Interactive Anomaly Detection on Attributed Networks

Kaize Ding
Arizona State University
kding9@asu.edu

Jundong Li
Arizona State University
jundong@asu.edu

Huan Liu
Arizona State University
huan.liu@asu.edu

ABSTRACT

Performing anomaly detection on attributed networks concerns with finding nodes whose patterns or behaviors deviate significantly from the majority of reference nodes. Its success can be easily found in many real-world applications such as network intrusion detection, opinion spam detection and system fault diagnosis, to name a few. Despite their empirical success, a vast majority of existing efforts are overwhelmingly performed in an unsupervised scenario due to the expensive labeling costs of ground truth anomalies. In fact, in many scenarios, a small amount of prior human knowledge of the data is often effortless to obtain, and getting it involved in the learning process has shown to be effective in advancing many important learning tasks. Additionally, since new types of anomalies may constantly arise over time especially in an adversarial environment, the interests of human expert could also change accordingly regarding to the detected anomaly types. It brings further challenges to conventional anomaly detection algorithms as they are often applied in a batch setting and are incapable to interact with the environment. To tackle the above issues, in this paper, we investigate the problem of anomaly detection on attributed networks in an interactive setting by allowing the system to proactively communicate with the human expert in making a limited number of queries about ground truth anomalies. Our objective is to maximize the true anomalies presented to the human expert after a given budget is used up. Along with this line, we formulate the problem through the principled multi-armed bandit framework and develop a novel collaborative contextual bandit algorithm, named GraphUCB. In particular, our developed algorithm: (1) explicitly models the nodal attributes and node dependencies seamlessly in a joint framework; and (2) handles the exploration-exploitation dilemma when querying anomalies of different types. Extensive experiments on real-world datasets show the improvement of the proposed algorithm over the state-of-the-art algorithms.

ACM Reference Format:

Kaize Ding, Jundong Li, and Huan Liu. 2019. Interactive Anomaly Detection on Attributed Networks. In *The Twelfth ACM International Conference on Web Search and Data Mining (WSDM '19), February 11–15, 2019, Melbourne, VIC, Australia*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3289600.3290964>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WSDM '19, February 11–15, 2019, Melbourne, VIC, Australia

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-5940-5/19/10...\$15.00

<https://doi.org/10.1145/3289600.3290964>

1 INTRODUCTION

Anomaly detection is a canonical research task in the data mining and machine learning community. It aims to identify noteworthy items, events or entities that do not conform to the expected patterns of majority in a dataset [2, 12]. Conventional anomaly detection methods were developed to handle attribute-value data which is often assumed to be independent and identically distributed (*i.i.d.*). However, in many real-world scenarios, data samples are often inherently connected due to a variety of compound reasons, which naturally form the so-called *attributed networks* [24, 31]. Attributed networks are increasingly used to model a wide range of complex systems, such as social media networks, collaboration networks and gene regulatory networks [5, 6, 40]. To this end, performing anomaly detection on such networks has broad impact on various domains, including network intrusion detection [19], opinion spam detection [41], system fault diagnosis [15], and information .

Normally, label information of anomalies on attributed networks is often costly and labor intensive to obtain while it is much easier to amass a vast amount of unlabeled data. As such, a vast majority of existing anomaly detection algorithms on attributed networks are predominately applied in an unsupervised setting by modeling the distribution of observed nodes and detect anomalies based on their deformations to the learned models [20, 30, 38, 44]. The aforementioned unsupervised approaches, however, establish the models in a batch-mode fashion without any interactions with the environment. In fact, humans can often provide invaluable information by depicting whether the discovered anomalies are instrumental or not. Hence, the failure of incorporating prior knowledge of anomalies often restricts the capability of these models in identifying actual anomalies of interest among a swarm of normal instances, leading to false positives or mismatching of user interests for specific applications. Additionally, new types of anomalies may arise in complex systems to combat with existing anomaly detectors, especially in an adversary environment. For example, in network intrusion detection, new threats and tactics are continuously being developed by the attackers with the change of the environment, which pose great challenges to conventional batch-mode anomaly detection methods. Fortunately, recent advances in interactive data exploration [1, 18, 26] and human-in-the-loop machine learning [23, 25, 27] show that by interactively involving prior human knowledge in the course of learning, the performance of many learning tasks can be remarkably improved, thus the anomaly detectors can quickly sense the system changes by interacting with the human expert. For this reason, in this paper, we make the initial investigation on the problem of *interactive anomaly detection on attributed networks* by making a limited number of queries from human expert about the ground truth anomalies.

Despite the importance of studying the anomaly detection on attributed networks in an interactive setting, it remains a daunting task due to the following challenges. Firstly, various types of

anomalies could appear together on attributed networks, such as contextual anomaly, structural anomaly and community anomaly [30]. On one hand, anomaly detectors may attempt to discover one particular known type of anomaly to reduce the potential risks of the systems; on the other hand, we are also interested in discovering unexplored types of anomalies as they can provide a global view of the underlying applications. The above two scenarios result in the well-known exploration-exploitation dilemma [9]. Specifically, we aim at seeking a balance between exploiting existing known anomaly types to improve the overall detection performance and exploring new anomaly types that leads to complementary insights. Fortunately, the multi-armed bandit, as a typical reinforcement learning algorithm, provides a potential solution to address the exploration-exploitation dilemma [33, 34]. To adapt it to our problem, at each trial, the multi-armed bandit algorithm will present one candidate anomalous node (along with the side information) to the human expert and query the expert to identify if it is an anomalous node or not. The feedback information from the expert will be integrated back into the multi-armed bandit model to update its strategy at the next round. The objective is to maximize the true anomalous nodes presented to the human expert given a limited number of queries. However, existing multi-armed bandit algorithms cannot be directly applied to the networked data problems as the data *i.i.d.* assumption becomes invalid. In other words, these algorithms may result in suboptimal solutions as they cannot well model the dependencies among nodes on attributed networks. In this regard, the second challenge centers around how to model the node dependencies and incorporate it into the multi-armed bandit framework for interactive anomaly discovery.

