

AURORA: Auditing PageRank on Large Graphs

Jian Kang

Arizona State University

jkang51@asu.edu

Meijia Wang

Arizona State University

mwang164@asu.edu

Nan Cao

Tongji University

nan.cao@gmail.com

Yinglong Xia

Huawei Research America

yinglong.xia@huawei.com

Wei Fan

Tencent

wei.fan@gmail.com

Hanghang Tong

Arizona State University

hanghang.tong@asu.edu

Abstract—Ranking on large-scale graphs plays a fundamental role in many high-impact application domains, ranging from information retrieval, recommender systems, sports team management, biology to neuroscience and many more. PageRank, together with many of its random walk based variants, has become one of the most well-known and widely used algorithms, due to its mathematical elegance and the superior performance across a variety of application domains. Important as it might be, state-of-the-art lacks an intuitive way to *explain* the ranking results by PageRank (or its variants), e.g., *why* it thinks the returned top- k webpages are the most important ones in the entire graph; *why* it gives a higher rank to actor John than actor Smith in terms of their relevance w.r.t. a particular movie?

In order to answer these questions, this paper proposes a paradigm shift for PageRank, from identifying *which* nodes are most important to understanding *why* the ranking algorithm gives a particular ranking result. We formally define the PageRank auditing problem, whose central idea is to identify a set of key graph elements (e.g., edges, nodes, subgraphs) with the highest influence on the ranking results. We formulate it as an optimization problem and propose a family of effective and scalable algorithms (AURORA) to solve it. Our algorithms measure the influence of graph elements and incrementally select influential elements w.r.t. their gradients over the ranking results. We perform extensive empirical evaluations on real-world datasets, which demonstrate that the proposed methods (AURORA) provide intuitive explanations with a linear scalability.

Index Terms—Graph mining, PageRank, explainability

I. INTRODUCTION

Ranking on graph data is a way to measure node importance and plays a fundamental role in many real-world applications, ranging from information retrieval [1], recommender systems [2], social networks [3], sports team management [4] to biology [5] and neuroscience [6]. Among others, PageRank [1], together with many of its random walk based variants, is one of the most well-known and widely used ones. Its mathematical elegance lies in that it only requires the topological structure and the associated edge weights as the input. Such a generality makes it applicable to networks¹ from many different application domains.

PageRank has a strong ability answering questions like *what* is the most important page in the World Wide Web;

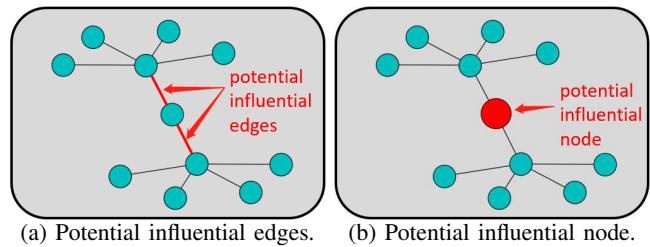


Fig. 1: Example of potential influential edges and node.

who is the most influential person in a collaboration network. Despite its superior performance on graph ranking, PageRank lacks intuitive ways to give answers to questions like *why* the top- k returned webpages are the most important ones; *why* actor John ranks higher than actor Smith in terms of their relevance w.r.t. a particular movie. *How* the ranking results are derived from the underlying graph structure has largely remained opaque to the end users, who are often not experts in data mining and mathematics.

To address this challenge, we aim to explain PageRank by finding the influential graph elements (e.g., edges, nodes and subgraphs), which we formulate as *PageRank auditing* problem. The key idea is to quantitatively understand how the ranking results would change if we perturb a specific graph element. To be specific, we measure the influence of each graph element by the rate of change in a certain loss function (e.g., L_p norm, etc.) defined over the ranking vector. We believe that auditing graph ranking can benefit many real-world applications. First, it can render the crucial explainability of such ranking algorithms, by identifying valuable information of influential graph elements. Thus, it can help answer questions like why a given node ranks on the top of the ranking list, which link leads to the ranking vector in a certain way. Furthermore, users can use the auditing results to optimize the network topology. In addition, it may help identify the vulnerabilities in the network (e.g. links between two clusters, cutpoints of clusters as shown in Figure 1a and Figure 1b). With the auditing result, users may find several links or nodes that have the greatest influence on the ranking results. It can help users identify if there exist suspicious individuals that manipulate the ranking results by linking heavily with unrelated, off-query topics. Finally, with

¹In this paper, we use ‘graph’ and ‘network’ interchangably.

TABLE I: Table of Symbols

Symbols	Definitions
$\mathbf{G} = (\mathcal{V}, \mathcal{E})$	the input network
(i, j)	edge from node i to node j
\mathbf{A}	adjacency matrix of the input network
$\mathbf{A}(i, j)$	the element at i^{th} row and j^{th} column
$\mathbf{A}(i, :)$	i^{th} row of matrix \mathbf{A}
$\mathbf{A}(:, j)$	j^{th} column of matrix \mathbf{A}
\mathbf{A}'	transpose of the matrix \mathbf{A}
\mathbf{A}^{-1}	inverse of the matrix \mathbf{A}
\mathbf{e}	the teleportation vector in PageRank
\mathbf{r}	PageRank of the input network
$\mathbf{r}(i)$	ranking score of i
$\text{Tr}(\mathbf{A})$	Trace of the matrix \mathbf{A}
$f(\mathbf{r})$	a loss function over ranking vector \mathbf{r}
n	number of nodes in the input network
m	number of edges in the input network
c	damping factor in PageRank

the knowledge of graph ranking auditing, we may design a more robust ranking algorithm that is hard to be manipulated by users with strategies like linking heavily [7].

The main contributions of the paper are summarized as follows.

- **Problem Definition.** We formally define PageRank auditing problem, and formulate it as an optimization problem, whose key idea is to measure the influence of different graph elements as the rate of change in the PageRank results.
- **Algorithms and Analysis.** We propose fast approximation algorithms to solve the PageRank auditing problem. The algorithms achieve a $(1 - 1/e)$ approximation ratio with a linear complexity.
- **Empirical Evaluations.** We perform extensive experiments on diverse, real-world datasets. The experimental results demonstrate that our proposed methods (a) provide reasonable and intuitive information to find influential graph elements, and (b) scale linearly w.r.t. the graph size.

The rest of the paper is organized as follows. Section II formally defines the PageRank auditing problem. Section III introduces our proposed algorithms. Then we provide experimental evaluations in Section IV. After reviewing related work in Section V, we conclude the paper in Section VI.

II. PROBLEM DEFINITION

In this section, we first present a table of symbols (Table 1) that contains the main notations used throughout the paper, and then review PageRank for ranking nodes on graphs, and finally give a formal definition of the PageRank auditing problem.

We use bold upper-case letters for matrices (e.g., \mathbf{A}), bold lower-case letters for vectors (e.g., \mathbf{a}), calligraphic fonts for sets (e.g., \mathcal{S}), and lower-case letters for scalars (e.g., c). For matrix indexing conventions, we use the rules similar to Matlab that are shown as follows. We use $\mathbf{A}(i, j)$ to denote the entry of matrix \mathbf{A} at i^{th} row and j^{th} column, $\mathbf{A}(i, :)$ to denote the i^{th} row of matrix \mathbf{A} , and $\mathbf{A}(:, j)$ to denote the j^{th} column of matrix \mathbf{A} . We use prime to denote the transpose of matrix (i.e., \mathbf{A}' is the transpose of matrix \mathbf{A}).

