned to m_i . Finally, the ω_i with the minimum m_i value is assigned to be the next set of parameters, θ .

There are two parameters in Algorithm 2: the number of subsamples k , which corresponds to the confidence parameter δ in Hsu and Sabato (2016), and the regularization parameter λ . The number of subsamples k should be set such that the size of the subsamples, $\frac{n}{k}$, is larger than $O(d \log(d))$, meaning each subsample has size at least the number of

dimensions in the feature vector (Hsu and Sabato 2016). In our experiments, we set k to be $\ln(n)$, where n is the number of previously queried nodes (equivalently the time step t), which allows k to grow slowly as more data is gathered through the querying process. We set the regularization parameter λ to 0. However, our feature matrices may be singular (since some subsamplings can result in feature matrices without full rank), so when computing the regression in Algorithm 2 we use the Moore–Penrose pseudoinverse, which corresponds to computing the ℓ_2 regularized parameters.

ϵ -greedy exploration

NOL* algorithms need to learn adaptively over time because the reward distribution may change as nodes are queried. Choosing a node based on V_{θ_t} exploits the current model by choosing the node with the maximum predicted reward. However, learning in networks is difficult precisely because the properties of nodes are diverse, thus it is desirable to introduce some randomness to the decision process in order to gather a diverse set of training examples. Therefore, we formulate NOL* algorithms as ϵ -greedy algorithms, meaning with probability ϵ the node to query is chosen uniformly at random from the set of unprobed nodes.

In order to increase the likelihood that our random samples are informative, we choose our random nodes from those that were present in the initial network \hat{G}_0 . The rationale behind this choice is that as a consequence of probing nodes sequentially nodes that have been in the network for longer have more “complete” information, since they have had more opportunity to be connected to in the $t - 1$ probes before time t . That is, if a node $j \in \hat{G}_0$ has only a few neighbors after many queries have been made, it could be that j has very few neighbors, but it could also be that j connects to a neighborhood that the algorithm has yet to discover. Therefore, to explore the possibility of learning a better model using different information, NOL* algorithms select the random node from $\hat{V}_0 - P_t$ to probe until all nodes in \hat{V}_0 are exhausted, when NOL* chooses any unprobed node in \hat{G}_t at random.

Since NOL* algorithms learn from all or most of the previous samples at every time step, it is not strictly necessary to continue random exploration through the entire querying process. This is consistent with the literature on ϵ -greedy algorithms, which often systematically lower the rate of exploration over time (Kirkpatrick et al. 1983; Cheng et al. 2010). NOL* adds an optional exponential decay to the initial random jump rate ϵ_0 , such that at time t the jump rate is computed as $\epsilon_t = \epsilon_0 e^{-\frac{t}{b}}$. We have also experimented with data-driven methods for adaptive- ϵ -greedy, such as in Tokic (2010), but did not find them to be advantageous and leave further investigation of their utility in this space for future work.

Scalability & guarantees

In general, the computational complexity of a NOL* algorithm is the product of (1) the budget, (2) the feature computation complexity and (3) the learning complexity. In symbols, $O(b \times O(\text{features}) \times O(\text{learning}))$.

In this section, we briefly discuss the complexity of the learning steps of NOL and NOL-HTR. We also note that the features used in NOL* algorithms are user defined and range from trivial to compute (e.g. degree) to computationally expensive (e.g. node embeddings). Due to this, we omit a detailed analysis of the specific features we use in

our experiments (described in [Experiments](#) section), but note that since the computations are to happen online, and only nodes whose features may have changed should be updated, the complexity depends not necessarily on the total number of nodes or edges in the graph, but on the size of the neighborhood around the queried node for which the features might have changed.

The complexity of a NOL online regression update depends only on the number of features, since the most expensive operation required is multiplication of the feature vector by a constant factor. Due to this, even with a relatively high-dimensional feature vector, the complexity of the learning step is trivial. In this case, the scalability of the entire process will likely hinge on the scalability of the feature value updates.

The procedure to learn parameters for NOL-HTR is more complicated. The data is first partitioned into k subsamples, the covariance of each subsample is computed and optimal regression parameters are estimated for each, and finally the median of the k parameters is computed. The most expensive operation is the regression parameter estimation, which requires computing the Moore-Penrose inverse of the regressor matrix. This can be computed via Singular Value Decomposition of the matrix, which has complexity $O(nd^2)$, where n is the number of nodes in the computation and d is the number of features.

Hsu and Sabato (2016) also derive a bound on the loss and show that our learned parameters θ are within an ϵ of the true parameters if we use Algorithm 2 with the number of samples $m \geq O\left(d \log\left(\frac{1}{\delta}\right)\right)$. In our case, m is the number of rows in X . The derived bound states that with probability $(1 - \delta)$ the empirical loss is bounded by

$$L(\theta) \leq \left(1 + O\left(\frac{d \log\left(\frac{1}{\delta}\right)}{m}\right)\right) L_*$$

where L_* is the true loss with the optimal parameters and $\delta = \frac{1}{e^k}$. In practice, we avoid wasted computation by limiting the total number of samples m to 2000, which is always larger than is necessary for the guarantees given the values of k in our experiments (which determines delta in our formulation of the algorithm).

Experiments

In this section, we show the utility of applying NOL* algorithms for network completion in a variety of datasets, including both synthetic and real world networks. We first describe the data we use to test our method, then explain our experimental methodology and present results.

Data

We test NOL* algorithms on a range of synthetic graphs, as well as five real world datasets.

Synthetic models

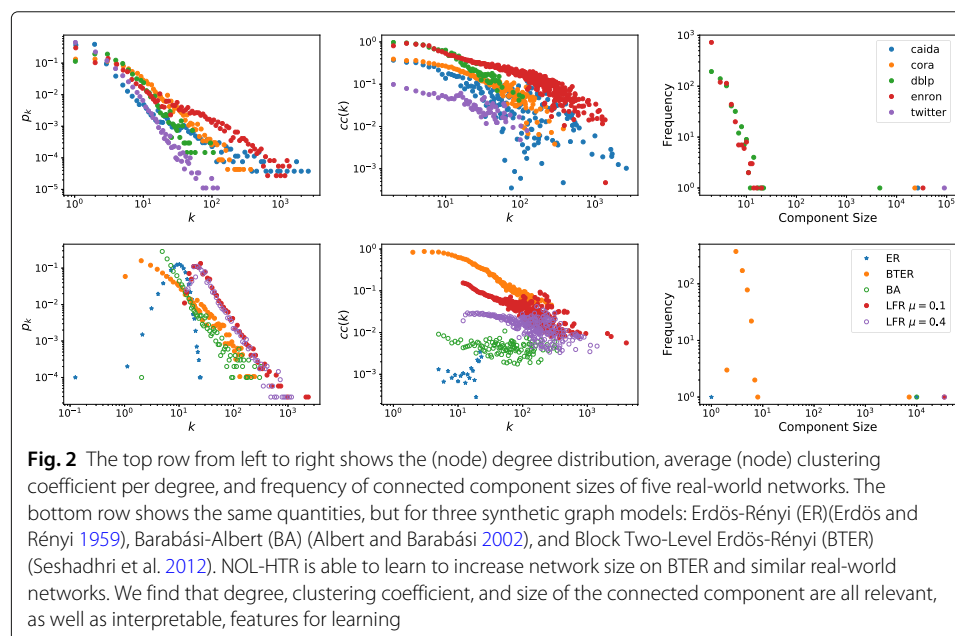
Synthetic network models are useful for generating test data that exhibits interesting properties often found in real world networks. We are interested in testing the performance of our algorithm on datasets that allow us to vary a few macroscopic properties of networks: the distribution of degrees, the extent of local clustering, the modularity of the global structure.