To tackle the above challenges, we formulate the interactive anomaly detection problem on attributed networks as a multi-armed bandit problem and develop a novel contextual bandit based framework called GraphUCB. Inheriting the merits of the multi-armed bandit framework, GraphUCB provides a principled solution to handle the exploration-exploitation dilemma when making queries of ground truth anomalies with human expert knowledge. Meanwhile, the developed model fuses nodal attributes and network structure synergistically in a joint framework to enable the discovery of anomalies in an interactive fashion. In particular, we propose to measure the abnormality of each node with its own nodal attributes and the contextual information from its neighborhood. In this way, it renders a more comprehensive measure of node abnormality and leads to better anomaly detection performance. The main contributions of this paper are summarized as follow:

- **Problem Formulation:** We formally define the problem of interactive anomaly detection on attributed networks and formulate the problem with the multi-armed bandit framework, which naturally provides a solution to address the exploration-exploitation dilemma.
- **Algorithm:** We develop a principled contextual multi-armed bandit framework GraphUCB that is able to model both the nodal attributes and node dependencies seamlessly in joint framework for interactive anomaly discovery on attributed networks. In particular, the proposed framework will make a limited number of queries from the human expert regarding

Symbols	Definitions
G	the input attributed network
$\mathbf{W} \in \mathbb{R}^{N \times N}$	the relation matrix that encodes node dependencies
r	the received payoff
a	the bandit arm
d	the dimension of nodal attributes
$\{\theta, \phi\} \in \mathbb{R}^d$	coefficient vectors of an arm
$\{\hat{\theta}, \hat{\phi}\} \in \mathbb{R}^d$	estimated coefficient vectors of an arm
$\mathbf{x} \in \mathbb{R}^d$	the node attribute vector
$\mathcal{A}$	the arm set of bandit model
T	the query budget
K	the number of arms
N	the number of nodes in the network

Table 1: Table of main symbols.

to the ground truth anomalies, and the answers are integrated back to the underlying system to improve the anomaly detection performance.

- **Evaluation:** We evaluate our proposed GraphUCB framework on various real-world attributed networks from different domains. The empirical experimental results demonstrate the improvement of our proposed framework against the state-of-the-art methods, including the conventional multi-armed bandit algorithms and the unsupervised anomaly detection methods for attributed networks.

2 PROBLEM DEFINITION

Following the commonly used notations, we use bold upper-case letters for matrices (e.g., $\mathbf{X}$), bold lowercase letters for vectors (e.g., θ), calligraphic fonts for sets (e.g., $\mathcal{V}$). In addition, we represent the identity matrix as $\mathbf{I}$, and the transpose of a matrix $\mathbf{X}$ is represented as $\mathbf{X}^T$. We summarize the main notations used throughout the paper in Table 1. For the other special notations, we will illustrate them in the corresponding sections.

Definition 2.1. Attributed Networks: An attributed network $G = (\mathcal{V}, \mathcal{E}, \mathbf{X})$ consists of: (1) the set of nodes $\mathcal{V}$, where $|\mathcal{V}| = N$; (2) the set of edges $\mathcal{E}$, where the $|\mathcal{E}| = e$; and (3) the node attributes $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N]$, where $\mathbf{x}_i \in \mathbb{R}^d$ ($i = 1, \dots, N$) is the attribute¹ information for the i -th node.

By considering an attributed network as input, we aim to interactively detect anomalies on attributed networks and formalize the studied problem as follows.

PROBLEM 1. Interactive Anomaly Detection on Attributed Networks: Given an attributed network G , the task is to maximize the number of true anomalies presented to the human expert with a budget of T queries. Specifically, at each interaction trial t , the anomaly detector presents one suspicious anomalous node i to the human expert and queries if it is anomalous or not. The human expert provides feedback and then the anomaly detector updates its detection strategy by incorporating the feedback. The interaction process continues until the query budget T is used up.

¹We use attribute and feature interchangeably.

Figure 1: The proposed interactive anomaly detection framework GraphUCB. As the first step, GraphUCB generates K node clusters from the input attributed network. Then at each trial t , GraphUCB selects one suspicious node and queries the human expert to identify if it is anomalous or not. The feedback from human expert will be integrated back into the model to update its selection strategy at the next round. This interaction process will iterate until the T queries budget is used up.

Next, we will introduce the proposed method which models nodal attributes and node dependencies coherently on top of the contextual multi-armed bandit framework.

3 THE PROPOSED GRAPHUCB FRAMEWORK

In this section, we first introduce some preliminary knowledge to facilitate the understanding of the contextual multi-armed bandit framework, and then on its basis, we carefully illustrate the proposed GraphUCB framework which seamlessly models the nodal attributes and network topological structure in a joint framework to enable the interactive discovery of anomalous nodes on attributed networks. The overall workflow of the proposed GraphUCB framework is shown in Figure 1.

3.1 Contextual Multi-armed Bandit

In many real-world scenarios, we often encountered with the scenario where it is necessary to seek for a balance between exploiting the current accumulated knowledge and exploring new knowledge through searching unknown space and this problem is often known as the exploration-exploitation dilemma. The K -armed bandit problem is a classic example, where the K -armed bandit can be seen as a set of real distributions $\{D_1, D_2, \dots, D_K\}$ and each distribution is associated with a reward delivered by one of the K arms. One arm can be pulled at each trial and the corresponding reward for that action will be returned. The objective is to maximize the sum of the collected rewards through a sequence of arm pulls. The crucial trade-off at each trial is between the exploitation of the arm that has the highest expected payoff and the exploration that offers more information about the expected payoffs of pulling the other arms. Such problems have attracted an increasing attention in recent years [8, 21, 28, 42]. Specifically, contextual multi-armed bandit algorithm [16, 29, 33] become a reference solution to address the exploration-exploitation dilemma with the companion side information, where the side information is used to infer the conditional expected payoff of an action. This class of models have been widely used in many real-world applications, such as recommender systems [10, 33, 47], display advertising [13, 34] and network embedding [25]. Among them, LinUCB [33] is one of the

most widely used algorithms which uses ridge regression to obtain the expected payoff of an action and its confidence interval based on the Upper Confidence Bound (UCB) framework [7]. LinUCB is first used to solve the personalized news recommendation problem where each article is considered as an arm to pull. The expected payoff of selecting a news article (pulling an arm) at trial t can be estimated based on the dot product of its contextual feature vector with an unknown item-dependent coefficient vector:

$$\mathbb{E}[r_a | \mathbf{x}_a] = \mathbf{x}_a^T \boldsymbol{\theta}_a^*, \quad (1)$$

where $\boldsymbol{\theta}_a^*$ denotes some unknown coefficient vector of arm a and $\mathbf{x}_a$ denotes the contextual feature vector of arm a .

The goal of the contextual multi-armed bandit algorithm is to update the arm-selection strategy with respect to new observations, such that after T interaction trials its accumulated regret is minimized. The accumulated regret after T interaction trials is defined as:

$$\mathbf{R}(T) = \sum_{t=1}^T R_t = \sum_{t=1}^T (r_{a_t^*} - r_{a_t}), \quad (2)$$

where a_t is the selected arm at trial t , a_t^* is the best arm to display to the user according to the selection strategy. r_{a_t} and $r_{a_t^*}$ are the corresponding payoffs respectively, R_t is the regret of the user at trial t . In the news recommendation problem, when a presented news article is clicked by a user, a payoff of 1 is incurred; otherwise, the observed payoff is 0. Equivalently, the target of LinUCB algorithm can be reformulated to maximize the expected payoff given a budget of trials T .