Given a graph \mathbf{G} with n nodes and m edges, PageRank essentially solves the following linear system,

$$\mathbf{r} = c\mathbf{A}\mathbf{r} + (1 - c)\mathbf{e} \quad (1)$$

where \mathbf{e} is the teleportation vector with length n and \mathbf{A} is the normalized adjacency matrix of the input graph. In PageRank, \mathbf{e} is chosen as the uniform distribution $\frac{1}{n}\mathbf{1}$; in personalized PageRank, \mathbf{e} is a biased vector which reflects user's preference (i.e., 'personalization') [8]; in random walk with restart [9], all the probabilities are concentrated on a single node. A default choice for the normalized adjacency matrix \mathbf{A} is the column-normalized matrix (i.e., the stochastic matrix). A popular alternative choice is the normalized graph Laplacian matrix. In fact, as long as the largest eigenvalue of \mathbf{A} is less than $1/c$, a fix-point solution of the above linear system will converge to $\mathbf{r} = (1 - c)(\mathbf{I} - c\mathbf{A})^{-1}\mathbf{e}$ [10]. In this paper, we use this general form for the adjacency matrix \mathbf{A} . For the ease of description, we also define $\mathbf{Q} = (\mathbf{I} - c\mathbf{A})^{-1}$, and $\mathbf{r} = \text{pg}(\mathbf{A}, \mathbf{e}, c)$ as the resulting PageRank vector with \mathbf{A} , \mathbf{e} , and c as the corresponding inputs.

Regarding explainable learning and mining techniques, Pang et al. [11] propose a novel notation of influence functions to quantify the impact of each training example on the underlying learning system (e.g., a classifier). The key idea is to trace the model's prediction back to its training data, where the model parameters were derived. In this way, it learns how a perturbation of a single training data will affect the resulting model parameters, and then identifies the training examples that are most responsible for model's predictions.

Our proposed method to explain the PageRank results is built upon the principle outlined in [11]. To be specific, we aim to find a set of graph elements (e.g., edges, nodes, a subgraph) such that, when we perturb/remove them, the ranking vector will have the greatest change. Formally, we define PageRank auditing problem as follows:

Problem 1: PageRank Auditing Problem.

Given: a graph with adjacency matrix \mathbf{A} , PageRank vector \mathbf{r} , a loss function f over its PageRank vector, user-specific element type (e.g. edges vs. nodes vs. subgraphs), and an integer budget k ;

Find: a set of k influential graph elements that has the largest impact on the loss function over its PageRank vector $f(\mathbf{r})$.

III. PROPOSED ALGORITHM

In this section, we propose a family of algorithms (AURORA), to solve PageRank auditing problem (Problem 1), together with some analysis in terms of effectiveness as well as efficiency.

A. Formulation

The intuition behind the proposed methods is to find a set of key graph elements (e.g., edges, nodes, subgraphs) whose perturbation/removal from the graph would affect the PageRank results most. To be specific, let $\mathbf{r} = \text{pg}(\mathbf{A}, \mathbf{e}, c)$ be the PageRank vector of the input graph \mathbf{A} , and $\mathbf{r}_{\mathcal{S}} = \text{pg}(\mathbf{A}_{\mathcal{S}}, \mathbf{e}, c)$ be the new PageRank vector after removing the graph elements

in set \mathcal{S} . We formulate the PageRank auditing problem as the following optimization problem.

$$\begin{aligned} \max_{\mathcal{S}} \quad & \Delta f = (f(\mathbf{r}) - f(\mathbf{r}_{\mathcal{S}}))^2 \\ \text{s.t.} \quad & |\mathcal{S}| = k \end{aligned} \quad (2)$$

In order to solve the above optimization problem, we need to answer two key questions including (Q1) how to quantitatively measure the influence of an individual graph element w.r.t. the objective function; and (Q2) how to collectively find a set of k graph elements with the maximal influence. We first present our proposed solution for Q1 in Subsection III-B, and then propose three different algorithms for Q2 in Subsections III-C, III-D and III-E, depending on the specific type of graph elements (i.e., edges vs. nodes vs. subgraphs).

B. Measuring Graph Element Influence

To measure how $f(\mathbf{r})$ will change if we perturb/remove a specific graph element, we define its influence as the rate of change in $f(\mathbf{r})$.

Definition 1: (Graph Element Influence). The influence of an edge (i, j) is defined as the derivative of the loss function $f(\mathbf{r})$ with respect to the edge, i.e., $\mathbb{I}(i, j) = \frac{df(\mathbf{r})}{d\mathbf{A}(i, j)}$. The influence of a node is defined as the aggregation of all inbound and outbound edges that connect to the node., i.e., $\mathbb{I}(i) = \sum_{j=1, j \neq i}^n [\mathbb{I}(i, j) + \mathbb{I}(j, i)] + \sum_{j=1, j=i}^n \mathbb{I}(i, j)$. And the influence of a subgraph is defined as the aggregation of all edges in the subgraph S , $\mathbb{I}(S) = \sum_{i, j \in S} \mathbb{I}(i, j)$.

We can see that the influence for both nodes and subgraphs can be naturally computed based on the edge influence. Therefore, we will focus on how to measure the edge influence. By the property of the derivative of matrices, we first rewrite the influence $\frac{df(\mathbf{r})}{d\mathbf{A}(i, j)}$ as

$$\frac{df(\mathbf{r})}{d\mathbf{A}} = \begin{cases} \frac{\partial f(\mathbf{r})}{\partial \mathbf{A}} + \left(\frac{\partial f(\mathbf{r})}{\partial \mathbf{A}}\right)' - \text{diag}\left(\frac{\partial f(\mathbf{r})}{\partial \mathbf{A}}\right), & \text{if undirected} \\ \frac{\partial f(\mathbf{r})}{\partial \mathbf{A}}, & \text{if directed} \end{cases} \quad (3)$$

Directly calculating $\frac{\partial f(\mathbf{r})}{\partial \mathbf{A}(i, j)}$ is hard, and we resort to the chain rule:

$$\frac{\partial f(\mathbf{r})}{\partial \mathbf{A}(i, j)} = \frac{\partial f(\mathbf{r})}{\partial \mathbf{r}} \frac{\partial \mathbf{r}}{\partial \mathbf{A}(i, j)} \quad (4)$$

Next, we present the details on how to solve each partial derivative in Eq. (4) one by one.

a) *Computing $\frac{\partial f(\mathbf{r})}{\partial \mathbf{r}}$:* Here we discuss the choices of $f(\cdot)$ function. In this paper, we choose $f(\cdot)$ to be squared L_2 norm for simplicity. However, it is worth mentioning that the proposed methods are applicable to a variety of other loss functions as well. We list some alternative choices and their corresponding derivatives in Table II. In the table, *L_p norm* is the most commonly-used vector norm that measures the overall sizes of the vector. Some commonly-used L_p norm includes L_1 norm and L_2 norm (also known as the Euclidean norm); *soft maximum* is used to approximate the maximum value of elements in the vector; *energy norm* is a measurement often used in system and control theory to measure the internal energy of vector.

TABLE II: Choices of $f(\cdot)$ functions and their derivatives

Descriptions	Functions	Derivatives
L_p norm	$f(\mathbf{r}) = \ \mathbf{r}\ _p$	$\frac{\partial f}{\partial \mathbf{r}} = \frac{\mathbf{r} \cdot \mathbf{r} ^{p-2}}{\ \mathbf{r}\ _p^{p-1}}$
Soft maximum	$f(\mathbf{r}) = \log\left(\sum_{i=1}^n \exp(\mathbf{r}(i))\right)$	$\frac{\partial f}{\partial \mathbf{r}} = \left[\frac{\exp(\mathbf{r}(i))}{\sum_{i=1}^n \exp(\mathbf{r}(i))}\right]$
Energy norm	$f(\mathbf{r}) = \mathbf{r}' \mathbf{M} \mathbf{r}$	$\frac{\partial f}{\partial \mathbf{r}} = (\mathbf{M} + \mathbf{M}') \mathbf{r}$

(\mathbf{M} in Energy Norm is a Hermitian positive definite matrix.)

b) 2. *Computing $\frac{\partial \mathbf{r}}{\partial \mathbf{A}(i, j)}$:* Recall that \mathbf{A} is the adjacency matrix of the input graph, and the PageRank solution can be written as solving the linear system in Eq. (1). Then we have that

$$\frac{\partial \mathbf{r}}{\partial \mathbf{A}(i, j)} = c \frac{\partial \mathbf{A}}{\partial \mathbf{A}(i, j)} \mathbf{r} + c \mathbf{A} \frac{\partial \mathbf{r}}{\partial \mathbf{A}(i, j)} \quad (5)$$