The degree distribution is an important factor in understanding when learning is possible or helpful. There are two simplified extremes that can characterize degree distributions in complex networks: *homogenous* and *heterogenous* (or heavy tailed) distributions. The important difference is that in a heterogeneous degree distribution some nodes accumulate far more connections than the majority of nodes in the network, resulting in these nodes becoming topological hubs. In the presence of hubs, the variance of the degree distribution can become large as the number of nodes in the network N grows, until it eventually diverges as $N \rightarrow \infty$. In contrast, a homogeneous degree distribution implies that hubs are not present, meaning the nodes are statistically equivalent with respect to degree. In the homogenous case, the degree of a randomly chosen node will be a random variable following a distribution with well-defined variance.

A second node characteristic that we conjecture is important for learning is the *clustering coefficient*. The clustering coefficient (more precisely, *local clustering coefficient*) is a measure of the extent to which a node's neighbors are connected to one another. Since the clustering coefficient is real-valued, it is often more intuitive to study the average clustering by degree, which is what we show in Fig. 2. Clustering is related to the local density of neighborhoods and can serve as a proxy for the amount a neighborhood has been explored.

Lastly, we are interested in studying the impact of the *modularity* of a network structure on algorithm performance, meaning the prevalence of within-community links relative to out-of-community links in a modularity-maximized partitioning of the nodes into communities. The extent of modularity in a network structure could be instructive in balancing exploration and exploitation, since learning to probe in a highly modular structure is more susceptible to settling on local minima by exploiting in a single community without discovering cross-community links.

To test the above conjectures on the limitations of learning in complex networks, we study five synthetic network models.



- 1 Erdős-Rényi (ER) (Erdős and Rényi 1959)⁴: In a network sampled from the ER model (specifically the ensemble G_{Np}), every possible (undirected) link between N nodes exists with probability p . Networks generated by the ER model have a homogeneous degree distribution (the exact distribution is Binomial, but it is often approximated by Poisson). Parameters: $N = 10000$, $p = 0.001$.
- 2 Barabási-Albert (BA) (Albert and Barabási 2002): The BA model generates networks through a growth and preferential attachment process where each node entering the network chooses a set of m neighbors to connect to with probability proportional to their relative degree. This process results in a heterogeneous degree distribution, which in the infinite limit follows a power law distribution with exponent 3. Parameters: $N = 10000$, $m = 5$, $m_0 = 5$. m_0 denotes the size of the initial connected network. m denotes the number of existing nodes to which a new node connects.
- 3 Block Two-level Erdős-Rényi (BTER) (Seshadhri et al. 2012): BTER is a flexible⁵ model that combines properties of the ER and BA model. It consists of two phases: (i) construct a set of disconnected *communities* made up of dense ER networks, with the size distribution of the communities following a heavy tailed distribution (i.e., a small number of large communities and many more small communities) and (ii) connect the communities to one another to achieve desired properties, such as a target value of global clustering coefficient. Parameters: $N = 10000$, target maximum clustering coefficient = 0.95, target global clustering coefficient = 0.15, target average degree $\langle k \rangle = 10$.
- 4 Lancichinetti-Fortunato-Radicchi (LFR) (Lancichinetti et al. 2008): In the LFR model, modular structure can be induced by varying the mixing parameter μ , which controls the extent to which nodes connect internally to a tight community (higher modularity) or loosely to the entire network (lower modularity), thus controlling the extent of modular community structure. Parameters: $N = 34546$, mixing parameter $\mu \in 0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7$, degree exponent $\gamma \in 2, 2.25, 2.5, 2.75, 3, 3.25, 3.5$, community size distribution exponent $\beta = 2$, average degree $\langle k \rangle = 12$, and maximum degree $k_{max} = 850$.
- 5 k -regular networks: In a k -regular network, every node is connected to k other nodes such that connections are random and every node has the same degree. Parameters: $N = 10000$, $k = 6$.

We note that we have left out analysis of the Watts-Strogatz (WS) model (Watts and Strogatz 1998), which is a model developed to study the small world property of random graphs. In WS, a rewiring parameter controls the trade-off between nodes clustering into triangles and the average path length between all nodes, a proxy for the small-world property. We do not study this model because (a) its degree distribution is homogeneous, following a Poisson distribution, and (b) the model does not result in networks with modular structure. Therefore, despite many uses in other contexts, the WS model is not well suited for studying the network completion problem in the present case.

⁴We omit almost all results on ER graphs from the paper because all methods perform indistinguishably on these graphs.

⁵There are many other random graph models that provide flexibility similar to the BTER that we could have used to generate networks for our experiments. We chose BTER out of convenience because it allows us to easily specify target values for average degree and clustering parameters directly.

Real-world networks

We evaluate NOL* on real world datasets whose characteristics are summarized in Table 1. We show the number of nodes (N), number of edges (E), number of triangles ($\#\Delta$), and the modularity (Q), computed by finding a modularity maximizing partition with the Leiden algorithm (Traag et al. 2019), a recently proposed improvement on the well-known Louvain algorithm (Blondel et al. 2008) for finding node partitions with high modularity. Degree distribution, average clustering by degree and component size distribution is shown for each network in Fig. 2.

Sampling methods

NOL* is agnostic to the method of sampling used to collect the initial sample graph, \hat{G}_0 . For the sake of continuity, all of the initial samples used in this paper were collected via *node sampling with induction*. In this technique, a set of nodes is chosen uniformly at random, then a subgraph is induced on the nodes (i.e. all of the links between them are included in the sample). Our samples are defined in terms of the proportion of the edges in the underlying network. To generate samples with the target proportion of edges, we choose a sample of nodes and induce a subgraph on them; if this subgraph has too many (or too few) edges, we repeat with a larger (or smaller) subset of nodes until we find an induced graph with an acceptable number of edges.

In the main text of this paper we present results on samples collected via node sampling with induction, but we have also verified many of the results using random walk with jump sampling (Ahmed et al. 2013), which can be found in [Appendix A.3](#).

Features

To accurately predict the number of unobserved neighbors of a partially observed node, NOL* algorithms require interpretable node features that are relevant across a variety of very different network structures. These features must be feasible to compute and update online for our algorithms to be scalable. We use the following features for each node i visible in the sample network:

- $\hat{d}(i)$: the normalized degree of node i in the sample network, which can be seen as a measure of the node's centrality in the sample. The inclusion of this feature assumes that the degree of a node in the sample is relevant to its total degree.
- $\hat{c}c(i)$: the clustering coefficient of node i in the sample network. This feature captures the local neighborhood density of a node, particularly the tendency of the node and its neighbors to form *triangles*.

Table 1 Basic Characterization of Real Networks

	Type	N	E	$\#\Delta$	Q
Caida	Internet Router	26.5k	53.4k	36k	0.68
Cora	Citation	23k	89k	78k	0.80
DBLP	Coauthorship	6.7k	17k	21k	0.89
Enron	Email Communication	36.7k	184k	72k	0.62
Twitter	Social interaction	90k	117k	9.4k	0.87

- *CompSize(i)*: the normalized size of the connected component *in the sample network* to which node *i* belongs, which can be used to facilitate exploration and exploitation based on where in the network the node is located.
- *pn(i)*: the fraction of node *i*'s neighbors which have already been probed. This feature indicates the extent to which the neighborhood of node *i* has been explored by the algorithm.
- *LostReward(i)*: the number of nodes that first connected to node *i* by being probed. Unlike *pn(i)*, *LostReward(i)* only counts nodes that were not connected to *i* before they were probed. This feature mitigates the ordering effects of probing nodes. If nodes *i* and *j* both have the same unobserved neighbors, for instance, probing *j* first would normally lower the total reward of node *i* when *i* would be probed. *LostReward(i)* gives credit to *i* upon its probing, since it could have brought in as many new nodes as *j*.