3.2 Attribute Information Modeling

As a rich network representation, the attributed networks encode abundant nodal attributes that describe the characteristics of a node. Thus we will first discuss how to model the attribute information of nodes for interactive anomaly detection on attributed networks. In fact, due to a compound of reasons, various types of anomalies could appear together on attributed networks [30]. For example, as shown in the illustrative example from Figure 1, three types of anomalies coexist on the input attributed network. Among them, node 1 and node 11 are structural anomalies as they do not belong

to any communities. Note that the value of the second attribute f_2 of node 10 significantly deviates from that of the other nodes on the network, thus it is regarded as a contextual anomaly. Meanwhile, node 2 is considered as a community anomaly since its attribute value of f_1 is relatively higher compared to the other nodes (nodes 2, 3, 4, 5 and 6) in the same community. As such, the crucial dilemma we try to address during interactive anomaly detection is to seek a balance between exploiting known anomaly types to improve the overall detection performance and exploring new and unknown anomaly types that leads to complementary insights. Therefore, it motivates us to formulate this problem within a contextual multi-armed bandit framework.

In order to formulate the studied problem within the K -armed contextual bandit framework, we first group the N nodes into a set of K different clusters. Normally, nodes in the same cluster share the similar features or conform to the same patterns as they can be seen as samples drawn from the same distribution. Therefore, we assume that the nodes within the same cluster share the same model that determines how they will get pulled and get rewarded. Specifically, for each node on the network, we can regard the cluster it belongs to as an arm to pull and its features as the contextual feature vector when pulling that particular arm. Note that in our framework, various methods can be applied to group nodes into different clusters. For example, we can apply the widely used clustering methods such as K-Means, K-Medoids to obtain the node clusters. In addition to that, various community detection methods [53, 54] can also be employed in our proposed framework. In this work, we apply K-Medoids as an attempt.

As we aim to model the attribute information of nodes into the contextual multi-armed bandit framework, we can directly take the nodal attributes as the contextual feature vector. When we try to estimate the payoff of a node selection, similar to LinUCB, the estimated payoff can be computed as the dot product of the contextual feature vector with the coefficient vector of the corresponding arm. As we take each cluster as an arm, the corresponding arm $a(i)$ of node i is the cluster where node i lies in. Then the expected payoff of selecting a node i can be formulated as:

$$r_i = \mathbf{x}_i^T \boldsymbol{\theta}_{a(i)}, \quad (3)$$

where $\mathbf{x}_i$ is the contextual feature vector of node i and $\boldsymbol{\theta}_{a(i)}$ is a coefficient vector of arm $a(i)$ that node i belongs to.

For now, the estimated payoff of selecting a node only associates with the nodal attributes, thus we call this part of the estimated payoff as the *attribute payoff*.

3.3 Node Dependency Modeling

In the previous subsection, we have discussed how to model the attribute information of nodes for interactive anomaly detection with the contextual multi-armed bandit framework. However, the abnormality of a node is not solely determined by its attributes. Apart from *i.i.d.* attribute-value data, an attributed network is commonly composed of several different communities, where the nodes are densely connected and share similar attributes with each other [54]. Even though the attribute information of a node is normal over the entire dataset, the node is still considered as an anomaly if it shows highly disparate attributes with its neighbors or community

Algorithm 1 The proposed GraphUCB algorithm.

Input: $\alpha, \beta, \rho \in \mathbb{R}_+, T \in \mathbb{N}_+, \lambda \in [0, 1], \mathbf{W} \in \mathbb{R}^{N \times N}, \mathbf{X} \in \mathbb{R}^{N \times d}$
Initialize: For each arm $a \in \{a_1, \dots, a_K\}$:
 $\mathbf{A}_a \leftarrow \lambda \mathbf{I}, \mathbf{b}_a \leftarrow 0, \hat{\boldsymbol{\theta}}_a \leftarrow \mathbf{A}_a^{-1} \mathbf{b}_a$
 $\mathbf{P}_a \leftarrow \lambda \mathbf{I}, \mathbf{q}_a \leftarrow 0, \hat{\boldsymbol{\phi}}_a \leftarrow \mathbf{P}_a^{-1} \mathbf{q}_a$
1: **for** $t = 1$ to T **do**
2: **for** node $i \in \mathcal{V}_t$ **do**
3: Observe feature vectors of node i and its neighbors $\mathcal{N}(i)$
4: Choose an anomalous node i_t based on:
5:
$$i_t = \arg \max_{i \in \mathcal{V}} \left(\mathbf{x}_i^T \hat{\boldsymbol{\theta}}_{a(i)} + \alpha \sqrt{\mathbf{x}_i^T \mathbf{A}_{a(i)}^{-1} \mathbf{x}_i} + \right.$$

6:
$$\left. \rho \left(\mathbf{y}_i^T \hat{\boldsymbol{\phi}}_{a(i)} + \beta \sqrt{\mathbf{y}_i^T \mathbf{P}_{a(i)}^{-1} \mathbf{y}_i} \right) \right)$$

7: Receive payoff $r_{i_t} \in \{0, 1\}$ for selecting the node i_t
8: $\mathbf{A}_{a(i_t)} = \mathbf{A}_{a(i_t)} + \mathbf{x}_{i_t} \mathbf{x}_{i_t}^T$
9: $\mathbf{b}_{a(i_t)} = \mathbf{b}_{a(i_t)} + (r_{i_t} - \rho \mathbf{y}_{i_t}^T \hat{\boldsymbol{\phi}}_{a(i_t)}) \mathbf{x}_{i_t}$
10: $\mathbf{P}_{a(i_t)} = \mathbf{P}_{a(i_t)} + \mathbf{y}_{i_t} \mathbf{y}_{i_t}^T$
11: $\mathbf{q}_{a(i_t)} = \mathbf{q}_{a(i_t)} + \frac{(r_{i_t} - \mathbf{x}_{i_t}^T \hat{\boldsymbol{\theta}}_{a(i_t)}) \mathbf{y}_{i_t}}{\rho}$
12: $\boldsymbol{\theta}_{a(i_t)} = \mathbf{A}_{a(i_t)}^{-1} \mathbf{b}_{a(i_t)}, \boldsymbol{\phi}_{a(i_t)} = \mathbf{P}_{a(i_t)}^{-1} \mathbf{q}_{a(i_t)}$
13: **end for**
14: $\mathcal{V}_{t+1} = \mathcal{V}_t - \{i_t\}$
15: **end for**

members [31], such as node 2 in Figure 1. Consequently, it is critical to exploit such node dependencies for discovering the anomalies on the attributed network.