Move the second term in Eq. (5) to the left, we have that

$$(\mathbf{I} - c\mathbf{A}) \frac{\partial \mathbf{r}}{\partial \mathbf{A}(i, j)} = c \frac{\partial \mathbf{A}}{\partial \mathbf{A}(i, j)} \mathbf{r} \quad (6)$$

$$\frac{\partial \mathbf{r}}{\partial \mathbf{A}(i, j)} = c(\mathbf{I} - c\mathbf{A})^{-1} \frac{\partial \mathbf{A}}{\partial \mathbf{A}(i, j)} \mathbf{r} \quad (7)$$

By the property of the first order derivative of matrix, we have that

$$\frac{\partial \mathbf{A}}{\partial \mathbf{A}(i, j)} = \begin{cases} \mathbf{S}^{ij} + \mathbf{S}^{ji}, & \text{if undirected} \\ \mathbf{S}^{ij}, & \text{if directed} \end{cases} \quad (8)$$

where \mathbf{S}^{ij} is the single-entry matrix with 1 at i^{th} row and j^{th} column and 0 elsewhere. Recall that $\mathbf{Q} = (\mathbf{I} - c\mathbf{A})^{-1}$, we can re-write Eq. (7) as the following equation.

$$\frac{\partial \mathbf{r}}{\partial \mathbf{A}(i, j)} = c \mathbf{r}(j) \mathbf{Q}(:, i) \quad (9)$$

Combine everything together, we get the closed-form solution for calculating the influence of an edge (i, j) as follows:

$$\frac{\partial f(\mathbf{r})}{\partial \mathbf{A}(i, j)} = 2c \mathbf{r}(j) \text{Tr}(\mathbf{r}' \mathbf{Q}(:, i)) \quad (10)$$

Following Eq. (10), we get the matrix of gradients for all edges as

$$\frac{\partial f(\mathbf{r})}{\partial \mathbf{A}} = 2c(1 - c) \mathbf{Q}' \mathbf{r} \mathbf{e}' \mathbf{Q}' \quad (11)$$

Two major computational challenges in calculating $\frac{\partial f(\mathbf{r})}{\partial \mathbf{A}}$ lie in (1) calculating \mathbf{Q} with $O(n^3)$ time complexity and (2) $O(n^2)$ space complexity to save the matrix of gradients. We address both challenges by exploring the low-rank structure of $\frac{\partial f(\mathbf{r})}{\partial \mathbf{A}}$. From Eq. (11), we can show that it can be re-written as the following low-rank form

$$\frac{\partial f(\mathbf{r})}{\partial \mathbf{A}} = 2c(1 - c) \mathbf{Q}' \mathbf{r} \mathbf{e}' \mathbf{Q}' = 2c(\mathbf{Q}' \mathbf{r}) \mathbf{r}' \quad (12)$$

where $\mathbf{Q}' \mathbf{r}$ is an $n \times 1$ vector and \mathbf{r} is an $n \times 1$ PageRank vector. Since $\mathbf{r} = (1 - c)\mathbf{Q}\mathbf{e}$, we have that $\mathbf{Q}' \mathbf{r}$ is a personalized PageRank on the reverse of the input graph with \mathbf{r} as teleportation vector with a constant scaling. With this in mind, we do not need to calculate \mathbf{Q} explicitly, or to save the entire matrix directly. Instead, we can use power method to calculate

\mathbf{r} and $\mathbf{Q}'\mathbf{r}$, each with $O(m)$ time and save these two vectors with $O(n)$ space. To extract the element of $\frac{\partial f(\mathbf{r})}{\partial \mathbf{A}}$ at the i^{th} row and the j^{th} column, we simply calculate the product of the i^{th} element in $\mathbf{Q}'\mathbf{r}$ and the j^{th} element in \mathbf{r} , and scale it by $2c$, which takes $O(1)$ time.

C. Auditing by Edges: AURORA-E

Due to its combinatorial nature, straight-forward methods for solving the optimization problem in Eq. (2) are not feasible. The key behind the proposed family of algorithms is due to the diminishing returns property of Problem 1, which is summarized in Theorem 1.

Theorem 1: (Diminishing Returns Property of Problem 1). For any loss function listed in Table 2, and for any set of graph elements \mathcal{S} , which could be either a set of edges, nodes or subgraphs, in the given graph, its influence measure $\mathbb{I}(\mathcal{S})$ defined in Definition 1 is (a) normalized; (b) monotonically non-decreasing; (c) submodular, where $\mathcal{S} \subseteq \mathcal{E}$.

Proof: We only prove the diminishing returns property in the edge case. The proofs for nodes and subgraphs are similar and thus is omitted due to the space limitation.

Let $\mathbb{I}(\mathcal{S}) = \sum_{(i,j) \in \mathcal{S}} \mathbb{I}(i,j)$. It is trivial that if there is no edge selected, the influence is 0. Thus it is normalized.

Let $\mathcal{I}, \mathcal{J}, \mathcal{K}$ be three sets and $\mathcal{I} \subseteq \mathcal{J}$. We further define three sets $\mathcal{S}, \mathcal{T}, \mathcal{K}$ as follows: $\mathcal{S} = \mathcal{I} \cup \mathcal{K}$, $\mathcal{T} = \mathcal{J} \cup \mathcal{K}$ and $\mathcal{R} = \mathcal{J} \setminus \mathcal{I}$, then we have

$$\begin{aligned} \mathbb{I}(\mathcal{J}) - \mathbb{I}(\mathcal{I}) &= \sum_{(i,j) \in \mathcal{J}} \mathbb{I}(i,j) - \sum_{(i,j) \in \mathcal{I}} \mathbb{I}(i,j) \\ &= \sum_{(i,j) \in \mathcal{J} \setminus \mathcal{I}} \mathbb{I}(i,j) \\ &= \sum_{(i,j) \in \mathcal{R}} \mathbb{I}(i,j) \\ &\geq 0 \end{aligned}$$

which proves that $\mathbb{I}(\mathcal{S})$ is monotonically non-decreasing.

Finally, we prove that it is submodular. Define $\mathcal{P} = \mathcal{T} \setminus \mathcal{S}$. We have that $\mathcal{P} = (\mathcal{J} \cup \mathcal{K}) \setminus (\mathcal{I} \cup \mathcal{K}) = \mathcal{R} \setminus (\mathcal{R} \cap \mathcal{K}) \subseteq \mathcal{R} = \mathcal{J} \setminus \mathcal{I}$. Then we have

$$\mathbb{I}(\mathcal{T}) - \mathbb{I}(\mathcal{S}) = \sum_{(i,j) \in \mathcal{P}} \mathbb{I}(i,j) \leq \mathbb{I}(\mathcal{J}) - \mathbb{I}(\mathcal{I})$$

which proves the submodularity of the edge influence. \blacksquare

The diminishing returns property naturally leads to a greedy algorithm to obtain a near-optimal solution for solving Problem 1. We first present the algorithm for auditing by edges in this subsection. The algorithms for auditing by nodes and by subgraphs will be presented in Subsections III-D and III-E, respectively.

With the diminishing returns property, we propose AURORA-E (Algorithm 1) algorithm to find top- k influential edges. The key idea of AURORA-E is to select one edge and update the gradient matrix at each of the k iterations.

The effectiveness and efficiency of the proposed AURORA-E are summarized in Lemma 1 and Lemma 2, respectively.

Algorithm 1: AURORA-E

```

Input : The adjacency matrix  $\mathbf{A}$ , integer budget  $k$ 
Output: A set of  $k$  edges  $\mathcal{S}$  with the highest influence
1 initialize  $\mathcal{S} = \emptyset$ ;
2 initialize  $c$  (e.g.,  $c = 1/2 \max \text{eigenvalue}(\mathbf{A})$ );
3 calculate PageRank  $\mathbf{r} = \text{pg}(\mathbf{A}, \mathbf{e}, c)$ ;
4 calculate partial gradients  $\frac{\partial f(\mathbf{r})}{\partial \mathbf{A}}$  by Eq. (12);
5 calculate gradients  $\frac{df(\mathbf{r})}{d\mathbf{A}}$  by Eq. (3);
6 while  $|\mathcal{S}| \neq k$  do
7   find  $(i, j) = \underset{(i,j)}{\text{argmax}} \mathbb{I}(i, j)$  with Eq. (3);
8   add edge  $(i, j)$  to  $\mathcal{S}$ ;
9   remove  $(i, j)$ , and remove  $(j, i)$  if undirected;
10  re-calculate  $\mathbf{r}$ ,  $\frac{\partial f(\mathbf{r})}{\partial \mathbf{A}}$  by Eq. (12), and  $\frac{df(\mathbf{r})}{d\mathbf{A}}$  by
    Eq. (3);
11 return  $\mathcal{S}$ ;

```

We can see that AURORA-E finds a $(1 - 1/e)$ near-optimal solution with a linear complexity.