Baseline methods

We compare the performance of NOL* with 4 heuristic baseline methods and a multi-armed bandit method. Our heuristic baselines are as follows:

1. High degree: Query the node with maximum observed degree in every step. This has been shown to optimal in some heavy tailed networks (Avrachenkov et al. 2014).
2. High degree with jump: With probability ϵ , query a node chosen uniformly at random from all partially observed nodes. With probability $1 - \epsilon$, query the node with the maximum observed degree. We set $\epsilon = 0.3$.
3. Random degree: Query a node chosen uniformly at random from all partially observed nodes.
4. Low degree: Query the node with minimum observed degree in every step. This method is approximately⁶ optimal for *k*-regular networks where every node has the same degree, since the lowest degree node is always furthest from the uniform degree *k* (see Appendix A.1).

We also compare our approach with the nonparametric multi-armed bandit model proposed in Madhawa and Murata (2019), which we refer to as KNN-UCB (k-nearest-neighbors upper confidence bound). This method similarly relies on computing a vector of structural features for each node, including degree, average neighbor degree, median neighbor degree and average fraction of probed neighbors. Each unprobed node is considered an arm in a Multi-armed Bandit formulation, and the next arm to pull (node to probe) is chosen by computing

$$\operatorname{argmax}_i \hat{f}(x_i) + \alpha \sigma(x_i)$$

Here, $\hat{f}(x_i)$ is the expected reward of probing node *i* and $\sigma(x_i)$ is the average distance to other points in the neighborhood. The expected reward is calculated as a weighted k-NN regression on nodes within the *k-NN radius* of node *i*, defined as the *k* nodes whose feature vectors have Euclidean distance less than *r* from the feature vector of node *i*. The term $\alpha \sigma(x_i)$ facilitates exploration by allowing the possibility of nodes without maximum

⁶Cases where low degree is not optimal, even in networks with uniform degree, can be constructed. Specifically, if the neighbors of the node with lowest degree are already in the network, querying it will result in reward 0. If in the same case the neighbors of the node with 2nd lowest degree are almost all outside of the network, reward for querying that node will be larger.

expected reward to be probed, assuming $\alpha > 0$. In our experiments we fixed the value of $k = 20$ and $\alpha = 2$, following the experiments in the original paper.

Experimental setup

Across all networks, our experiments are run over 20 independent initial node samples of the network. In synthetic networks, the underlying network is a realization of the model using the parameters described in [Synthetic models](#) section.

Our experiments aim to (1) investigate how network properties impact the performance of NOL-HTR, (2) exhibit the performance tradeoffs between NOL-HTR and NOL, (3) show that NOL* algorithms outperform the baseline methods in settings where learning is possible and approximate the heuristic methods elsewhere, (4) analyze the prediction error of NOL-HTR, and (5) analyze the evolution of the feature weights learned by NOL-HTR over time.

Performance metrics

After each probe of the network, NOL* earns a *reward*, defined in this work as the number of previously unobserved nodes included in the network after a probe. Formally, the reward at time t is defined as $r_t = |V_{t+1}| - |V_t|$. We study the performance of each method by showing the *cumulative reward*, $\hat{c}_r(T)$, where T is a time step between $0, 1, 2, \dots, b$. Formally, $\hat{c}_r(T) = \sum_{t=1}^T r_t$. We also want to quantify the utility of decisions made by NOL*. For this purpose, we study the prediction error of the model. This quantifies the extent to which our prediction, $\mathcal{V}_{\theta_t}(i) = \theta_t \phi_t(i)$, differs from the true reward value r_t . Thus, we calculate $E(t) = \mathcal{V}_{\theta_t}(i) - r_t$.

NOL-HTR parameter search

We ran a two dimensional parameter search over values of k (1-16, 32, 64, 128, $\log_{10}(n)$, $\log_e(n)$, $\log_2(n)$) and ϵ (0, 0.1, 0.2, 0.3, 0.4, and exponential decay versions of each). We ran this search on all of the networks in [Fig. 2](#), as well as 56 LFR networks (Lancichinetti et al. 2008) spanning a wide variety of network structures. We varied two parameters in the LFR networks: μ , which controls modular structure by adjusting the probability of cross-community links; and γ , which controls the exponent of the degree distribution of the entire network (Lancichinetti et al. 2008).

Across all of these networks, we did not observe a general trend in which some parameter settings performed best consistently across experiments. There was not a single best choice or regime of parameters through the experiments for any of the networks we searched on individually, nor was there a standout choice of parameter across the networks.

However, we have found some evidence to suggest that performance on more modular structures is improved by more randomization, meaning a non-zero value of ϵ . We ran a linear regression using ϵ as the target variable and global clustering, modularity, and degree exponent as covariates. Results are presented in [Table 2](#). All three covariates positively effected ϵ , with modularity having the strongest effect. This result is consistent with a network scientific understanding of the querying process: when querying a modular network structure, the likelihood of finding a local minima within one highly connected and clustered module is high, meaning that adding more randomness to the algorithm will increase the likelihood that the algorithm sees examples that allow it to avoid settling on such a minima.

Table 2 Coefficients of a linear regression using the output of our NOL-HTR parameter search

	Coefficient	Standard Error	<i>p</i> -value
Intercept	-0.2163	0.116	0.066
Global Clustering	0.1361	0.478	0.777
<i>Q</i>	0.3836	0.135	0.006
γ	0.0458	0.034	0.186

The regressor matrix was made up of the following statistics for each network in the search: global clustering, modularity (*Q*), and estimated degree exponent (γ). The response variable was the best performing ϵ . Results suggest that higher modularity predicts higher ϵ , suggesting that networks with modular structure benefit most from exploration

We report results using a set of parameters that performed reasonably well across networks. These parameters are $\epsilon = 0.3$, decaying as $\epsilon_t = 0.3e^{-\frac{t}{b}}$ and $k = \ln(t)$, where *t* represents the number of nodes probed thus far in the experiment.

Results

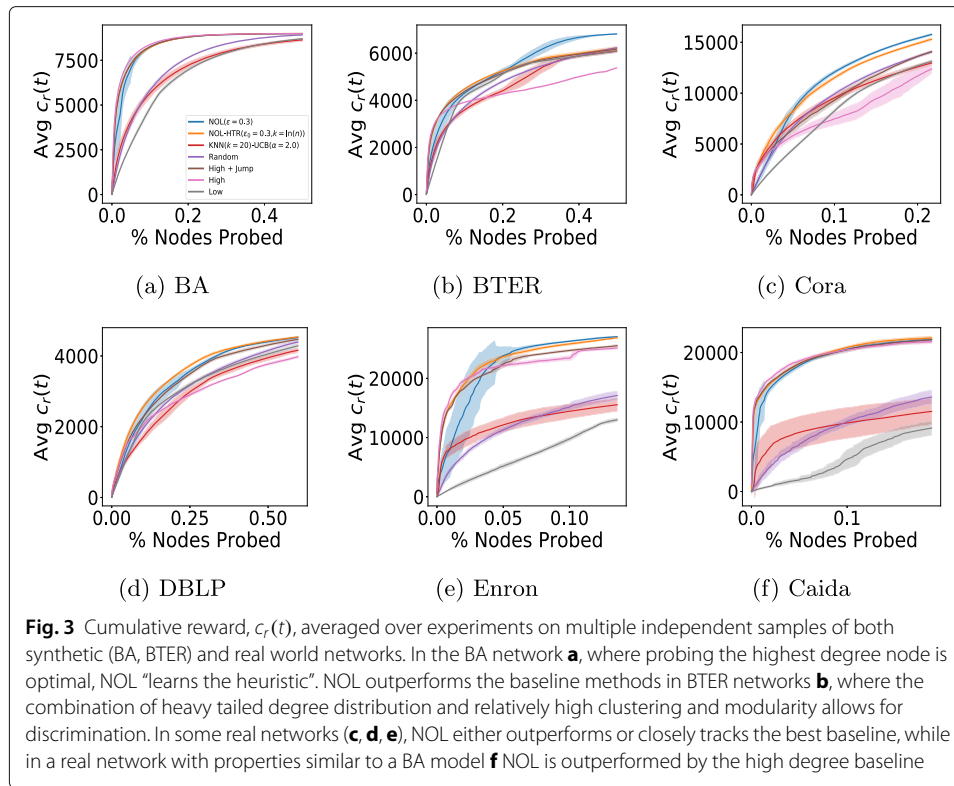
In this section, we compare the performance of the heuristic baseline methods, the KNN-UCB baseline approach, NOL, and NOL-HTR. Broadly, we find that NOL-HTR and NOL perform similarly in terms of average cumulative reward, but that the performance of NOL-HTR is more consistent, indicated by standard deviations that are tighter around the mean across experiments.