To model the node dependencies, we introduce a relational matrix $\mathbf{W}$ that encodes the dependencies among nodes. In this matrix, each element $\mathbf{w}_{i,j}$ is nonnegative and its value quantifies the influence that node j has on node i when determining its payoff. Also, $\mathbf{w}_{i,j} = 0$ if there is no edge between node i and node j , which means no mutual influence exists between these two nodes. For any node i , let $\mathcal{N}(i)$ denote its neighbor nodes in the given network, and for node $j \in \mathcal{N}(i)$, we normalize $\mathbf{w}_{i,j} = \frac{1}{|\mathcal{N}(i)|}$ as a direct representation of edge weight between node i and j . With the relation matrix $\mathbf{W}$, influence from the neighborhood can be captured when estimating the payoff of a node selection. Consequently, we extend Eq (1) by adding a new coefficient vector $\boldsymbol{\phi}$ to model such node dependencies from the neighborhood:

$$r_i = \mathbf{x}_i^T \boldsymbol{\theta}_{a(i)} + \rho \mathbf{y}_i^T \boldsymbol{\phi}_{a(i)} \quad (4)$$

$$s.t. \quad \mathbf{y}_i = \text{AGGR}(\{\mathbf{w}_{i,j} \mathbf{x}_j, \forall j \in \mathcal{N}(i)\}),$$

where ρ is an important parameter that controls the impact of node dependency information on payoff estimation. AGGR() is an aggregator function which aggregates the weighted features from neighbors. Here the aggregator function we propose to use is the mean operator, where we simply take the elementwise mean of the vectors in $\{\mathbf{w}_{i,j} \mathbf{x}_j, \forall j \in \mathcal{N}(i)\}$. Any other aggregator function can also be explored in our framework. Correspondingly, we name the latter part of the estimated payoff as the *dependency payoff*.

It is worth noting that the joint payoff estimation is clearly depicted in Eq (4): the estimated payoff of node i is not only determined

by the estimation of node i , but also is deeply influenced by nodes $\mathcal{N}(i)$ in its neighborhood. Namely, when our GraphUCB framework tries to determine the payoff of presenting one node to the human expert, instead of considering the attribute information of that node individually, our model will get a more comprehensive decision by leveraging both the nodal attributes and the node dependencies.

3.4 Model Parameter Learning for GraphUCB

Until now, we have discussed how to estimate the payoff a node selection for interactive anomaly detection with both the node attribute information and node dependencies. In particular, at each trial t , we will receive a payoff r_{i_t} of selecting node i_t from the human expert, indicating whether the selected node is anomalous or not based on the domain knowledge and interest of human expert. With the received payoff, we can jointly learn the contextual bandit model for detecting anomalies on the given attributed network. As the estimated payoff can be separated into two independent parts, if we treat the *dependency payoff* part as constant, the expected payoff is linear with respect to the coefficient vector θ . Likewise, the expected payoff becomes linear with respect to the coefficient vector ϕ when we take the *attribute payoff* part as constant. Thus according to the ridge regression [16, 49], we have a closed-form estimation of $\hat{\theta}_a = \mathbf{A}_a^{-1} \mathbf{b}_a$ and $\hat{\phi}_a = \mathbf{P}_a^{-1} \mathbf{q}_a$ for any arm $a \in \mathcal{A}$. Specifically, in each trial t , when node i_t is selected, $\mathbf{A}_{a(i_t)}$, $\mathbf{b}_{a(i_t)}$ and $\mathbf{P}_{a(i_t)}$, $\mathbf{q}_{a(i_t)}$ can be respectively updated by the following equations:

$$\mathbf{A}_{a(i_t)} = \mathbf{A}_{a(i_t)} + \mathbf{x}_{i_t}^T \mathbf{x}_{i_t}, \quad (5)$$

$$\mathbf{b}_{a(i_t)} = \mathbf{b}_{a(i_t)} + (r_{i_t} - \rho \mathbf{y}_{i_t}^T \phi_{a(i_t)}) \mathbf{x}_{i_t}, \quad (6)$$

$$\mathbf{P}_{a(i_t)} = \mathbf{P}_{a(i_t)} + \mathbf{y}_{i_t}^T \mathbf{y}_{i_t}, \quad (7)$$

$$\mathbf{q}_{a(i_t)} = \mathbf{q}_{a(i_t)} + \frac{(r_{i_t} - \mathbf{x}_{i_t}^T \theta_{a(i_t)}) \mathbf{y}_{i_t}}{\rho}, \quad (8)$$

where for any $a \in \mathcal{A}$, we initialize $\mathbf{A}_a = \lambda \mathbf{I}$ at the beginning [50], and $\mathbf{I}$ is a $d \times d$ identity matrix.

Different from the LinUCB framework, by virtue of incorporating the node dependency information into the payoff estimation, the uncertainty of estimation for the coefficient vectors comes from both the *attribute payoff* part and the *dependency payoff* part. As shown in [49], the uncertainty of estimation for coefficient vectors in the *attribute payoff* part can be formulated as:

$$\|\mathbf{x}_i\|_{\mathbf{A}_{a(i)}^{-1}} = \sqrt{\mathbf{x}_i^T \mathbf{A}_{a(i)}^{-1} \mathbf{x}_i}, \quad (9)$$

where $\mathbf{A}_{a(i)}^{-1}$ is the inverse covariance matrix for the *attribute payoff* part of arm $a(i)$ in the round t .