Lemma 1: (Approximation Ratio of AURORA-E). Let $\mathcal{S}_k = \{s_1, s_2, \dots, s_k\}$ represents the set formed by AURORA-E, \mathcal{O} is the optimal solution of Problem 1, $\mathbb{I}(\mathcal{S})$ is the influence defined in Definition 1.

$$\mathbb{I}(\mathcal{S}_k) \geq (1 - 1/e)\mathbb{I}(\mathcal{O})$$

Proof: Omitted for space. \blacksquare

Lemma 2: (Time and Space Complexities of AURORA-E). Algorithm 1 is $O(mk)$ in time and $O(m+n)$ in space, where m and n are the numbers of edges and nodes in the input graph; and k is the budget.

Proof: It takes $O(m)$ time complexity to calculate \mathbf{r} and $\frac{\partial f(\mathbf{r})}{\partial \mathbf{A}}$ by applying power iterations. In the while-loop, we find the edge with the greatest influence by traversing all edges, which takes $O(m)$ time. Time spent to re-calculate \mathbf{r} and $\frac{\partial f(\mathbf{r})}{\partial \mathbf{A}}$ remains the same as $O(m)$. Since the body inside the loop will run k times, the overall time complexity is $O(mk)$. In Algorithm 1, it takes $O(m)$ space to save the sparse adjacency matrix \mathbf{A} and $O(n)$ space to save the PageRank vector \mathbf{r} and column vector $\mathbf{Q}'\mathbf{r}$ in Eq. (12). Therefore it has $O(m+n)$ space complexity. \blacksquare

D. Auditing by Nodes: AURORA-N

By Theorem 1, the influence of nodes also enjoys the diminishing returns property. Following this, we propose a greedy algorithm AURORA-N (Algorithm 2) to find a set of k influential nodes with $(1 - 1/e)$ approximation ratio with a linear complexity. The efficiency of the proposed AURORA-N is summarized in Lemma 3.

Lemma 3: (Time and Space Complexities of AURORA-N). Algorithm 2 is $O(mk)$ in time and $O(m+n)$ in space, where m and n are the numbers of edges and nodes in the input graph; and k is the budget.

Proof: Omitted for space. \blacksquare

Algorithm 2: AURORA-N

Input : The adjacency matrix \mathbf{A} , integer budget k
Output: A set of k nodes \mathcal{S} with highest influence

- 1 initialize $\mathcal{S} = \emptyset$;
- 2 initialize c (e.g., $c = 1/2 \max \text{eigenvalue}(\mathbf{A})$);
- 3 calculate PageRank $\mathbf{r} = \text{pg}(\mathbf{A}, \mathbf{e}, c)$;
- 4 calculate partial gradients $\frac{\partial f(\mathbf{r})}{\partial \mathbf{A}}$ by Eq. (12);
- 5 calculate gradients $\frac{df(\mathbf{r})}{d\mathbf{A}}$ by Eq. (3);
- 6 **while** $|\mathcal{S}| \neq k$ **do**
- 7 find $v_i = \underset{i}{\operatorname{argmax}} \mathbb{I}(i)$;
- 8 add v_i to \mathcal{S} ;
- 9 remove all inbound and outbound edges of v_i ;
- 10 re-calculate \mathbf{r} , $\frac{\partial f(\mathbf{r})}{\partial \mathbf{A}}$ by Eq. (12), and $\frac{df(\mathbf{r})}{d\mathbf{A}}$ by Eq. (3);
- 11 **return** \mathcal{S} ;

E. Auditing by Subgraphs: AURORA-S

Here, we discuss how to select an influential subgraph with k nodes and we focus on the vertex-induced subgraph. With the diminishing returns property (Theorem 1) in mind, we propose AURORA-S (Algorithm 3) to greedily identify the influential subgraph with $(1 - 1/e)$ approximation ratio with a linear complexity. The efficiency of the proposed AURORA-S is summarized in Lemma 4.

Algorithm 3: AURORA-S

Input : The adjacency matrix \mathbf{A} , output size k
Output: A vertex-induced subgraph of k nodes \mathcal{S} with highest influence

- 1 initialize $\mathcal{S} = \emptyset$;
- 2 initialize c (e.g., $c = 1/2 \max \text{eigenvalue}(\mathbf{A})$);
- 3 calculate PageRank $\mathbf{r} = \text{pg}(\mathbf{A}, \mathbf{e}, c)$;
- 4 calculate partial gradients $\frac{\partial f(\mathbf{r})}{\partial \mathbf{A}}$ by Eq. (12);
- 5 calculate gradients $\frac{df(\mathbf{r})}{d\mathbf{A}}$ by Eq. (3);
- 6 **while** $|\mathcal{S}| \neq k$ **do**
- 7 find $(i, j) = \underset{(i, j)}{\operatorname{argmax}} \mathbb{I}(i, j)$;
- 8 **if** $|\mathcal{S}| + 2 \leq k$ **then**
- 9 add v_i and v_j to \mathcal{S} ;
- 10 **else**
- 11 find the endpoint v with higher gradient;
- 12 **if** $v \notin \mathcal{S}$ **then**
- 13 add v to \mathcal{S} ;
- 14 **else**
- 15 add the other endpoint to \mathcal{S} ;
- 16 remove all edges in \mathcal{S} ;
- 17 re-calculate \mathbf{r} , $\frac{\partial f(\mathbf{r})}{\partial \mathbf{A}}$ by Eq. (12), and $\frac{df(\mathbf{r})}{d\mathbf{A}}$ by Eq. (3);
- 18 **return** \mathcal{S} ;

Lemma 4: (Time and Space Complexities of AURORA-S). Algorithm 3 is $O(mk)$ in time and $O(m + n)$ in space, where m and n are the numbers of edges and nodes in the input graph; and k is the budget. ■

Proof: Omitted for space. ■

F. Generalization and Variants

The proposed family of AURORA algorithms assume the input graph is a plain network. However, it is worth pointing out that AURORA algorithms also work on different types of networks and other random walk-based techniques.

a) AURORA on Normalized PageRank: Recall that in Section II, we remove the constraint of \mathbf{A} being a normalized adjacency matrix and use the general form instead. The effect of removing that constraint will cause the L_1 norm of PageRank to be not equal to 1. We show that AURORA is also able to work on L_1 normalized PageRank. Let $S(\mathbf{r}) = \sum_i \mathbf{r}(i)$, we have

$$\frac{\partial f(\mathbf{r})}{\partial \mathbf{A}} = c \mathbf{Q}' \left(-\frac{2f(\mathbf{r})}{S(\mathbf{r})} \mathbf{1} + \frac{2}{S(\mathbf{r})} \mathbf{r}' \right) \quad (13)$$

Then we can apply AURORA algorithms by replacing Eq. (12) with Eq. (13)

b) AURORA on Network of Networks: Network of Networks (NoN) is a type of networks first introduced in [12] with the ability to leverage the within-network smoothness in the domain-specific network and the cross-network consistency through the main network. An NoN is usually defined as the triplet $\mathcal{R} = \langle \mathbf{G}, \mathcal{A}, \theta \rangle$, where \mathbf{G} is the main network, \mathcal{A} is a set of domain-specific network and θ is a mapping function to map the main node to the corresponding domain-specific network. In [12], CROSSRANK and CROSSQUERY are two ranking algorithms proposed to solve ranking on NoN. The authors have proved that they are actually equivalent to the well-known PageRank and random walk with restart on the integrated graph. Thus, AURORA algorithms also have the ability to audit CROSSRANK and CROSSQUERY on Network of Networks.