Cumulative reward We report the average cumulative reward $\hat{c}_r(T)$ over a budget of thousands of probes on 6 networks in Fig. 3. The average and standard deviation of $\hat{c}_r(T)$ are computed over experiments on 20 independent samples. We omit results on ER networks, noting that because all nodes are statistically equivalent in terms of structural properties, every probing method performs equivalently and neither learning or heuristics provide any significant advantage (see Appendix A.1).

In BA networks (Fig. 3a), NOL-HTR performs on par with the High Degree baseline, which is known to be near optimal in networks with heavy tailed degree distributions (Avrachenkov et al. 2014). Further, NOL-HTR outperforms NOL by achieving both higher average reward and smaller standard deviation. NOL-HTR also consistently outperforms the baseline methods in networks generated by the BTER model (Fig. 3b). Although NOL-HTR outperforms NOL in both reward and variance towards the beginning of the experiment, after around 25% of the nodes have been probed NOL begins to outperform NOL-HTR. See Comparing NOL and NOL-HTR section for a discussion of some potential explanations for this observation. NOL-HTR performed as well as NOL and the best heuristics in every real world network we experimented on. Comparing directly with NOL, NOL-HTR is able to achieve similar or better performance, always with smaller standard deviation, on every network, regardless of the underlying distributions (compare networks in Fig. 2). We note that across experiments, the KNN-UCB method was unable to match the performance of our model and often underperformed the baseline methods.⁷

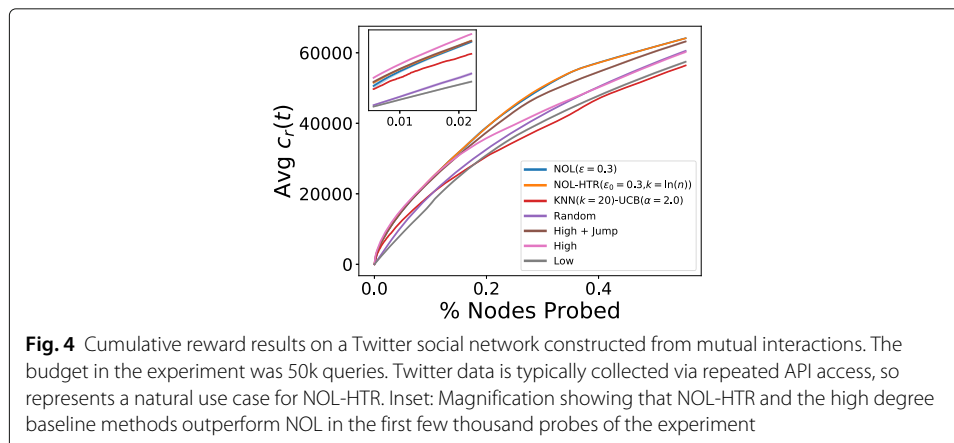
Use case: twitter social network We present results on a social interaction network sampled from Twitter in Fig. 4. In this setting, querying a node corresponds to calling the

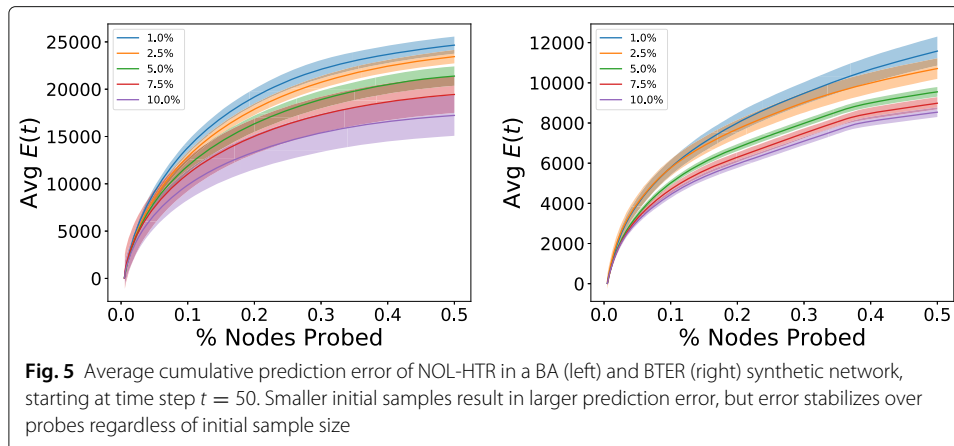
⁷We further note that the sample collection techniques used in the experiments in the KNN-UCB paper (Madhawa and Murata 2019) were defined differently from those we employ here and that all of our experiments ran over larger probing budgets.



Twitter API to obtain more data about a specific account. This is an example of a natural use case for NOL-HTR, both because the network must be expanded by repeatedly accessing an API and because there is evidence that the distribution of social interactions is heavy-tailed (see e.g. Kossinets and Watts 2006 and Morales et al. 2012).

This Twitter Social Network data set consists of Twitter users connected to one another with an undirected edge if they mutually retweeted or mentioned each other throughout the course of a single day in 2009. We focus our experiment on the largest connected component of this interaction network. While this is only a small fraction of the Twitter network, the data was collected via the Twitter Firehose, and is therefore complete for that time span, which allows us to accurately simulate the process of growing a network by querying Twitter users. Some other characteristics of the dataset are outlined in Table 1 and the distributions of degree and clustering are shown in Fig. 2. Notably in our context,





though the maximum degree in this network is not very large (≈ 100), choosing a random node with degree near the maximum is much less likely than choosing a node with low degree.

As shown in Fig. 4, NOL outperforms the heuristic baseline methods and NOL-HTR.⁸ However, the inset plot shows that, similar to the experiments on other networks, NOL-HTR outperforms NOL early in the experiment before eventually being overtaken. We discuss our conjectures about why we observe this trade-off behavior in [Comparing NOL and NOL-HTR](#) section.

Prediction error We analyze the ability of NOL* to learn by showing the measure of *prediction error*, $E(t)$, defined in [Performance metrics](#) section, over time and across different initial sample sizes.

Cumulative prediction error, $E(T)$, as a function of time is shown in Fig. 5. The error is averaged over 10 independent samples of the network for initial sample sizes of 1%, 2.5%, 5%, 7.5% and 10% of the complete network (as percent of edges in the network).

We show results in both the BA and BTER models. The predictions of NOL-HTR are noisy in the beginning of the experiment and the resulting outliers skew the analysis of cumulative error, therefore we present the error starting from $t = 50$. In both cases the prediction error is relatively stable over time, indicated by the slow growth of the curves. We also observe that the average prediction error is lowest when the initial sample is largest, and highest when the initial sample is smallest; This is intuitive, since when the initial sample is large, the model is learning its predictions from more accurate information, whereas when the initial sample is small, the training information is very noisy.