Similarly, if we take *attribute payoff* part as constant, the uncertainty of the estimation for coefficient vectors $\phi_{a(i)}$ can be computed as:

$$\|\mathbf{y}_i\|_{\mathbf{P}_{a(i)}^{-1}} = \sqrt{\mathbf{y}_i^T \mathbf{P}_{a(i)}^{-1} \mathbf{y}_i}, \quad (10)$$

where $\mathbf{P}_{a(i)}^{-1}$ is the inverse covariance matrix for the *dependency payoff* of arm $a(i)$. Hence, a tight upper confidence bound(UCB) for the expected payoff of selecting a particular arm can be derived based on Eq (9) and Eq (10). At each trial t , our GraphUCB framework will choose an anomalous node according to the following

	BlogCatalog	Flickr	ACM
# nodes	5,196	7,575	16,484
# edges	171,743	239,738	71,980
# attributes	8,189	12,047	8,337
# labels	6	9	9
# anomalies	300	450	600

Table 2: Details of the used datasets.

selection strategy:

$$i_t = \arg \max_{i \in \mathcal{V}} \left(\mathbf{x}_i^T \hat{\theta}_{a(i)} + \alpha \sqrt{\mathbf{x}_i^T \mathbf{A}_{a(i)}^{-1} \mathbf{x}_i} + \rho \left(\mathbf{y}_i^T \hat{\phi}_{a(i)} + \beta \sqrt{\mathbf{y}_i^T \mathbf{P}_{a(i)}^{-1} \mathbf{y}_i} \right) \right). \quad (11)$$

The detailed description of the proposed GraphUCB framework is illustrated in Algorithm 1. In this algorithm, the nodal attributes and node dependencies are clearly modeled into the payoff estimation, which are indispensable for identifying the abnormality of nodes. It is worth noting that, with the feedback from human expert continuously integrated back into our model, the node selection strategy can be updated along with the interests of human expert over time. For GraphUCB, another advantage of integrating the *dependency payoff* part into the payoff estimation is that the confidence interval of the expected payoff also comes from two different parts, which will reduce the estimation uncertainty.

4 EXPERIMENTS

In this section, we perform extensive empirical evaluations on real-world attributed networks to verify the effectiveness of the proposed GraphUCB framework in interactive anomaly detection. Before presenting the detailed experimental results, we first introduce the used datasets and then present the experimental settings.

4.1 Datasets

In the experiments, we collect and use three real-world attributed networks that have been widely used in prior research [25, 32] to evaluate the performance of different methods:

BlogCatalog: BlogCatalog is a blog sharing website. The bloggers in blogcatalog can follow each other forming a social network. Users are associated with a list of tags to describe themselves and their blogs, which are regarded as node attributes. The predefined groups that the bloggers subscribed are taken as the class labels.

Flickr: Flickr is an image hosting and sharing website. Similar to BlogCatalog, users can follow each other and form a social network. Node attributes of users are defined by their specified tags that reflect their interests. Users could also join some predefined groups and are taken as the class labels.

ACM: This dataset collects a citation network for published papers before 2016. Each paper is regarded as a node in the network, and the links are the citation relations among different papers. We apply the bag-of-words model on the paper abstract to obtain node attributes. We select papers from nine different research areas and they are taken as the class labels of papers.

To evaluate the effectiveness of the proposed GraphUCB framework in conducting interactive anomaly detection, we propose to

Figure 2: Anomaly discovery curve results.

inject anomalies into the attributed networks. In particular, we refer to two widely used methods [45, 46] to generate a combined set of anomalies for each dataset from both the network structural and nodal attribute information perspectives. In real-world networked data, anomalous substructures are usually created by suspicious activities or behaviors. For example, small clique is a typical anomalous substructure in which a small set of nodes are much more closely linked to each other than average [45]. Thus in terms of the injected structural anomalies, we choose to add links to make a number of nodes to form small cliques. Specifically, if we specify the clique size as m , we first randomly select m nodes from the network, and then make those nodes fully connected with each other, and then all the m nodes in the clique are considered as anomalies. We iterate this process until a number of n cliques are generated, then the number of structural anomalies is $m \times n$. In our experiments, we fix the clique size m to 15 and n varies for different datasets. In addition, we adopt the attribute perturbation schema introduced by Song et al. [46] to generate the other type of anomaly. Specifically, for a selected node i , we randomly pick another k nodes from the data and select the node j with the most different attributes from node i among the k nodes, i.e., maximize the Euclidean distance $\|x_i - x_j\|_2$. Afterwards, we then replace the attributes x_i of node i by x_j . The value of k is set to be 25 here. In the experiments, we inject an equal number of anomalies from both the structural perspective and the attribute perspective, the details of these three used datasets are shown in Table 3.

4.2 Experiments Settings

In this section, we introduce the detailed experimental settings, including the compared baseline methods, data preprocessing and evaluation metrics.

4.2.1 Compared Methods. We compare the proposed GraphUCB with the following two categories of methods, including conventional multi-armed bandit approaches ϵ -greedy, LinUCB, LinTS and the state-of-the-art anomaly detection methods on attributed networks Radar and ANOMALOUS.

- ϵ -greedy is one of the most popular exploration-exploitation strategies in literature. In our problem setting, at trial t , it picks the node with the highest estimated reward based on

the current knowledge with probability $1 - \epsilon$ and randomly picks a node with probability ϵ .

- **Linear UCB (LinUCB)** [33] is a linear model under the UCB framework by combining linear bandit and contextual bandit together. In our scenario, the estimated payoff r_{a_t} can be obtained through the dot product of the node-dependent coefficient with the node attribute information.
- **Contextual Thompson Sampling (LinTS)** [3] is also a contextual multi-armed bandit algorithm. It is based on Thompson Sampling and is designed for the stochastic problem with linear payoff functions.
- **Radar** [30] is one of the state-of-the-art unsupervised anomaly detection frameworks for attributed networks. It detects anomalies whose behaviors are singularly different from the majority by characterizing the residuals of attribute information and its coherence with network information.
- **ANOMALOUS** [37] performs joint anomaly detection and attribute selection to detect anomalies on attributed networks based on the CUR decomposition and residual analysis.

4.2.2 Data Preprocessing. The feature dimension of the three aforementioned datasets could be exceedingly large. Directly making use of the data of high dimensionality would be problematic as the feature representation is often very sparse and poses great challenges to learning and inference due to the curse of dimensionality. To facilitate the learning process, in our experiments, we employ Principal Component Analysis (PCA) to reduce the dimensionality of node features, and then use the low-dimensional node representation as the contextual vectors for contextual multi-armed bandit models. The dimensionality of the contextual feature vector is set to as 20.