c) AURORA on Attributed Networks: Given a large attributed network, it is important to learn the most influential node-attribute or edge-attribute w.r.t. a query node. We show that AURORA algorithms have the ability to find top- k influential edge-attributes and node-attributes on attributed networks. The central idea is to treat attributes as *attribute nodes* and form an augmented graph with those attribute nodes. To support node attributes, let \mathbf{A} be the $a \times a$ node-to-node adjacency matrix, and \mathbf{W} be the $w \times a$ node-to-attribute matrix, then we can form an augmented graph $\mathbf{G} = \begin{pmatrix} \mathbf{A} & \mathbf{0} \\ \mathbf{W} & \mathbf{0} \end{pmatrix}$. To support edge attributes, similar to [13], we embed an edge-node for each edge in the input graph and define a mapping function ψ that maps each edge-node to edge attribute in the original graph. We assume \mathbf{A} is the $a \times a$ node-to-node adjacency matrix, and b is the number of different edge attribute values. By embedding edge-node, it creates a $(a + b) \times (a + b)$ augmented graph \mathbf{Y} . Then to find the top- k influential node-attributes and edge-attributes, we can easily

TABLE III: Statistics of the datasets

Category	Network	Type	Nodes	Edges
SOCIAL	Karate	U	34	78
	Dolphins	U	62	159
	WikiVote	D	7,115	103,689
	Pokec	D	1,632,803	30,622,564
COLLABORATION	GrQc	U	5,242	14,496
	DBLP	U	42,252	420,640
	NBA	U	3,923	127,034
	cit-DBLP	D	12,591	49,743
	cit-HepTh	D	27,770	352,807
PHYSICAL	cit-HepPh	D	34,546	421,578
	Airport	D	1,128	18,736
OTHERS	Lesmis	U	77	254
	Amazon	D	262,111	1,234,877

(In Type, U means undirected graph; D means directed graph.)

run AURORA-N on the augmented attributed graphs \mathbf{G} and \mathbf{Y} , respectively.

IV. EXPERIMENTAL EVALUATION

In this section, we evaluate the proposed AURORA algorithms. All experiments are designed to answer the following two questions:

- **Effectiveness.** How effective are the proposed AURORA algorithms in identifying key graph elements w.r.t. the PageRank results?
- **Efficiency.** How efficient and scalable are the proposed AURORA algorithms?

A. Setup

Datasets. We test our algorithms on a diverse set of real-world network datasets. All datasets are publicly available. The statistics of these datasets are listed in Table III.

- **SOCIAL NETWORKS.** Here, nodes are users and edges indicate social relationships. Among them, *Karate* [14] is a well-known network dataset of a university karate club collected by Wayne Zachary in 1977. *Dolphins* [15] is an undirected social network of frequent associations between dolphins in a community living off Doubtful Sound, New Zealand. *WikiVote* [16] is generated by Wikipedia voting data from the inception of Wikipedia till January 2008. *Pokec* [17] is a popular online social network in Slovakia.
- **COLLABORATION NETWORKS.** Here, nodes are individuals and two people are connected if they have collaborated together. We use the collaboration network in the field of General Relativity and Quantum Cosmology (*GrQc*) in Physics from arXiv preprint archive². *DBLP*³ is a co-authorship network from DBLP computer science bibliography. And *NBA* [18] is a collaboration network of NBA players from 1946 to 2009. *cit-DBLP* [19] is the citation network of DBLP, a database of scientific publications such as papers and books. Each node in the network is a publication, and each edge represents

a citation of a publication by another publication. *cit-HepTh* [20] is an ArXiv HEP-TH (High Energy Physics - Theory) citation network. The data covers papers from January 1993 to April 2003. If a paper i cites paper j , there is a directed edge from i to j . *cit-HepPh* [20] is an ArXiv HEP-PH (High Energy Physics - Phenomenology) citation network. The data covers papers from January 1993 to April 2003. If a paper i cites paper j , there is a directed edge from i to j .

- **PHYSICAL INFRASTRUCTURE NETWORKS.** This category refers to the networks of physical infrastructure entities. Nodes in them correspond to physical infrastructure, and edges are connections. *Airport*⁴ is a dataset of airline traffic. Each node represents an airport in the United States, an edge (i, j) represents the airline from i to j while the edge weight stands for the normalized number of passengers.
- **OTHERS.** This category contains networks that do not fit into the above categories. *Lesmis* [21] is a network of co-appearances of characters in Victor Hugo's novel "*Les Misérables*". A node represents a character and an edge connects a pair of characters if they both appear in the same chapter of the book. *Amazon* [22] is a co-purchasing network collected by crawling Amazon website. It is based on the *Customers Who Bought This Item Also Bought* feature.

Baseline Methods. We compare our proposed methods with several baseline methods, which are summarized as follows.

- *Brute force (brute force).* Calculate the changes by iterating all possible combinations of graph elements.
- *Random Selection (random).* Randomly select k elements and calculate the change by removing them.
- *Top- k Degrees (degree).* We first define the degree of an edge (u, v) as follows,

$$d(u, v) = \begin{cases} (d(u) + d(v)) \times \max_{i \in \{u, v\}} d(i), & \text{if undirected} \\ (d(u) + d(v)) \times d(u), & \text{if directed} \end{cases}$$

where $d(u)$ represents the degree of node u .

To audit by graph elements, we select k elements with the highest degrees. For edges, we select k edges with the highest edge degrees defined above; for nodes, we select k nodes with the highest node degrees; for subgraphs, we form a vertex-induced subgraph from k nodes with the highest degrees.

- *PageRank.* We first define the PageRank score of an edge (u, v) as follows,

$$\mathbf{r}(u, v) = \begin{cases} (\mathbf{r}(u) + \mathbf{r}(v)) \times \max_{i \in \{u, v\}} \mathbf{r}(i), & \text{if undirected} \\ (\mathbf{r}(u) + \mathbf{r}(v)) \times \mathbf{r}(u), & \text{if directed} \end{cases}$$

where $r(u)$ is the PageRank score of node u .

To audit by graph elements, we select k elements with the highest PageRank scores. That is, for edges, we select

²<https://arxiv.org/>

³<http://dblp.uni-trier.de/>

⁴<https://www.transtats.bts.gov/>

k edges with the highest PageRank scores defined above; for nodes, we select k nodes with highest PageRank scores; for subgraphs, we form a vertex-induced subgraph from k nodes with the highest PageRank scores.

- **HITS.** We first define HITS score of an edge (u, v) and node u as follows,

$$HITS(u, v) = hub(u) \times hub(v) + auth(u) \times auth(v)$$

$$HITS(u) = hub(u) + auth(u)$$

where $hub(u)$ and $auth(u)$ represent the hub score and authority score of node u , respectively.

To audit by graph elements, we select k elements with the highest HITS scores. That is, for edges, we select k edges with the highest HITS scores defined above; for nodes, we select k nodes with the highest HITS scores; for subgraphs, we form a vertex-induced subgraph from k nodes with the highest HITS scores.

Metrics. Here, we choose the loss function to be squared L_2 norm. We quantify the performance of auditing by the goodness score Δf (defined in Eq. (2)) of the graph elements \mathcal{S} found by the corresponding algorithms.

Repeatability and Machine Configuration. All datasets are publicly available. We will release the code of our proposed algorithms upon the publication of the paper. All experiments are performed on a virtual machine with 4 Intel i7-8700 CPU cores at 3.4GHz and 32GB RAM. The operating system is Windows 10. All codes are written in Python 3.6.

B. Effectiveness Results

a) **Quantitative Comparison:** We perform effectiveness experiments with the baseline methods. We set k from 1 to 10 and find k influential edges and nodes, respectively. We set k only from 2 to 10 to find an influential subgraph of size- k , respectively. This is because a vertex-induced subgraph with only 1 node does not contain any edge and therefore is meaningless for the purpose of PageRank auditing. We conduct the experiments with *brute force* on relatively small datasets to obtain the ground-truth of Eq. (2), including *Karate*, *Dolphins* and *Lesmis*. On other larger datasets, searching a ground-truth with k most influential elements is prohibitively expensive due to its combinatorial nature. For example, even if we use the small *Lesmis* dataset, it will take over a day to find ground-truth with $k = 5$. Therefore, the results by *brute force* are absent on the remaining larger datasets.