Feature weight analysis We qualitatively analyze the feature weights learned by NOL-HTR over time (see Fig. 8 for visualization). Since the weights were computed on random subsamplings of the observed data at every time step, they fluctuated considerably between probes. Still, we found that across most networks (all but Caida and Enron), the degree of a node was weighted positively and large in magnitude, implying that it is the most salient feature to predict reward. This is intuitive, since we expect the target variable

⁸Our attempt to run the KNN-UCB baseline on the Twitter network did not finish in a reasonable amount of time, therefore we omit it.

to correlate highly with the sample degree. There were no features with consistently highly negative weights, which would imply a negative predictor of reward. Instead, the other features typically fluctuated around 0, meaning they were not consistently predictive of either high or low rewards, but were non-zero so did contribute to the prediction.

The Caida and Enron networks were exceptions to the above general trends, though NOL-HTR exhibits similar performance on these networks in terms of cumulative reward. In these experiments, the weight of the degree feature was centered around 0 along with the other features, but with very large fluctuations in both positive and negative directions. This suggests that the importance assigned to degree depended strongly on the particular subsampling of the data in an individual time step. It may also have implications for the impact of degree correlations (i.e. average neighbor degree), since the fluctuations are consistent with both low and high degree nodes predicting high reward.

Comparing NOL and NOL-HTR

The above analysis illustrates some differences and tradeoffs between NOL and NOL-HTR: **(1)** NOL-HTR achieves similar performance to NOL and is more consistent. This is evidenced by the tighter standard deviations around the average cumulative reward. **(2)** NOL-HTR outperforms NOL in the beginning of every experiment. This is consistent with our expectation based on how each of the algorithms work: NOL-HTR is able to leverage outlier, high reward queries early on because it computes maximum likelihood parameters at every step, while NOL updates its parameters with online gradient descent using a fixed learning rate, and is therefore not able to adapt as effectively to high reward nodes. **(3)** As the number of probes increases, NOL often begins to outperform NOL-HTR. There are a few possibilities for why this is happening. First, NOL is learning through gradient descent, so adjusting the learning rate of the algorithm could impact the amount of examples in the tail it takes to reach highly predictive parameters, explaining the lag. Second, as the number of samples grows, the NOL-HTR parameter setting of $k = \ln(n)$ grows very slowly. This means that more data is being subsampled into the same number of bins, thus each subsample may become more noisy and the outliers become less distinguishable. This could be alleviated by increasing the value of k more quickly as the number of samples grow, or by capping the size of the subsamples, or by choosing the sample size so that data from the tail is more likely to have an impact on the parameters.

Limitations of learning

We are interested in understanding the limitations of our approach in this setting. We experiment across a wide variety of network structures using the LFR community benchmark (Lancichinetti et al. 2008). We generated 56 LFR networks spanning two parameters: μ , which controls modular structure by adjusting the probability of cross-community links; and γ , which controls the exponent of the degree distribution of the entire network. Then, we ran the NOL-HTR parameter search (see [NOL-HTR parameter search](#) section) on each network, as well as random and high degree baseline methods. For each network, we found the set of NOL-HTR parameters that resulted in the maximum average cumulative reward and computed the percent gain

in performance using either of the baselines as

$$c_r^\Delta = \frac{c_r^{HTR} - c_r^{base}}{c_r^{base}} \times 100$$

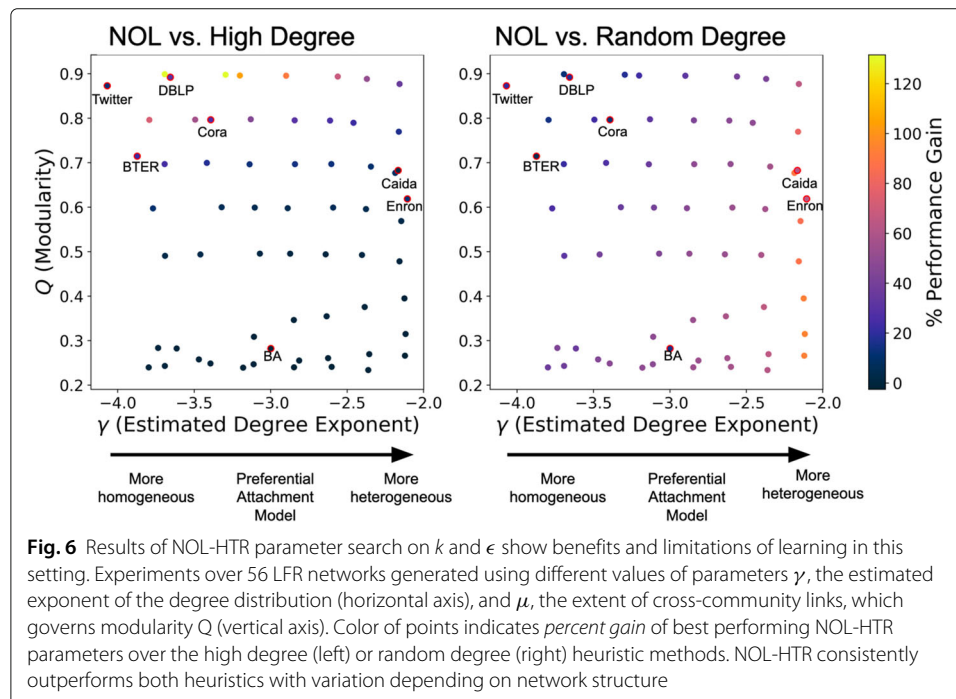
Results are shown in Fig. 6. Compared with high degree probing (left plot), NOL-HTR is always able to gain in performance in higher modularity networks ($Q > 0.6$), with performance gains spanning from 15-30% up to about 130%. When modularity is lower, including in the BA model realization, NOL-HTR approximates high degree across degree exponents, with the maximum performance loss of less than 3% (-2.54%) compared to the heuristic.

Comparing with random degree probing (right plot), NOL-HTR is always able to gain in performance, with minimum performance gain of 5% and maximum of 93%. The maximum gains are in networks with the most heterogeneous degree distributions. This corresponds to the most star-like network structures, meaning most randomly chosen nodes will have very low degree.

The performance compared to random degree probing in the real networks seem to fit with performance on LFR networks with similar parameters, allowing for some noise in either direction.

Comparing with high degree on BTER, Caida, Cora and Enron, performance appears to be about on par with similar LFR networks, again allowing for some noise. Performance gain by using NOL-HTR rather than high degree in the DBLP network is substantial, but smaller than an LFR network with similar properties.

The degree exponent estimate for our Twitter sample is an outlier in this dataset. However, there is also substantial performance gain, though it might be smaller than what we would expect given an LFR network with similar properties.



Taken together, these results show that the ability of our model to learn in this setting is tied to the structural properties of the network. In some cases, a low computational cost heuristic such as high or random degree can perform well, but our experiments show that there is almost never a disadvantage to using a NOL* algorithm, and that the advantages can be substantial, up to 130% gains in performance.

Conclusion and future work

We proposed and evaluated algorithms to address the problem of reducing the incompleteness of a partially observed network via successive queries as an online learning problem. We presented two algorithms in the NOL* family and highlighted NOL-HTR for learning heavy-tailed reward distributions. We showed that NOL-HTR is able to consistently outperform other methods, especially early-on in the process of querying nodes when the extreme values are yet to be discovered. We also showed that macroscopic properties of the underlying network structure, specifically the degree distribution and extent of modularity, are important factors in understanding when learning will be relatively easy, difficult, or nearly impossible. Alongside experiments on multiple synthetic and real world networks, we presented experiments on a Twitter interaction network, a realistic use case for a network growth algorithm such as NOL-HTR.