4.2.3 Evaluation Metrics. Under the interactive anomaly detection setting, for any multi-armed bandit models, the detector presents one node according to its node selection strategy at each trial, then we can verify the result with the ground truth anomalies and return the feedback. Thus we adopt several widely used metrics to compare different methods. The evaluation metrics used in the experiments include:

- **Anomaly Discovery Curve** - Since our objective is to maximize the number of true anomalies presented to the human expert

Cumulative Precision@T															
	BlogCatalog					Flickr					ACM				
Round T	50	100	150	200	250	50	100	150	200	250	50	100	150	200	250
ϵ -greedy	0.500	0.440	0.327	0.250	0.204	0.340	0.320	0.253	0.270	0.260	0.400	0.350	0.300	0.265	0.272
LinUCB	0.540	0.620	0.573	0.520	0.432	0.460	0.430	0.380	0.375	0.380	0.240	0.350	0.400	0.405	0.368
LinTS	0.560	0.570	0.573	0.528	0.488	0.440	0.410	0.407	0.395	0.376	0.300	0.290	0.293	0.300	0.292
Radar	0.540	0.570	0.540	0.485	0.444	0.440	0.390	0.327	0.345	0.340	0.180	0.230	0.247	0.263	0.264
ANOMALOUS	0.560	0.550	0.460	0.425	0.400	0.420	0.410	0.380	0.370	0.348	—	—	—	—	—
GraphUCB	0.760	0.730	0.693	0.630	0.552	0.580	0.530	0.520	0.550	0.504	0.540	0.500	0.480	0.485	0.456
Cumulative Recall@T															
	BlogCatalog					Flickr					ACM				
Round T	50	100	150	200	250	50	100	150	200	250	50	100	150	200	250
ϵ -greedy	0.090	0.147	0.163	0.167	0.170	0.057	0.107	0.127	0.180	0.217	0.033	0.058	0.075	0.088	0.113
LinUCB	0.090	0.207	0.287	0.347	0.360	0.077	0.143	0.190	0.250	0.317	0.020	0.058	0.100	0.135	0.153
LinTS	0.093	0.190	0.287	0.343	0.407	0.073	0.137	0.203	0.263	0.313	0.025	0.048	0.073	0.100	0.122
Radar	0.090	0.190	0.270	0.323	0.370	0.073	0.130	0.163	0.230	0.283	0.015	0.038	0.062	0.090	0.110
ANOMALOUS	0.093	0.183	0.230	0.283	0.333	0.070	0.143	0.190	0.247	0.290	—	—	—	—	—
GraphUCB	0.127	0.243	0.347	0.420	0.460	0.097	0.177	0.260	0.367	0.420	0.045	0.083	0.120	0.162	0.190

Table 3: Results of cumulative precision@T and recall@T.

given a budget, we plot the total number of true anomalies discovered against the number of queries presented to the user. Ideally, this curve should climb as quickly as possible [17].

- *Cumulative Precision@T* - We evaluate the proportion of true anomalies that we discovered in the top T queries:

$$\text{Precision@T} = \frac{|\text{true anomalies discovered}| \cap |\text{queried anomalies}|}{|\text{queried anomalies}|}.$$

- *Cumulative Recall@T* - It measures the proportion of true anomalies that we discovered in the total number of ground truth anomalies:

$$\text{Recall@T} = \frac{|\text{true anomalies discovered}| \cap |\text{queried anomalies}|}{|\text{all true anomalies}|}.$$

In addition to these multi-armed bandit algorithms that are interactive in nature, we also compare the proposed GraphUCB framework with several unsupervised anomaly detection methods. To have a fair comparison with these methods and make them adaptable to our interactive setting, we assume that the anomalous nodes ranked by these unsupervised methods will be presented to the human expert one by one according to the descending order of the anomaly scores. Thus we can also compute the results in terms of the aforementioned three evaluation metrics.

4.3 Evaluation Results

In the experiments, we evaluate the performance of our proposed GraphUCB framework by comparing it with the aforementioned baseline methods. The budget number T is specified as 250. Figure 2 presents the anomaly discovery curve of all the algorithms on all three datasets from $t = 1$ to T . Meanwhile, we also report the cumulative precision and recall results in Table 3. From the evaluation results, we make the following observations:

- The proposed GraphUCB framework outperforms all the baseline methods on all the three datasets. We also perform a pairwise Wilcoxon signed-rank test between GraphUCB and these baseline methods. The comparison results indicate that

the proposed GraphUCB framework is significantly better than others, with a significance level of 0.05.

- The contextual bandit based methods including the proposed GraphUCB framework, LinUCB, LinTS are superior to the conventional unsupervised anomaly detection methods. It verifies that the attributed network anomaly detection performance can be remarkably improved through incorporating the expert domain knowledge in an interactive fashion.
- GraphUCB shows considerable stronger ability in detecting anomalies on attributed networks than LinUCB and LinTS, especially when the given query budget is small. The reason is that both LinUCB and LinTS neglect the node dependency information on attributed networks when making the node selection strategy. This observation supports the assumption that node dependency information is indispensable for anomaly detection on attributed networks.
- Note that we do not report the results of ANOMALOUS on the ACM dataset due to the issue of running out of memory. Since ANOMALOUS is based on CUR decomposition, it cannot be easily scaled to large-scale networks.

4.4 Parameter Analysis

Next, we investigate the impacts of three controlling parameters in our proposed GraphUCB framework, and report the performance variance results (cumulative precision and recall when $T = 250$) on the BlogCatalog dataset in Figure 3. Among the three parameters, ρ controls the impact of *dependency payoff* when determining the estimated payoff of selecting a node, α and β are the parameters that control the balance between exploration and exploitation when the model makes the node selection decision. In Figure 3(a), as we set $\alpha = 10$, $\beta = 0.01$, we observe that the anomaly detection performance is not sensitive when ρ is in the range of 10^{-4} to 1. With the growth of ρ , the performance reaches the peak when ρ is around 10 and then gradually decreases if ρ increases. The result indicates that it is necessary to find a balance between the *attribute payoff* and the *dependency payoff* when estimating the payoff of selecting a node for the problem of interactive anomaly detection on attributed networks. Similar to the results of parameter ρ , the plots

Figure 3: Results of parameter analysis on Blogcatalog.

corresponding to parameter α and β also conform to a convex shape in the experiments. The optimal value of α is around 10 when we set $\rho = 10, \beta = 0.01$, and the best performance is observed when β is set around 0.01 with both ρ and α specified as 10. Therefore, seeking a balance to address the exploration-exploitation dilemma is also critical for enhancing the overall anomaly detection performance. In addition, the parameter analysis results on the datasets of Flickr and ACM reveal the similar trend, thus we omit the results here.

5 RELATED WORK

We briefly review related work from two perspectives: (1) graph based anomaly detection; and (2) multi-armed bandit algorithms.