The results of quantitative comparison across 9 different datasets are shown from Figure 2 to Figure 4. From those figures, we have the following observations: (1) our family of AURORA algorithms consistently outperform other baseline methods (except for brute force) on all datasets; (2) on the three small datasets with ground-truth *brute force*, the performance of the proposed AURORA algorithms is very close to the ground-truth. Figure 5 and Figure 6 shows the effect of k on auditing results. We can observe the following findings from those figures: (1) our family of AURORA algorithms incrementally find influential graph elements w.r.t. the budget k ; (2) the proposed AURORA algorithms consistently outperform baseline methods on different budgets.

b) **Case Studies on Airport Dataset:** A natural use case of our AURORA algorithms is to find influential edges and nodes in a given graph. To demonstrate that our algorithms are indeed able to provide intuitive information, we test our algorithms on the *Airport* dataset. This dataset was manually created from commercial airline traffic data in 2017, which is provided by United States Department of Transportation. More detailed description and statistics of this dataset can be found in Subsection IV-A. We perform AURORA-E and AURORA-N to find the most influential airlines (edges in the graph) and airports (nodes in the graph) across United States with $k = 7$.

Edges selected by AURORA-E are ATL-LAX, LAX-ATL, ATL-ORD, ORD-ATL, ATL-DEN, DEN-ATL and LAX-ORD. In contrast, PageRank selects ATL-LAS instead of DEN-ATL and ATL-DFW instead of LAX-ORD. DEN-ATL plays a more important role in determining the centrality (e.g., PageRank) of other airports. This is because DEN serves as one of the busiest hub airports that connects West coast and East coast; while ATL-LAS is less important in that regard, considering the existence of ATL-LAX and ATL-PHX. Comparing LAX-ORD and ATL-DFW, LAX-ORD directly connects Los Angeles and Chicago, both of them are largest cities in the United States.

In the scenario of node-auditing, AURORA-N selects ATL, LAX, ORD, DFW, DEN, LAS and CLT. In contrast, PageRank selects SFO instead of CLT. CLT seems to be a more reasonable choice because it serves as a major hub airport, the 6th busiest airport by FAA statistics, to connect many regional airports around States like North Carolina, South Carolina, Virginia, West Virginia, etc. Compared with CLT, SFO is less influential in that regard, mainly due to the following two reasons: (1) it ranks after CLT (7th vs. 6th) in the list of busiest airports by FAA statistics; (2) due to the location proximity of SFO to LAX and SJC, even if this node is perturbed (i.e., absent), many surrounding airports (especially regional airports in California) could still be connected via LAX and SJC.

c) **Case Studies on DBLP Dataset:** Another interesting use case of AURORA algorithms is sense-making in graph proximity. We construct a co-authorship network from DBLP computer science bibliography to test our algorithms. We perform AURORA-N and PageRank with $k = 6$. Different from the previous case study, here we use a personalized PageRank with the query node is Christos Faloutsos. In this case, the top-ranked scholars in the resulting ranking vector \mathbf{r} form the proximity (i.e., ‘neighborhood’) of the query node (i.e., who are most relevant to Christos Faloutsos). Consequently, the nodes selected by an auditing algorithm indicate those important nodes in terms of making/maintaining the neighborhood of the query node. Comparing the results by AURORA-N and PageRank, 6 of them are the same while AURORA-N selects Jure Leskovec instead of Yannis Ioannidis. This result is consistent with the intuition, since Jure Leskovec, as the former student of Christos Faloutsos with lots of joint publications, plays a more prominent role in the neighborhood of Christos Faloutsos by sharing more common collabora-

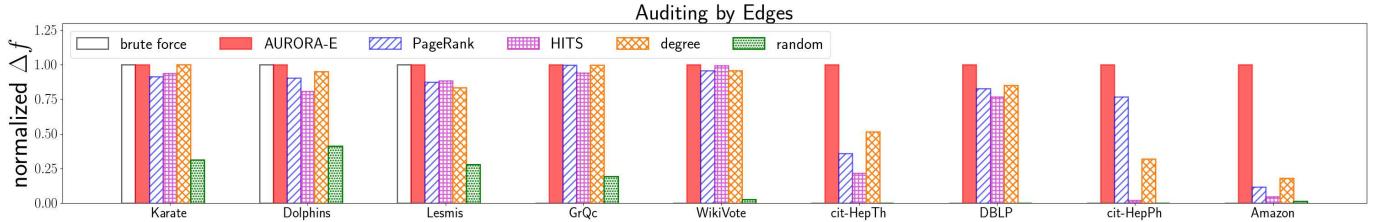


Fig. 2: Auditing results by edges. Budget $k = 10$. Higher is better. Best viewed in color.

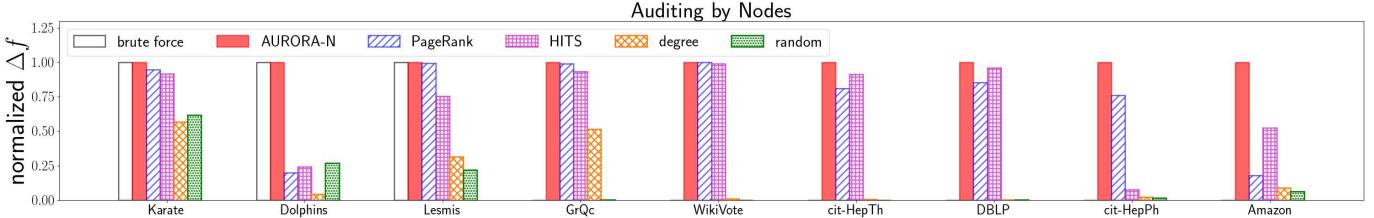


Fig. 3: Auditing results by nodes. Budget $k = 10$. Higher is better. Best viewed in color.

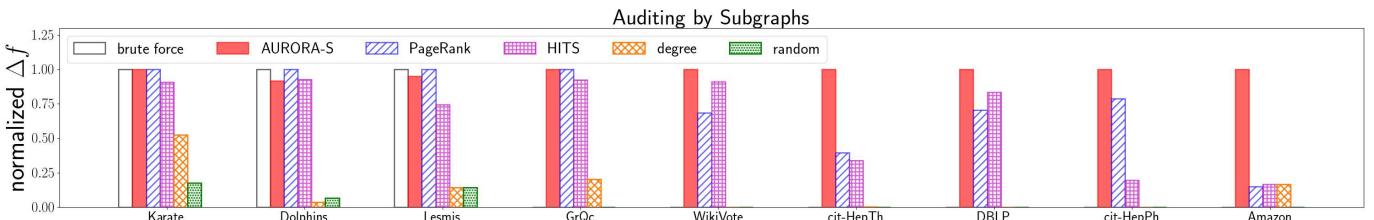


Fig. 4: Auditing results by subgraphs. Budget $k = 10$. Higher is better. Best viewed in color.

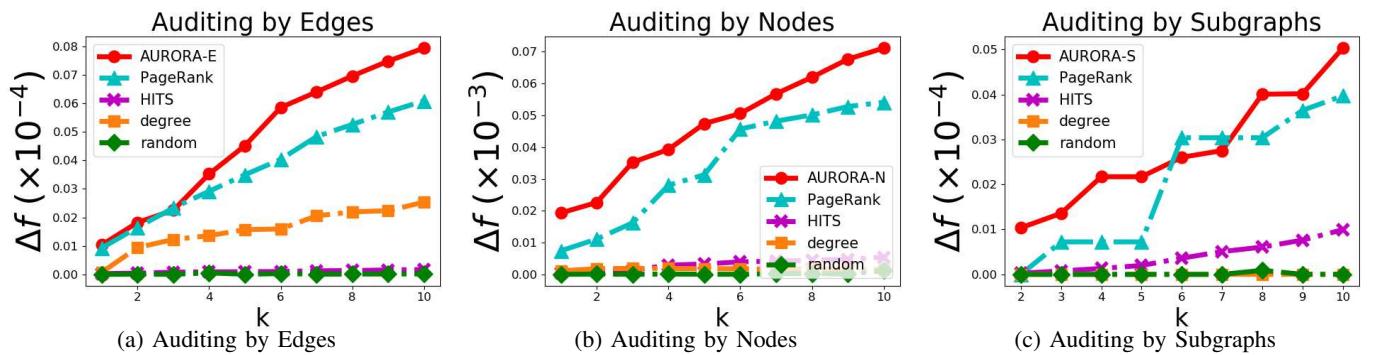


Fig. 5: Effect of k on auditing results (cit-HepPh Dataset). Higher is better. Best viewed in color.

rators.

d) Case Studies on NBA Dataset: In a collaboration network, a subgraph can be naturally viewed as a team (e.g. sports team). From this perspective, AURORA-S has the potential to find teammates of a player. We set the query node as Allen Iverson. Since there are 5 players for each team on the court, we set $k = 5$. The subgraph selected by AURORA-S consists of Allen Iverson, Larry Hughes, Theo Ratliff, Joe Smith and Tim Thomas. In contrast, in the resulting subgraph by PageRank, Tim Thomas is replaced by Drew Gooden. Though Drew Gooden and Allen Iverson shares many common neighbors in the network, which might cause its ranking in PageRank to be higher, they have never played in the same team together. Compared with the subgraph by PageRank, all players in the subgraph by AURORA-S have played with Iverson in Philadelphia Sixers during the time period from 1997 to 1999.