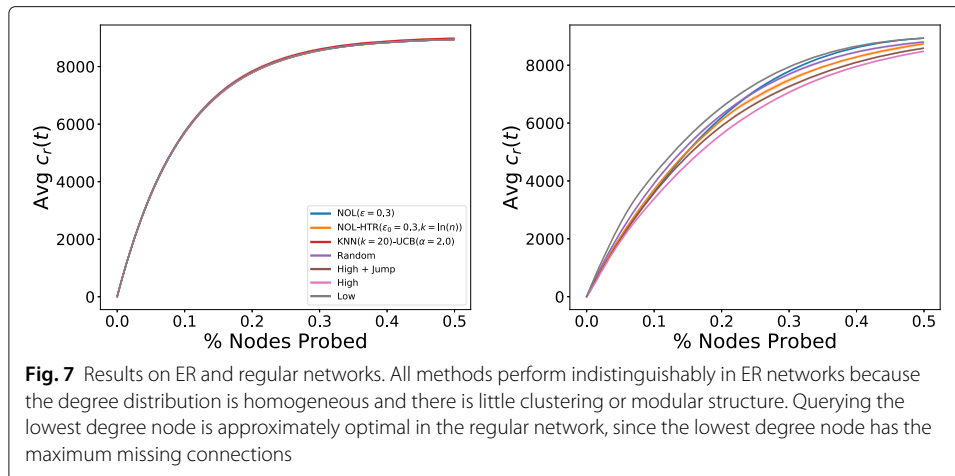
The problem of online network discovery remains fruitful for future work. The case of noisy observations from the query model has yet to be fully addressed. For example, a query may return only a sample of the neighbors, or a list of potential neighbors that may include false positives. The specific models we presented here are not sensitive to this type of noise, but could be addressed in extensions to NOL*. Similarly, an adversarial version of the problem can be formulated, where an adversary is intentionally poisoning the queries (via either the query that is sent or the data that is returned) and the model must include the adversary in order to make appropriate adjustments to decision making. Such noisy network discovery tasks could be formulated as Partially Observed Markov Decision Processes (POMDPs), which differ from our MDP formulation by the fact that the agent is uncertain about its current state (for example, we could model some error on the feature vectors).

Appendix A:

A.1 Results on ER and regular networks

In Fig. 7, we present cumulative reward results on an ER network (left) and a k -regular network (right). The result on the ER network shows that every querying strategy performs indistinguishably from the others. This is due to the fact that the degree distribution of the network is homogeneous, meaning that the expected number of neighbors of each node is well defined and the same for all nodes. Since the expected number of neighbors is the same for every node, but the exact number of neighbors for a given node is random, no querying strategy is able to outperform any other.

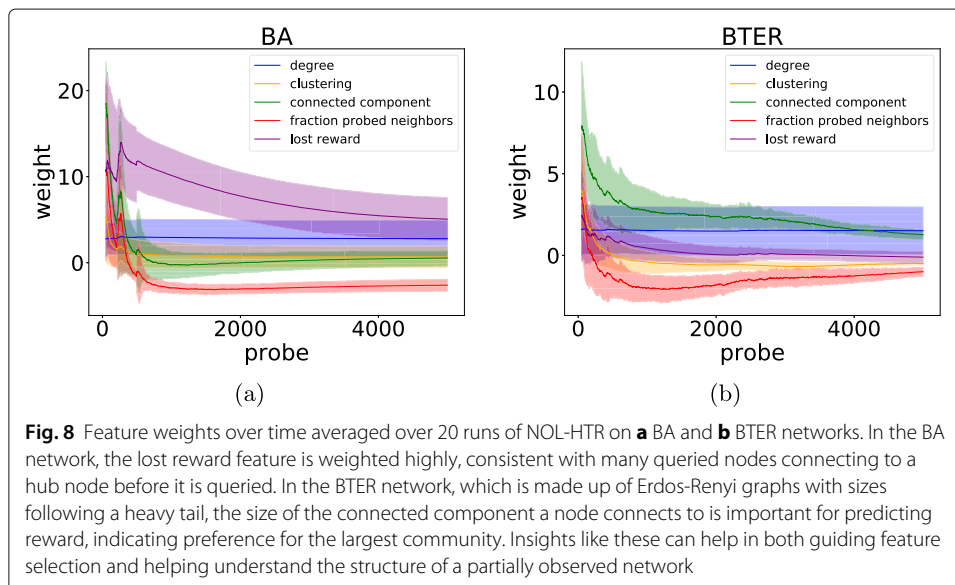
In the k -regular network, every node has exactly the same degree, but the particular neighbors are randomly chosen. Since every node has the same degree, choosing the lowest degree node is approximately optimal when maximizing the number of newly observed connections. For the same reason, choosing the highest degree node is usually far from optimal, since the maximum degree node in the sample will have degree nearly k , thus the minimum reward. This explains why high degree is the worst performer.



A.2 Feature weight analysis

In Fig. 8, we present analysis of feature weights from a NOL-HTR experiment, showing how the feature weights change as querying continues, averaged over 20 trials. The purpose is to show that it is possible to analyze the learned feature weights to help understand why NOL* algorithms perform the way they do. We choose NOL-HTR here, but in principle any parameterized model could be temporally analyzed in a similar fashion.

In the BA example (Fig. 8a), the LostReward feature, which takes order effects of querying into account (see Experiments section), appears as the most positively weighted feature. This makes some intuitive sense: hubs accumulate larger values of lost reward over time, since many nodes brought in by other queries are connected after they eventually get queried. This means the hubs are “missing out” on reward and since they are in fact the best nodes to probe, this is reflected in the weights of the features. This implies that in a BA network, we expect that the degree and lost reward features will be correlated. This correlation provides a potential explanation for why degree is, somewhat counterintuitively, one of the least important features when querying a BA network: the degree of



hub nodes is being accounted for in the lost reward feature, so degree itself does not have a strong impact.

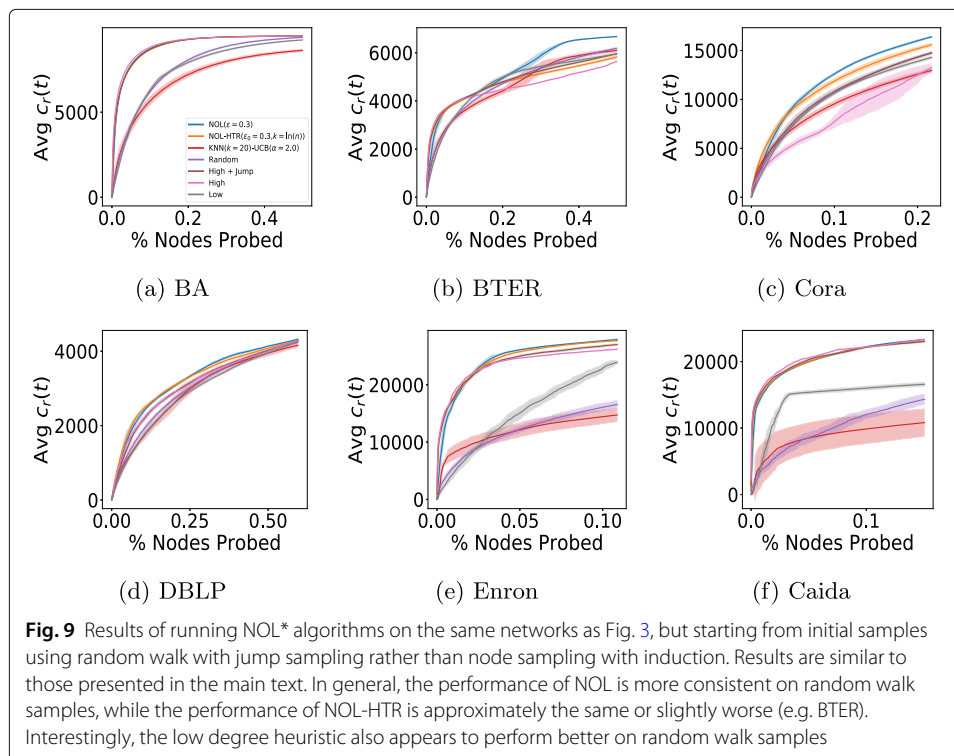
In the BTER example (Fig. 8b), the normalized size of the connected component that a node is in is the most highly weighted feature. As a reminder, BTER networks are generated by constructing many ER networks (“communities”) of different sizes, where the size of the networks follows a power law distribution, then connecting them to one another. This means there are a small number of very large and relatively well connected “communities”, which are the best place to query to bring in new nodes.