5.1 Graph Based Anomaly Detection

Graph based anomaly detection methods can be generally divided into two categories: 1) anomaly detection on *plain networks*; 2) anomaly detection on *attributed networks* [5]. For a given plain network, the only available information we can leverage for detecting anomalous nodes is the network structure. Therefore, this category of anomaly detection methods aim to find patterns and spot anomalies by exploiting the network structure information from different perspectives [4, 22, 52], either at the structural level or at the community level [5]. Different from plain networks, rich attributes information can be easily observed on attributed networks and could have a strong connection with the network structure. Hence, recent years we have witnessed an increasingly amount of efforts in leveraging the structure information as well as its coherence with nodal attributes to spot anomalies on attributed networks [20, 35, 37–39, 43, 44]. CODA [20] is one of the first attempts that simultaneously finds communities as well as spots community anomalies within a unified probabilistic model. Later on, researchers found that complex anomalies in attributed networks could be revealed in only a subset of relevant attributes, thus they proposed to integrate subspace selection and anomaly detection together as a solution, which proves to achieve superior performance. Among them, Consub+CODA [44] carries out subspace selection first and then use CODA to detect anomalies. ConOut [43] identifies the local context for each node and performs anomaly ranking within the local context. ANOMALOUS [37] performs anomaly detection and attribute selection with CUR decomposition in a joint manner. As different types of anomalies could coexist on attributed networks, Radar [31] proposes to detect anomalous nodes in a more general way with the residual analysis.

5.2 Multi-armed Bandit Algorithms

Multi-armed bandit algorithms provide principled solutions to address the well-known exploration-exploitation dilemma, which is about to make a trade-off to obtain new knowledge and the need to use that knowledge to improve the learning performance. As opposed to the traditional context-free multi-armed bandit algorithms [8, 21, 28, 42], contextual multi-armed bandit algorithms utilize the companion side information to infer the expected pay-off, which have attracted lots of attention as they often lead to better learning performance than the context-free multi-armed bandit algorithms in many applications, including recommender systems [16, 33, 47] and network embedding [25]. As another class of related works, prior research also explored the idea of modeling dependency among bandits [11, 49, 50], and showed that the exploitation of such dependency could lead to a dramatic performance increase of conventional bandit algorithms. For instance, the GOB.Lin algorithm [11] utilizes a graph Laplacian term to regularize the model so that users and their friends have similar bandit parameters. On the other hand, CoLin [50] assumes that the reward in bandit is generated through an additive model, indicating that friends' feedback on their recommendations can be passed via the network to explain the target user's feedback. Recently, researchers also tried to capture the arm dependency by organizing different arms into different clusters [36, 48] and it is very similar to our studied problem. However, simply performing clustering on attributed networks may ignore the inherent node dependencies and may lead to suboptimal results in interactive anomaly discovery.

6 CONCLUSION AND FUTURE WORK

In this paper, we make an initial investigation of the research problem of interactive anomaly detection on attributed networks and present a novel contextual multi-armed bandit algorithm GraphUCB. Particularly, by interactively incorporating the feedbacks from the human expert about the queried anomalies, our proposed algorithm can remarkably enhance the detection performance of conventional anomaly detectors on attributed networks within a limited number of query budget. In addition, GraphUCB provides a systematic way to fuse both the nodal attribute information and node dependencies seamlessly in a joint framework to address the exploration-exploitation dilemma. In this way, it leads to a balance between exploiting unknown anomaly types and exploring

new and unknown anomaly types. To corroborate the effectiveness of the proposed GraphUCB framework, we perform extensive experiments on multiple real-world attributed networks, the experimental results demonstrate the superiority of GraphUCB over the state-of-the-art approaches.

In our current setting, some open issues are valuable for future study. First of all, we assumed the relation matrix that encodes the node dependencies is given a prior and is fed into the bandit algorithm. However, the strength of links could vary remarkably and treat all links equally cannot fully capture the node dependencies. Hence, it is meaningful to characterize the tie strength [51] when we perform interactive anomaly detection. In addition to that, the method developed in this paper measures the node abnormality based on its local neighborhood structure. Incorporating the measurement of node abnormality from a global view is another topic that needs deeper investigation. Another direction is to investigate the interplay between anomaly detection and network connectivity optimization [14] to make the whole system more robust.

ACKNOWLEDGEMENTS

This material is based upon work supported by, or in part by, the NSF grant 1614576, and the ONR grant N00014-16-1-2257. The authors also wish to thank Qinghai Zhou for the help of the experiments.