C. Efficiency Results

We show the running time vs. number of edges m and budget size k on *Pokec* dataset in Figure 7 and Figure 8. We can see that the proposed AURORA algorithms scale linearly with respect to m and k , respectively. This is consistent with our complexity analysis that the family of AURORA algorithms are linear with respect to the number of edges and the budget.

D. Visualization

To better understand the auditing results, we developed a prototype system with D3.js to represent the influence of graph elements visually. In the system, we use the strength of line to represent the gradient of an edge and use the size and color for the gradient of nodes. An example of visualizing hand-crafted toy graph is shown in Figure 9. It is easy to see in the figure that Node 5 is the most influential node, and edges around Node 5 is more influential than other edges, both of which are consistent with our intuition.

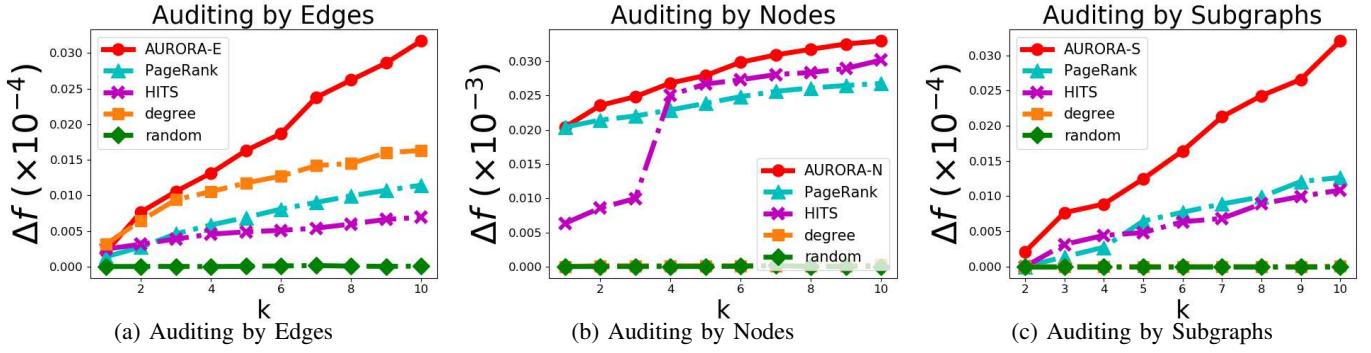


Fig. 6: Effect of k on auditing results (cit-HepTh Dataset). Higher is better. Best viewed in color.

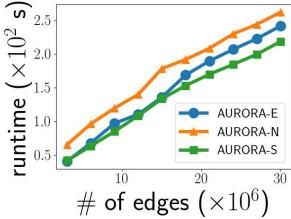


Fig. 7: running time vs. the number of edges on Pokec dataset

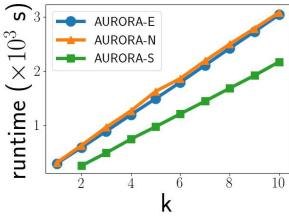


Fig. 8: running time vs. number of k on Pokec dataset

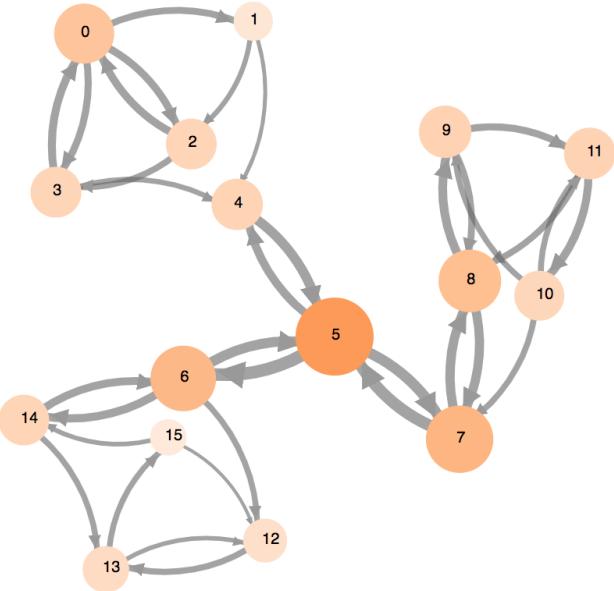


Fig. 9: Visualization of toy graph on the visualization system. Best viewed in color.

V. RELATED WORK

In this section, we briefly review the related work, from the following two perspectives, including (1) graph ranking and (2) explainability.

Graph ranking. Regarding graph ranking, PageRank [1] and HITS [23] are probably the most well-known and widely used algorithms. PageRank measures the importance of nodes as a stationary distribution of random walks. HITS assumes that each node has two scores: *hub* and *authority*. A node has a high hub score if it links to many nodes with high authority

scores, and a node has a high authority score if it links to many nodes with high hub scores. Many variants of PageRank and HITS have been developed in the literature. To name a few, in [24], the author studies the stability of PageRank and HITS, based on which they propose two new algorithms (Randomized HITS and Subspace HITS). Ding et al. [25] provide a unified ranking method for HITS and PageRank. In [8], Haveliwala et al. propose the well known *personalized PageRank* by replacing the uniform teleportation vector with a biased personalized topic-specific vector; while random walk with restart [9] concentrates all teleportation probabilities to a single node. Other random walk based graph ranking methods include [26] and many more.

Explainability. How to enhance interpretability of machine learning and data mining models has been attracting a lot of research interests in recent years. Two representative ways to explain such black-box predictions consist of (a) using other interpretable models to provide interpretable representation of model's predictions [27]–[29], and (b) quantifying the influence by perturbing either features or training data [11], [30]–[33]. The proposed AURORA algorithms follow the generic principle in [11] in the context of graph ranking, and has been integrated into a prototype system [34]. Finding influential nodes on graphs is a very active research area. The central theme of these research is to find nodes to maximize the spread of influence in social networks (i.e., influence maximization) [35]–[37]. The seminal work in influence maximization problem is attributed to Kempe et al. [38], [39], in which they have discovered its diminishing returns property. The focus of this paper is fundamentally different from the classic influence maximization problem, in the sense that we want to find most influential nodes, edges or subgraphs w.r.t. the graph ranking results, as opposed to the size of ‘infected’ nodes during the influence propagation process.

VI. CONCLUSION

In this paper, we study the problem of auditing PageRank, where we aim to find the most influential graph elements (e.g., edges, nodes, subgraphs) w.r.t. graph ranking results. We formally define the PageRank auditing problem by measuring influence of each graph element as the rate of change in a certain loss function defined over the ranking vector, and formulate it as an optimization problem. We further

propose a family of fast approximation algorithms, named AURORA, with $(1 - 1/e)$ approximation ratio and a linear complexity in both time and space. The extensive experimental evaluations on more than 10 datasets demonstrate that the proposed AURORA algorithms are able to identify influential graph elements, consistently outperform baseline methods, and scale linearly to large graphs. In the future, we would like to generalize this auditing paradigm to other graph ranking methods (e.g., HITS) as well as dynamic networks.