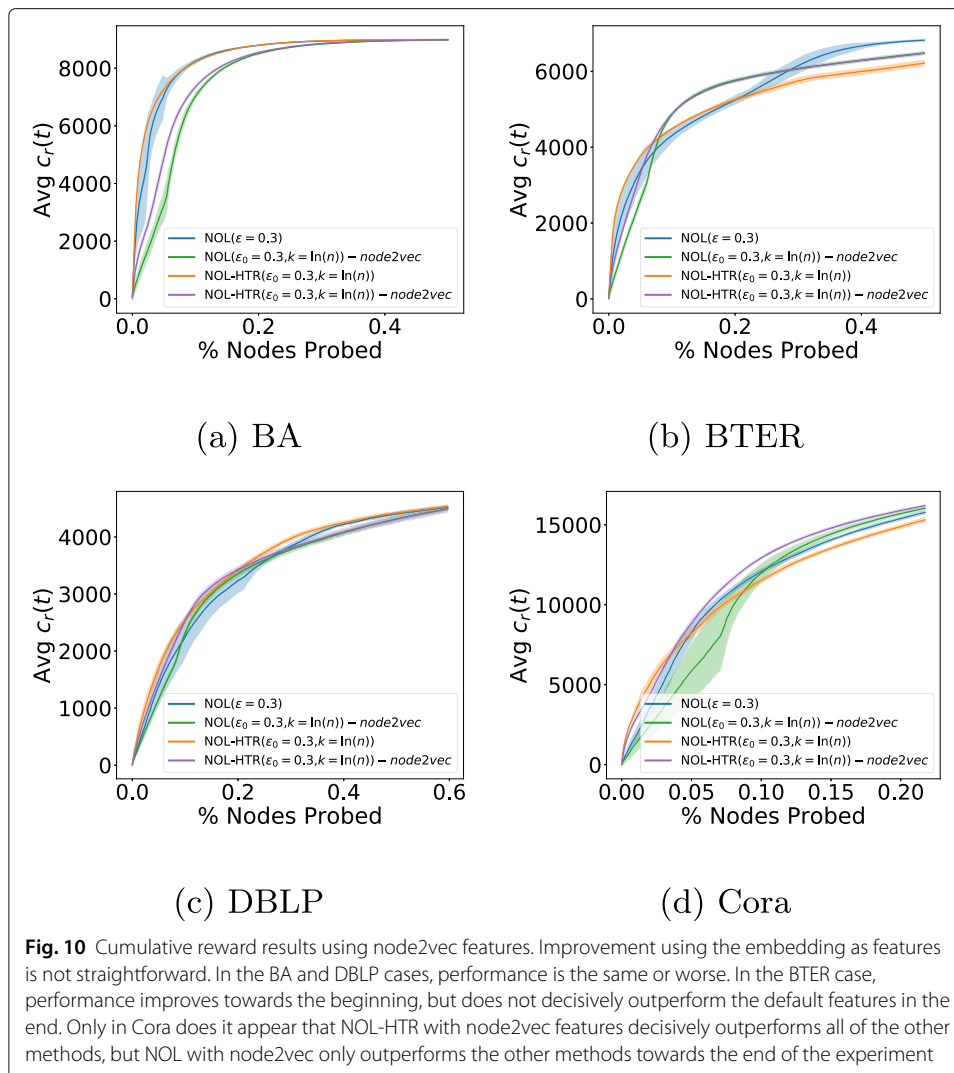
A.3 Random walk sampling

In Fig. 9 we show results of NOL* algorithms starting from random walk samples. The sampling works as follows: a node is chosen uniformly at random, then a random walk from that node proceeds until the desired proportion of edges is discovered, jumping randomly 15% of the time. The resulting performance is similar to what we saw in the node sampling with induction samples: NOL* algorithms are able to learn to match or outperform the heuristic methods in almost every case. However, the performance of NOL appears to be more consistent on random walk samples, as evidenced by the fact that the standard deviation around the mean performance tends to be tighter. The performance of NOL-HTR is approximately the same or even slightly worse (e.g. BTER) on the random walk samples. The low degree heuristic also appears to perform better on random walk samples.

A.4 Alternative features

In this section, we show some experiments using node2vec (Grover and Leskovec 2016) as features, rather than our hand selected features. Since node2vec is significantly more





computationally expensive, we expect a tradeoff between computation time and performance increases due to more expressive features. However, the results shown in Fig. 10 present a mixed picture: In some cases, using the embedding features makes no improvement or can even degrade performance. Only in one case, the Cora network, do the node2vec features truly outperform the others, and even then only with one of the learning algorithms (NOL-HTR). These results are inconclusive on the benefit of using embeddings as features and show that more experimentation and research needs to be done to understand the tradeoff between complexity of features and improvement in performance.

Abbreviations

NOL: Network Online Learning; HTR: Heavy Tail Regression; API: Application Programming Interface; MDP: Markov Decision Process; POMDP: Partially-observed Markov Decision Process; WWW: World Wide Web; MEUD: Maximum Expected Uncovered Degree; ER: Erdős-Rényi BTER: Block Two-level Erdős-Rényi BA: Barabási-Albert; LFR: Lancichinetti-Fortunato-Radicchi; WS: Watts-Strogatz; KNN-UCB: k-nearest-neighbors upper confidence bound; # Δ : Number of Triangles

Acknowledgements

The authors acknowledge Brennan Klein for his help designing Fig. 1.

Authors' contributions

This manuscript would not exist without TL; he was the lead contributor in all aspects of the manuscript. TS was involved in generating the synthetic data sets, preparing the Twitter data, and generating the feature importance results. SB was involved in the design of algorithms presented in the manuscript. TER supervised the work and contributed to methodology and experimental design. TL was the lead developer of the code and conducted all the experiments. TL was the lead contributor in writing the manuscript. All author(s) read and approved the final manuscript.

Funding

This research was sponsored in part by (1) the National Science Foundation grants NSF CNS-1314603 and NSF IIS-1741197, (2) the Combat Capabilities Development Command Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-13-2-0045 (ARL Cyber Security CRA), and (3) the Under Secretary of Defense for Research and Engineering under Air Force Contract No. FA8702-15-D-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Combat Capabilities Development Command Army Research Laboratory or the Under Secretary of Defense for Research and Engineering or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes not withstanding any copyright notation here on.

Availability of data and materials

Code and data to reproduce the results in this paper are available at <https://github.com/tlarock/nol/>. Datasets used in our experiments were downloaded from SNAP (Leskovec and Krevl 2014).

Competing interests

The authors declare that they have no competing interests.

Author details

¹Network Science Institute, Northeastern University, 360 Huntington Ave, Boston MA, 02115-5000 USA. ²Khoury College of Computer Sciences, Northeastern University, 360 Huntington Ave, Boston MA, 02115-5000, USA. ³Department of Computer Science and Engineering, Indian Institute of Technology Palakkad, Palakkad Kerala, 678557, India.