REFERENCES

- [1] E. Aichert, H.-P. Kriegel, E. Schubert, and A. Zimek. Interactive data mining with 3d-parallel-coordinate-trees. In *SIGMOD*, pages 1009–1012, 2013.
- [2] C. C. Aggarwal. Outlier analysis. In *Data Mining*, pages 237–263, 2015.
- [3] S. Agrawal and N. Goyal. Thompson sampling for contextual bandits with linear payoffs. In *ICML*, pages 127–135, 2013.
- [4] L. Akoglu, M. McGlohon, and C. Faloutsos. Oddball: Spotting anomalies in weighted graphs. In *PAKDD*, pages 410–421, 2010.
- [5] L. Akoglu, H. Tong, and D. Koutra. Graph based anomaly detection and description: A survey. *Data Mining and Knowledge Discovery*, 29(3):626–688, 2015.
- [6] L. Akoglu, H. Tong, B. Meeder, and C. Faloutsos. PICS: Parameter-free identification of cohesive subgroups in large attributed graphs. In *SDM*, pages 439–450, 2012.
- [7] P. Auer. Using confidence bounds for exploitation-exploration trade-offs. *Machine Learning Research*, 3(Nov):397–422, 2002.
- [8] P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2-3):235–256, 2002.
- [9] M. J. Benner and M. L. Tushman. Exploitation, exploration, and process management: The productivity dilemma revisited. *Academy of Management Review*, 28(2):238–256, 2003.
- [10] D. Bouneffouf, A. Bouzeghoub, and A. L. Gançarski. A contextual-bandit algorithm for mobile context-aware recommender system. In *ICONIP*, pages 324–331, 2012.
- [11] N. Cesa-Bianchi, C. Gentile, and G. Zappella. A gang of bandits. In *NIPS*, pages 737–745, 2013.
- [12] V. Chandola, A. Banerjee, and V. Kumar. Anomaly detection: A survey. *Computing Surveys*, 41(3):15, 2009.
- [13] O. Chapelle and L. Li. An empirical evaluation of thompson sampling. In *NIPS*, pages 2249–2257, 2011.
- [14] C. Chen, H. Tong, B. A. Prakash, T. Eliassi-Rad, M. Faloutsos, and C. Faloutsos. Eigen-optimization on large graphs by edge manipulation. *TKDD*, 10(4):49, 2016.
- [15] W. Cheng, K. Zhang, H. Chen, G. Jiang, Z. Chen, and W. Wang. Ranking causal anomalies via temporal and dynamical analysis on vanishing correlations. In *KDD*, pages 805–814, 2016.
- [16] W. Chu, L. Li, L. Reyzin, and R. Schapire. Contextual bandits with linear payoff functions. In *AISTATS*, pages 208–214, 2011.
- [17] S. Das, W.-K. Wong, T. Dietterich, A. Fern, and A. Emmott. Incorporating expert feedback into active anomaly discovery. In *ICDM*, pages 853–858, 2016.
- [18] K. Dimitriadou, O. Papaemmanouil, and Y. Diaio. Explore-by-example: An automatic query steering framework for interactive data exploration. In *SIGMOD*, pages 517–528, 2014.
- [19] Q. Ding, N. Katanka, P. Barford, E. Kolaczyk, and M. Crovella. Intrusion as (anti) social communication: Characterization and detection. In *KDD*, pages 886–894, 2012.
- [20] J. Gao, F. Liang, W. Fan, C. Wang, Y. Sun, and J. Han. On community outliers and their efficient detection in information networks. In *KDD*, pages 813–822, 2010.
- [21] J. C. Gittins. Bandit processes and dynamic allocation indices. *Royal Statistical Society*, pages 148–177, 1979.
- [22] K. Henderson, B. Gallagher, L. Li, L. Akoglu, T. Eliassi-Rad, H. Tong, and C. Faloutsos. It's who you know: Graph mining using recursive structural features. In *KDD*, pages 663–671, 2011.
- [23] A. Holzinger. Interactive machine learning for health informatics: when do we need the human-in-the-loop? *Brain Informatics*, 3(2):119–131, 2016.
- [24] X. Huang, J. Li, and X. Hu. Label informed attributed network embedding. In *WSDM*, pages 731–739, 2017.
- [25] X. Huang, Q. Song, J. Li, and X. Hu. Exploring expert cognition for attributed network embedding. In *WSDM*, 2018.
- [26] A. Kalinin, U. Cetintemel, and S. Zdonik. Interactive data exploration using semantic windows. In *SIGMOD*, pages 505–516, 2014.
- [27] W. Karwowski. *International encyclopedia of ergonomics and human factors*, volume 3. 2001.
- [28] T. L. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.
- [29] J. Langford and T. Zhang. The epoch-greedy algorithm for multi-armed bandits with side information. In *NIPS*, pages 817–824, 2008.
- [30] J. Li, H. Dani, X. Hu, and H. Liu. Radar: Residual analysis for anomaly detection in attributed networks. In *IJCAI*, pages 2152–2158, 2017.
- [31] J. Li, H. Dani, X. Hu, J. Tang, Y. Chang, and H. Liu. Attributed network embedding for learning in a dynamic environment. In *CIKM*, pages 387–396, 2017.
- [32] J. Li, X. Hu, J. Tang, and H. Liu. Unsupervised streaming feature selection in social media. In *CIKM*, pages 1041–1050, 2015.
- [33] L. Li, W. Chu, J. Langford, and R. E. Schapire. A contextual-bandit approach to personalized news article recommendation. In *WWW*, pages 661–670, 2010.
- [34] W. Li, X. Wang, R. Zhang, Y. Cui, J. Mao, and R. Jin. Exploitation and exploration in a performance based contextual advertising system. In *KDD*, pages 27–36, 2010.
- [35] E. Muller, P. I. Sánchez, Y. Mulle, and K. Böhm. Ranking outlier nodes in subspaces of attributed graphs. In *ICDE Workshops*, pages 216–222, 2013.
- [36] S. Pandey, D. Chakrabarti, and D. Agarwal. Multi-armed bandit problems with dependent arms. In *ICML*, pages 721–728, 2007.
- [37] Z. Peng, M. Luo, J. Li, H. Liu, and Q. Zheng. Anomalous: A joint modeling approach for anomaly detection on attributed networks. In *IJCAI*, pages 3513–3519, 2018.
- [38] B. Perozzi and L. Akoglu. Scalable anomaly ranking of attributed neighborhoods. In *SDM*, pages 207–215, 2016.
- [39] B. Perozzi, L. Akoglu, P. Iglesias Sánchez, and E. Müller. Focused clustering and outlier detection in large attributed graphs. In *KDD*, pages 1346–1355. ACM, 2014.
- [40] J. J. Pfeiffer III, S. Moreno, T. La Fond, J. Neville, and B. Gallagher. Attributed graph models: Modeling network structure with correlated attributes. In *WWW*, pages 831–842, 2014.
- [41] S. Rayana and L. Akoglu. Collective opinion spam detection: Bridging review networks and metadata. In *KDD*, pages 985–994, 2015.
- [42] H. Robbins. Some aspects of the sequential design of experiments. In *Herbert Robbins Selected Papers*, pages 169–177, 1985.
- [43] P. I. Sánchez, E. Müller, O. Irmeler, and K. Böhm. Local context selection for outlier ranking in graphs with multiple numeric node attributes. In *SSDBM*, page 16, 2014.
- [44] P. I. Sánchez, E. Muller, F. Laforet, F. Keller, and K. Böhm. Statistical selection of congruent subspaces for mining attributed graphs. In *ICDM*, pages 647–656, 2013.
- [45] D. B. Skillicorn. Detecting anomalies in graphs. In *ISI*, pages 209–216, 2007.
- [46] X. Song, M. Wu, C. Jermaine, and S. Ranka. Conditional anomaly detection. *TKDE*, 19(5):631–645, 2007.
- [47] S.-Y. Teng, J. Li, L. P.-Y. Ting, K.-T. Chuang, and H. Liu. Interactive unknowns recommendation in e-learning systems. In *ICDM*, 2018.
- [48] Q. Wang, C. Zeng, W. Zhou, T. Li, L. Shwartz, and G. Y. Grabarnik. Online interactive collaborative filtering using multi-armed bandit with dependent arms. *arXiv preprint arXiv:1708.03058*, 2017.
- [49] X. Wang, S. C. Hoi, C. Liu, and M. Ester. Interactive social recommendation. In *CIKM*, pages 357–366, 2017.
- [50] Q. Wu, H. Wang, Q. Gu, and H. Wang. Contextual bandits in a collaborative environment. In *SIGIR*, pages 529–538, 2016.
- [51] R. Xiang, J. Neville, and M. Rogati. Modeling relationship strength in online social networks. In *WWW*, pages 981–990, 2010.
- [52] X. Xu, N. Yuruk, Z. Feng, and T. A. Schweiger. Scan: A structural clustering algorithm for networks. In *KDD*, pages 824–833, 2007.
- [53] Z. Xu, Y. Ke, Y. Wang, H. Cheng, and J. Cheng. A model-based approach to attributed graph clustering. In *SIGMOD*, pages 505–516, 2012.
- [54] J. Yang, J. McAuley, and J. Leskovec. Community detection in networks with node attributes. In *ICDM*, pages 1151–1156, 2013.