ACKNOWLEDGMENT

This work is supported by NSF (IIS-1651203, IIS-1715385 and IIS-1743040), DTRA (HDTRA1-16-0017), ARO (W911NF-16-1-0168), DHS (2017-ST-061-QA0001), NSFC (61602306, Fundamental Research Funds for the Central Universities), and gifts from Huawei and Baidu.

REFERENCES

- [1] L. Page, S. Brin, R. Motwani, and T. Winograd, “The pagerank citation ranking: Bringing order to the web.” Stanford InfoLab, Technical Report 1999-66, November 1999.
- [2] M. Gori and A. Pucci, “Itemrank: A random-walk based scoring algorithm for recommender engines,” in *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, San Francisco, CA, USA, 2007, pp. 2766–2771.
- [3] J. Weng, E.-P. Lim, J. Jiang, and Q. He, “Twitterrank: Finding topic-sensitive influential twitters,” in *Proceedings of the Third ACM International Conference on Web Search and Data Mining*. New York, NY, USA: ACM, 2010, pp. 261–270.
- [4] F. Radicchi, “Who is the best player ever? a complex network analysis of the history of professional tennis,” *PloS one*, vol. 6, no. 2, p. e17249, 2011.
- [5] R. Singh, J. Xu, and B. Berger, “Pairwise global alignment of protein interaction networks by matching neighborhood topology,” in *Research in computational molecular biology*. Springer, 2007, pp. 16–31.
- [6] J. J. Crofts and D. J. Higham, “Gooling the brain: Discovering hierarchical and asymmetric network structures, with applications in neuroscience,” *Internet Mathematics*, vol. 7, no. 4, pp. 233–254, 2011.
- [7] D. F. Gleich and M. W. Mahoney, “Using local spectral methods to robustify graph-based learning algorithms,” in *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, New York, NY, USA, 2015, pp. 359–368.
- [8] T. H. Haveliwala, “Topic-sensitive pagerank,” in *Proceedings of the 11th international conference on World Wide Web*. ACM, 2002, pp. 517–526.
- [9] H. Tong, C. Faloutsos, and J.-y. Pan, “Fast random walk with restart and its applications,” in *ICDM*. IEEE, 2006, pp. 613–622.
- [10] L. Li, Y. Yao, J. Tang, W. Fan, and H. Tong, “Quint: On query-specific optimal networks,” in *KDD*. ACM, 2016, pp. 985–994.
- [11] P. W. Koh and P. Liang, “Understanding black-box predictions via influence functions,” in *Proceedings of the 34th International Conference on Machine Learning*, International Convention Centre, Sydney, Australia, 2017, pp. 1885–1894.
- [12] J. Ni, H. Tong, W. Fan, and X. Zhang, “Inside the atoms: Ranking on a network of networks,” in *KDD*, New York, NY, USA, 2014, pp. 1356–1365.
- [13] R. Pienta, A. Tamersoy, H. Tong, and D. H. Chau, “Mage: Matching approximate patterns in richly-attributed graphs,” in *Big Data (Big Data), 2014 IEEE International Conference on*. IEEE, 2014, pp. 585–590.
- [14] W. W. Zachary, “An information flow model for conflict and fission in small groups,” *Journal of anthropological research*, vol. 33, no. 4, pp. 452–473, 1977.
- [15] D. Lusseau, K. Schneider, O. J. Boisseau, P. Haase, E. Slooten, and S. M. Dawson, “The bottlenose dolphin community of doubtful sound features a large proportion of long-lasting associations,” *Behavioral Ecology and Sociobiology*, vol. 54, no. 4, pp. 396–405, 2003.
- [16] J. Leskovec, D. Huttenlocher, and J. Kleinberg, “Signed networks in social media,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, New York, NY, USA, 2010, pp. 1361–1370.
- [17] L. Takac and M. Zabovsky, “Data analysis in public social networks.”
- [18] L. Li, H. Tong, N. Cao, K. Ehrlich, Y.-R. Lin, and N. Buchler, “Replacing the irreplaceable: Fast algorithms for team member recommendation,” in *WWW*, 2015, pp. 636–646.
- [19] M. Ley, “The dblp computer science bibliography: Evolution, research issues, perspectives,” in *Proceedings of the 9th International Symposium on String Processing and Information Retrieval*, London, UK, 2002, pp. 1–10.
- [20] J. Leskovec, J. Kleinberg, and C. Faloutsos, “Graphs over time: densification laws, shrinking diameters and possible explanations,” in *KDD*. ACM, 2005, pp. 177–187.
- [21] D. E. Knuth, *The Stanford GraphBase: a platform for combinatorial computing*. Addison-Wesley Reading, 1993, vol. 37.
- [22] J. Leskovec, J. Kleinberg, and C. Faloutsos, “Graph evolution: Densification and shrinking diameters,” *ACM Transactions on Knowledge Discovery from Data (TKDD)*, vol. 1, no. 1, p. 2, 2007.
- [23] J. M. Kleinberg, “Authoritative sources in a hyperlinked environment,” *Journal of the ACM (JACM)*, vol. 46, no. 5, pp. 604–632, 1999.
- [24] A. Y. Ng, A. X. Zheng, and M. I. Jordan, “Stable algorithms for link analysis,” in *Proceedings of the 24th annual international ACM SIGIR conference on Research and development in information retrieval*. ACM, 2001, pp. 258–266.
- [25] C. Ding, X. He, P. Husbands, H. Zha, and H. Simon, “PageRank, hits and a unified framework for link analysis,” in *Proceedings of the 2003 SIAM International Conference on Data Mining*. SIAM, 2003, pp. 249–253.
- [26] G. Jeh and J. Widom, “Simrank: a measure of structural-context similarity,” in *Proceedings of the 8th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2002, pp. 538–543.
- [27] R. Caruana, Y. Lou, J. Gehrke, P. Koch, M. Sturm, and N. Elhadad, “Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission,” in *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2015, pp. 1721–1730.
- [28] M. T. Ribeiro, S. Singh, and C. Guestrin, “Why should i trust you?: Explaining the predictions of any classifier,” in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2016, pp. 1135–1144.
- [29] O. Bastani, C. Kim, and H. Bastani, “Interpreting blackbox models via model extraction,” *arXiv preprint arXiv:1705.08504*, 2017.
- [30] M. Feldman, S. A. Friedler, J. Moeller, C. Scheidegger, and S. Venkatasubramanian, “Certifying and removing disparate impact,” in *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2015, pp. 259–268.
- [31] P. Adler, C. Falk, S. A. Friedler, G. Rybeck, C. Scheidegger, B. Smith, and S. Venkatasubramanian, “Auditing black-box models for indirect influence,” in *Data Mining (ICDM). 2016 IEEE 16th International Conference on*. IEEE, 2016, pp. 1–10.
- [32] A. Datta, S. Sen, and Y. Zick, “Algorithmic transparency via quantitative input influence: Theory and experiments with learning systems,” in *Security and Privacy (SP), 2016 IEEE Symposium on*. IEEE, 2016, pp. 598–617.
- [33] J. Li, W. Monroe, and D. Jurafsky, “Understanding neural networks through representation erasure,” *arXiv preprint arXiv:1612.08220*, 2016.
- [34] J. Kang, S. Freitas, H. Yu, Y. Xia, N. Cao, and H. Tong, “X-rank: Explainable ranking in complex multi-layered networks,” in *CIKM*. New York, NY, USA: ACM, 2018, pp. 1959–1962.
- [35] D. Gruhl, R. Guha, D. Liben-Nowell, and A. Tomkins, “Information diffusion through blogspace,” in *WWW*. ACM, 2004, pp. 491–501.
- [36] W. Chen, Y. Wang, and S. Yang, “Efficient influence maximization in social networks,” in *KDD*. ACM, 2009, pp. 199–208.
- [37] M. Gomez Rodriguez, J. Leskovec, and A. Krause, “Inferring networks of diffusion and influence,” in *KDD*. ACM, 2010, pp. 1019–1028.
- [38] D. Kempe, J. Kleinberg, and É. Tardos, “Maximizing the spread of influence through a social network,” in *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2003, pp. 137–146.
- [39] D. Kempe, J. M. Kleinberg, and É. Tardos, “Influential nodes in a diffusion model for social networks.” Springer.