Received: 24 January 2020 Accepted: 29 July 2020

Published online: 09 September 2020

References

- Ahmed NK, Neville J, Kompella RR (2013) Network sampling: from static to streaming graphs. *TKDD* 8(2):7–1756
- Albert R, Barabási AL (2002) Statistical mechanics of complex networks. *Rev Mod Phys* 74(1):47–97
- Alves LGA, Aleta A, Rodrigues FA, Moreno Y, Nunes Amaral LA (2020) Centrality anomalies in complex networks as a result of model over-simplification. *New J Phys* 22(1):013043
- Avrachenkov K, Basu P, Neglia G, Ribeiro BF, Towsley DF (2014) Pay few, influence most: online myopic network covering. In: *INFOCOM Workshops*. IEEE, Toronto. pp 813–818
- Blondel VD, Guillaume J-L, Lambiotte R, Lefebvre E (2008) Fast unfolding of communities in large networks. *J Stat Mech Theory Exp* 2008(10):P10008
- Breza E, Chandrasekhar AG, McCormick TH, Pan M (2017) Using aggregated relational data to feasibly identify network structure without network data. *NBER Working Paper* (23491)
- Chen S, Mira A, Onnela J-P (2019) Flexible model selection for mechanistic network models. *J Complex Net* 8(2). <https://doi.org/10.1093/comnet/cnz024>
- Cheng R, Lo E, Yang XS, Luk M, Li X, Xie X (2010) Explore or exploit? effective strategies for disambiguating large databases. *PVLDB* 3(1):815–825
- Cho J, Garcia-Molina H, Page L (1998) Efficient crawling through UR ordering. *Comput Netw* 30(1-7):161–172
- Erdős P, Rényi A (1959) On random graphs I. *Publ Math* 6:290–297
- Ghosh S, Zafar MB, Bhattacharya P, Sharma NK, Ganguly N, Gummadi PK (2013) On sampling the wisdom of crowds: random vs. expert sampling of the twitter stream. In: *CIKM'13: 22nd ACM International Conference on Information and Knowledge Management San Francisco California USA October*. Association for Computing Machinery, New York. pp 1739–1744
- Gile KJ (2011) Improved inference for respondent-driven sampling data with application to HIV prevalence estimation. *JASA* 106(493):135–146
- González-Bailón S, Wang N, Rivero A, Borge-Holthoefer J, Moreno Y (2014) Assessing the bias in samples of large online networks. *Soc Networks* 38:16–27
- Grover A, Leskovec J (2016) node2vec: Scalable feature learning for networks. In: *KDD'16: 22nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining San Francisco California USA August*. Association for Computing Machinery, New York. pp 855–864
- Hanneke S, Xing EP (2019) Network completion and survey sampling. In: *AISTATS 2019: The 22nd International Conference on Artificial Intelligence and Statistics*. Proceedings of Machine Learning Research. pp 209–215
- Hsu D, Sabato S (2016) Loss minimization and parameter estimation with heavy tails. *JMLR* 17:1–40
- Pfeiffer III JJP, Neville J, Bennett PN (2014) Active exploration in networks: using probabilistic relationships for learning and inference. In: *CIKM '14: 2014 ACM Conference on Information and Knowledge Management Shanghai China November*. Association for Computing Machinery, New York. pp 639–648
- Kim M, Leskovec J (2011) The network completion problem: inferring missing nodes and edges in networks. In: *SDM*. pp 47–58
- Kirkpatrick S, Gelatt CD, Vecchi MP (1983) Optimization by simulated annealing. *Science* 220(4598):671–680
- Kossinets G, Watts DJ (2006) Empirical analysis of an evolving social network. *Science* 311(5757):88–90

- Lancichinetti A, Fortunato S, Radicchi F (2008) Benchmark graphs for testing community detection algorithms. *Phys Rev E* 78(4):46110. <https://doi.org/10.1103/PhysRevE.78.046110>, 0805.4770
- LaRock T, Sakharov T, Bhadra S, Eliassi-Rad T (2018) Reducing network incompleteness through online learning: A feasibility study. In: MLG '18. http://www.mlgworkshop.org/2018/papers/MLG2018_paper_40.pdf
- Leskovec J, Chakrabarti D, Kleinberg JM, Faloutsos C, Ghahramani Z (2010) Kronecker graphs: an approach to modeling networks. *JMLR* 11:985–1042
- Leskovec J, Krevl A (2014) SNAP Datasets: stanford large network dataset collection. <http://snap.stanford.edu/data>
- Madhawa K, Murata T (2019) A multi-armed bandit approach for exploring partially observed networks. *Appl Netw Sci* 4(1):26. <https://doi.org/10.1007/s41109-019-0145-0>
- Morales AJ, Losada JC, Benito RM (2012) Users structure and behavior on an online social network during a political protest. *Physica A* 391(21):5244–5253
- Murai F, Rennó D, Ribeiro B, Pappa GL, Towsley D, Gile K (2018) Selective harvesting over networks. *Data Min Knowl Discov* 32(1):187–217
- Peixoto TP (2018) Reconstructing networks with unknown and heterogeneous errors. *Phys Rev X* 8(4):041011
- Sampson J, Morstatter F, Maciejewski R, Liu H (2015) Surpassing the limit: keyword clustering to improve Twitter sample coverage. In: HT '15: 26th ACM Conference on Hypertext and Social Media Guzelyurt Northern Cyprus September. Association for Computing Machinery, New York. pp 237–245
- Sanz J, Cozzo E, Borge-Holthoefer J, Moreno Y (2012) Topological effects of data incompleteness of gene regulatory networks. *BMC Syst Biol* 6(1):110
- Seshadhri C, Kolda TG, Pinar A (2012) Community structure and scale-free collections of Erdős-Rényi graphs. *Phys Rev E* 85(5):056109
- Soundarajan S, Eliassi-Rad T, Gallagher B, Pinar A (2015) Maxoutprobe: a algorithm for increasing the size of partially observed networks. *CoRR abs/1511.06463*
- Soundarajan S, Eliassi-Rad T, Gallagher B, Pinar A (2016) Maxreach: reducing network incompleteness through node probes. In: ASONAM. IEEE, San Francisco. pp 152–157
- Soundarajan S, Eliassi-Rad T, Gallagher B, Pinar A (2017) ϵ -WGX: adaptive edge probing for enhancing incomplete networks. In: Proceedings of the 2017 ACM on Web Science Conference, WebSci 2017. Association for Computing Machinery, New York. pp 161–170
- Strehl AL, Littman ML (2007) Online linear regression and its application to model-based reinforcement learning. In: Advances in Neural Information Processing Systems 20 (NIPS 2007). Neural Information Processing Systems, San Diego. pp 1417–1424
- Sutton R, Barto A (2018) Reinforcement Learning: An Introduction. 2nd edn. MIT Press, Cambridge, MA
- Tokic M (2010) Adaptive epsilon-greedy exploration in reinforcement learning based on value difference. In: Kl. Springer, Karlsruhe. pp 203–210
- Traag VA, Waltman L, van Eck NJ (2019) From Louvain to Leiden: Guaranteeing well-connected communities. *Sci Rep* 9(1):5233
- Vázquez A, Pastor-Satorras R, Vespignani A (2002) Internet topology at the router and autonomous system level. *CoRR cond-mat/0206084*
- Wang YJ, Wong GY (1987) Stochastic blockmodels for directed graphs. *J Am Stat Assoc* 82(397):8–19
- Watts DJ, Strogatz SH (1998) Collective dynamics of 'small-world' networks. *Nature* 393(6684):440
- Wejnert C, Heckathorn DD (2008) Web-based network sampling: efficiency and efficacy of respondent-driven sampling for online research. *Social Methods Res* 37(1):105–134

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► [springeropen.com](https://www.springeropen.com)
